Exhibit B-11

Invalidity Claim Chart for U.S. Patent No. 11,012,252 (the "'252 Patent")
in view of Obviousness References

As detailed in the cover pleading, and in addition to the reasons set forth in Exs. B-1 through B-10, the claims of the '252 Patent are invalid under 35 U.S.C. § 103 based on the prior art references identified in these invalidity contentions, alone or in combination with the obviousness references cited below. As explained in the cover pleading, a person of ordinary skill in the art as of the claimed priority date of the '252 Patent would have been motivated to combine the prior art identified in these invalidity contentions with the obviousness references cited below.

Defendants reserve the right to amend or supplement this claim chart at a later date as more fully set forth in the invalidity contentions.

Common Management Interface Specification for 8X/16X Pluggable Transceivers Rev 3.0 ("CMIS") was published September 18, 2018. CMIS qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

A 4-32 Gb/s Bidirectional Link With 3-Tap FFE/6-Tap DFE and Collaborative CDR in 22 nm CMOS ("Musah") was published December 2014 in IEEE Journal of Solid-State Circuits, Volume 49, No. 12. Musah qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

A 12.5-Gbps, 7-bit Transmit DAC with 4-Tap LUT-based Equalization in 0.13um CMOS ("Chung") was published September 21, 2008. Chung qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (pre-AIA).

U.S. Patent 10,659,337 to Jamal Riani et al. ("Riani") was filed on August 28, 2018. Riani qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

A 40-Gb/s Quarter Rate SerDes Transmitter and Receiver Chipset in 65 nm CMOS ("Zheng") was published November 2017 in IEEE Journal of Solid-State Circuits, Volume 52, No. 11. Zheng qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

U.S. Patent 10,461,918 issued to Kobayashi et al. ("Kobayashi") was filed on August 31, 2018 and published on October 29, 2019. Kobayashi qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

U.S. Patent Application 2016/0099795 filed by Lusted et al. ("Lusted") was published on April 4, 2016. Lusted qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

U.S. Patent Application 2009/0274170 filed by Yu et al. ("Yu") was published on November 5, 2009. Yu qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2)(AIA).

1
CREDO 2012
Amphenol v. Credo
IPR2025-00834

Chinese Patent 106953818 A to Yin ("Yin") was published July 14, 2017. Yin qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2)(AIA).

U.S. Patent 8,251,721 to Ikeda et al. was filed on August 23, 2010 and issued August 28, 2012. Ikeda qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

U.S. Patent 10,498,562 to Norimatsu ("Norimatsu") was filed April 8, 2016 and published on December 3, 2019. Norimatsu qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

Texas Instruments DS125DF1610 retimer chip datasheet ("DS125DF1610") was published on March 2017. DS125DF1610 qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

U.S. Patent No. 7,762,727 to Aronson ("Aronson") was filed on October 31, 2008, and was issued on July 27, 2010. Aronson qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

"High-speed Serial Bus Repeater Primer" Rev. 1.2 was published in October 2015 by Samie Samaan et al. ("Samaan"). Samaan qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

2

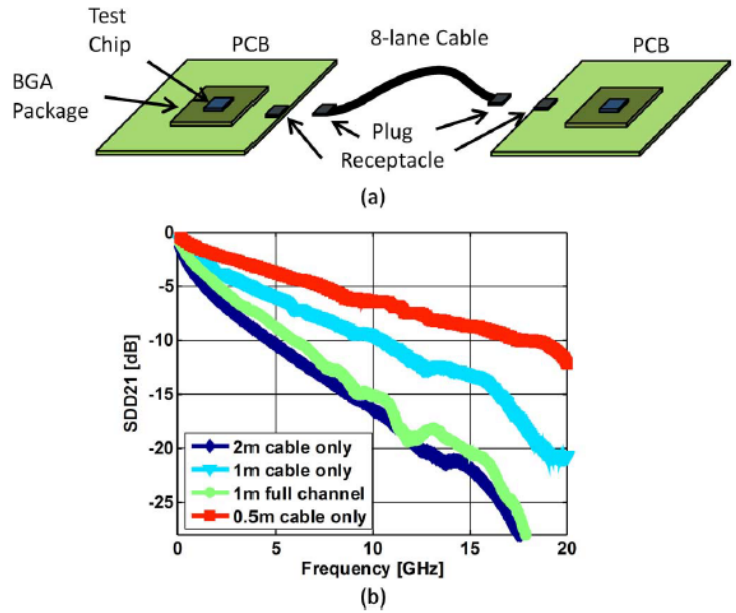
A. INDEPENDENT CLAIM 1

Claim 1	Disclosure
1[pre] An active Ethernet cable that comprises:	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: **1.1 Introduction** This specification defines the Common Management Interface Specification (CMIS) for QSFP-DD, OSFP and COBO in order to enable flexible use of modules and cable assemblies (hereafter referred to as modules unless cable assemblies are specifically mentioned) by the user. It shares some similarities with the management interfaces used in other form factors such as QSFP, SFP, and CFP. This specification supports a module with 8 host electrical lanes. Support for modules with more host lanes is possible (in increments of up to 8 lanes) by a bank selection feature (which allows several instances of an adressable page to exist in the paged register map). Due to this new mechanism this specification is not backwards compatible with SFF-8636. Some host ports may support both 4-channel and 8-channel devices; when a 4-channel device is inserted into such a port the appropriate memory map (e.g. SFF-8636) must be used. This case is outside the scope of this document. The host-to-module interactions needed to initialize active modules are presented in two state machines. A module-level state machine defines the initialization of the management interface as well as control of the module-wide power mode. A data path-specific state machine defines the host and module interactions and behaviors needed for power up and initialization of the data path(s) in the module. The CMIS is intended to cover a wide range of possible module functionalities and applications, ranging from cable assemblies to coherent DWDM modules. The CMIS is designed with the intent not to burden the simple cases but to provide a common set of management interactions and a common register map. In some modules, multiplexing of the high speed signals between the host electrical interface and the module media interface may occur in the module. In this specification, all references to lane numbers are based on the electrical connector interface lanes, unless otherwise indicated. In cases where a status or control aspect is applicable only to lanes after multiplexing or demultiplexing has occurred, the status or control is intended to apply to all lanes in the data path, unless otherwise indicated. Symmetry is expected between the Tx and the Rx hardware structure; for example Tx lanes that are multiplexed in the Tx are demultiplexed in the Rx. Each Application selected by the host is applied to the same Tx and Rx host electrical lanes. CMIS, 5-6.

Claim 1	Disclosure
	Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. Table 1 and Table 2 provide a summary of the high-level behaviors and properties of each module state for passive copper cable assembly and active module implementations, respectively. Implementers of passive copper cable assemblies should note that this specification was designed to allow for flat-memory EEPROM implementations for that cable type. Active modules and active cable assemblies shall use the paged memory model. For all module types, refer to sections 1.4.3-1.4.9 for detailed requirements for each state. CMIS, 18. ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.**

4

Claim 1	Disclosure
	Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling.

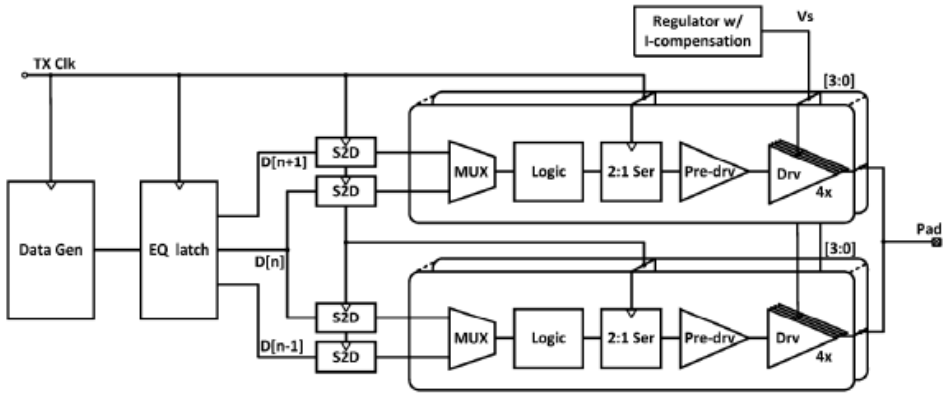
Claim 1	Disclosure
9	Musah, 3079-3080.  Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

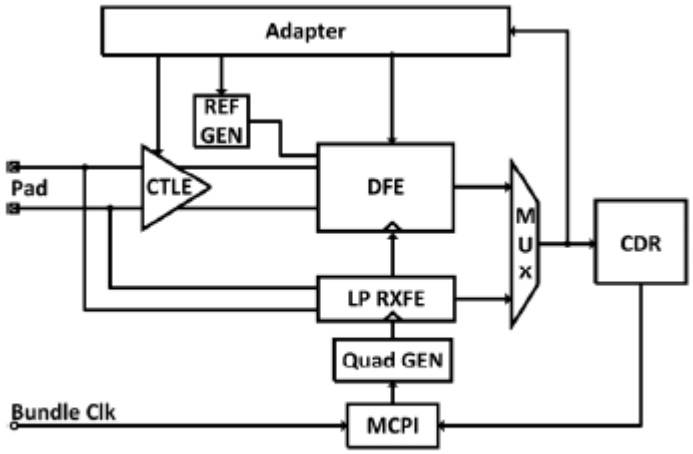
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

7

Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

8

Claim 1	Disclosure
6	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

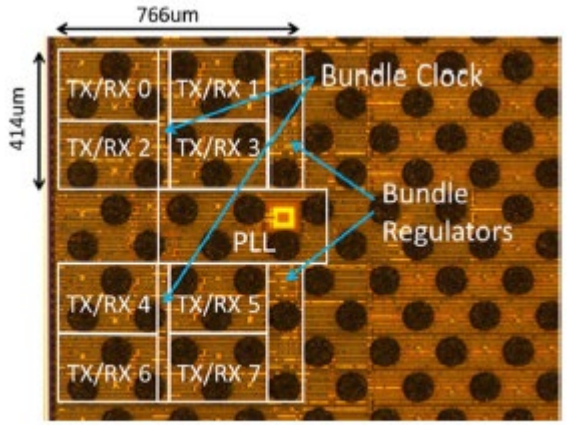
Claim 1	Disclosure
10	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 1	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

11

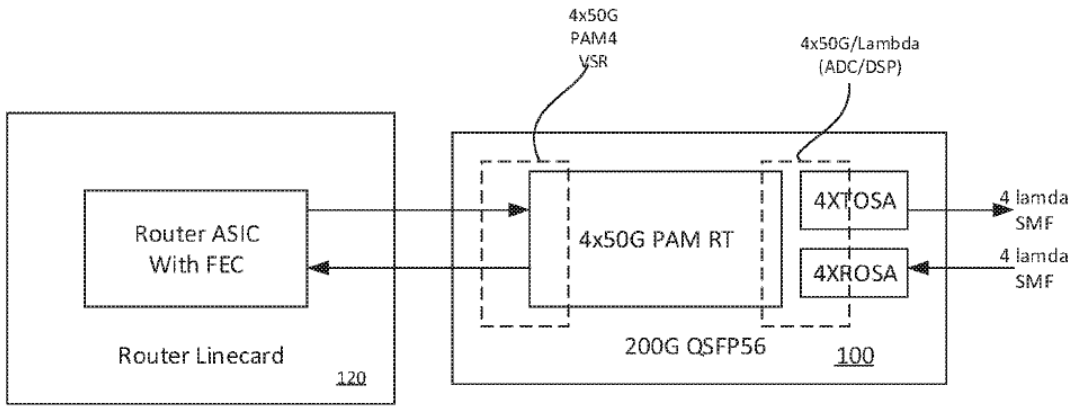
Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

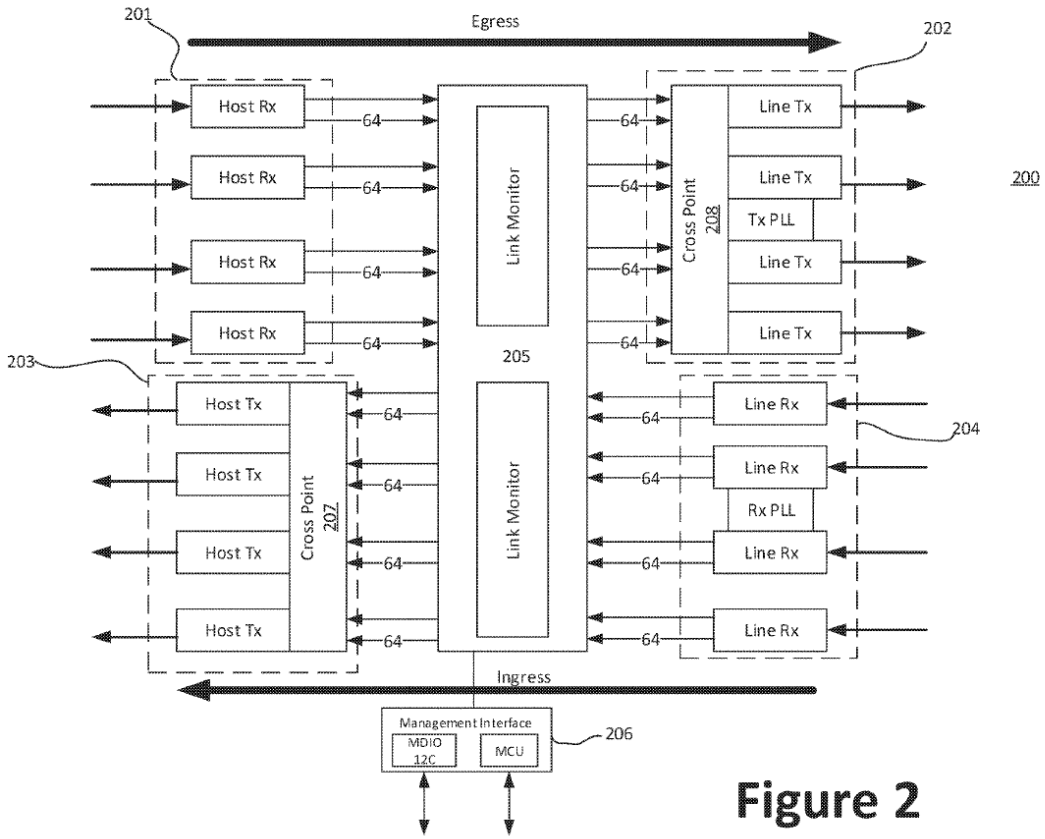
12

Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 1	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free $(BER < 10^{-12})$ transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088.

14

Claim 1	Disclosure
15	 **Figure 1** Riani, Figure 1.

Claim 1	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

17

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 1	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G

19

Claim 1	Disclosure
	data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below. Riani, 4:43-5:5. ***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER < 10^{-12} over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

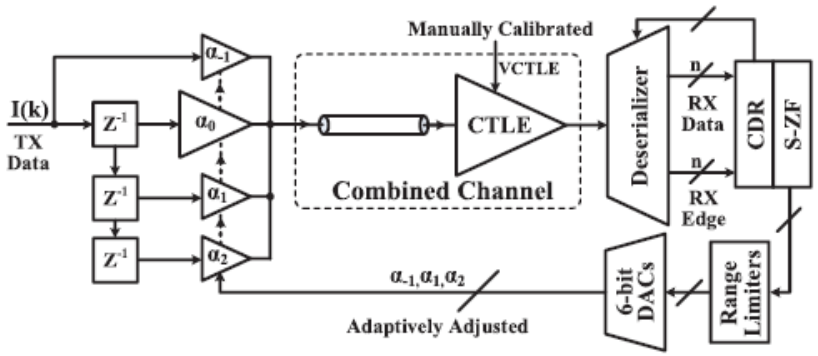
Claim 1	Disclosure
21	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

22

Claim 1	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 1	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

Claim 1	Disclosure
25	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

26

Claim 1	Disclosure

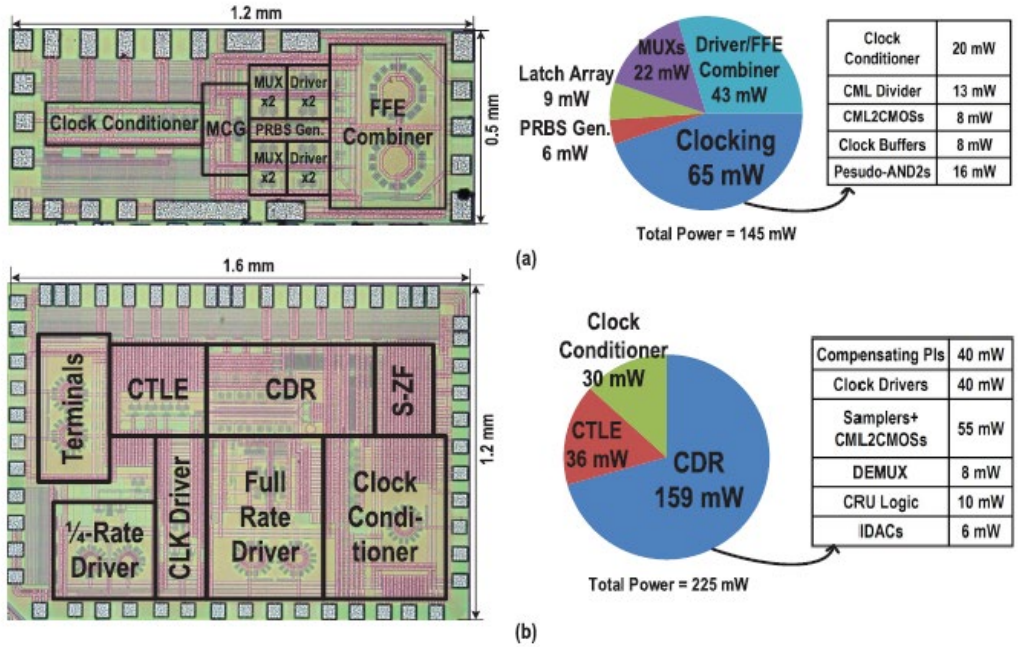


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

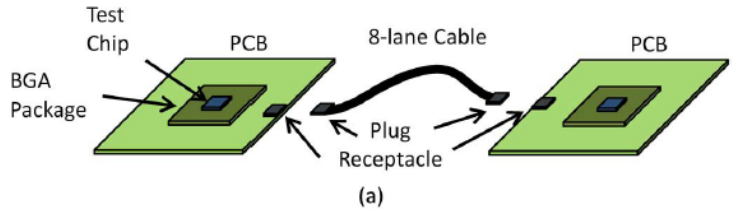
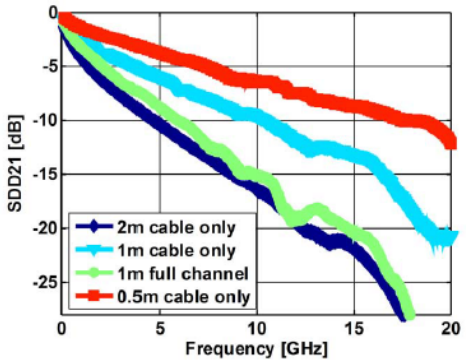
Claim 1	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 1, the data transmission system 1 is in accordance with DP1.4, for example, and includes an upstream device 2, a downstream device 4, and an active cable 3 connecting the upstream device 2 and downstream device 4." Kobayashi, 2:41-45.
1[a] electrical conductors connected between a first connector and a second connector,	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

28

Claim 1	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

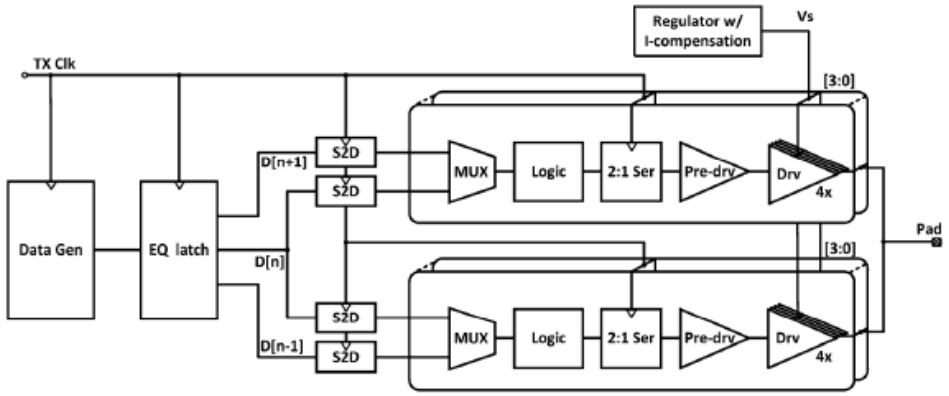
29

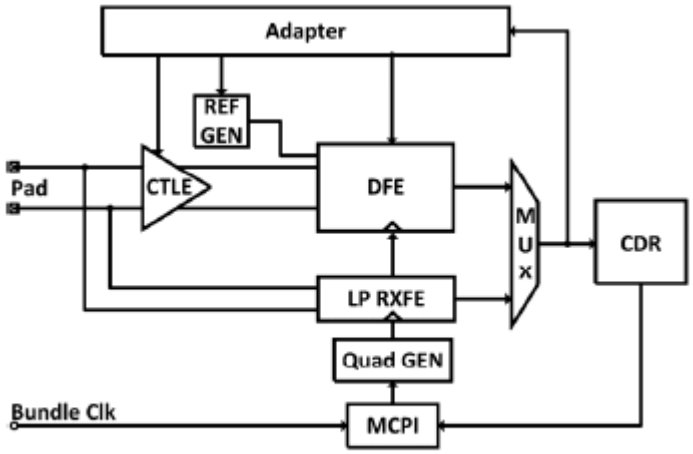
Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

Claim 1	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

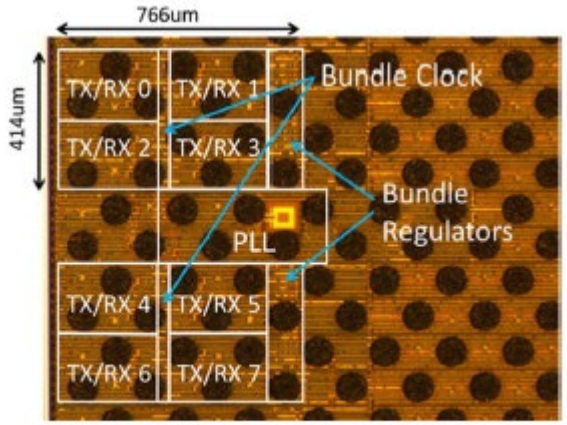
Claim 1	Disclosure
34	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 1	Disclosure
35	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 1	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

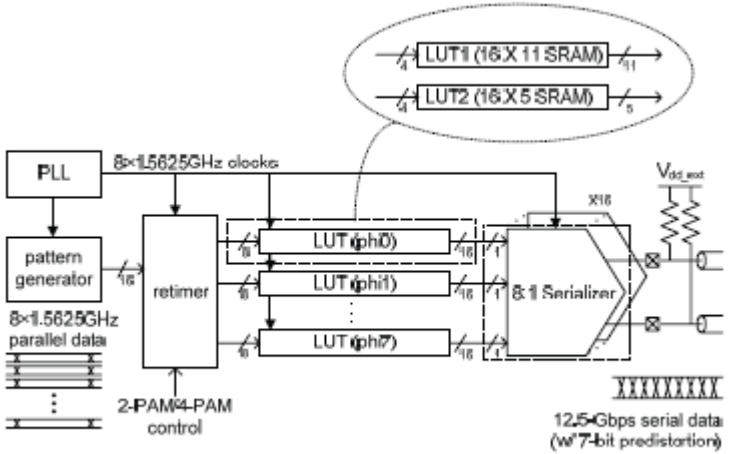
37

Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 1	Disclosure
39	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

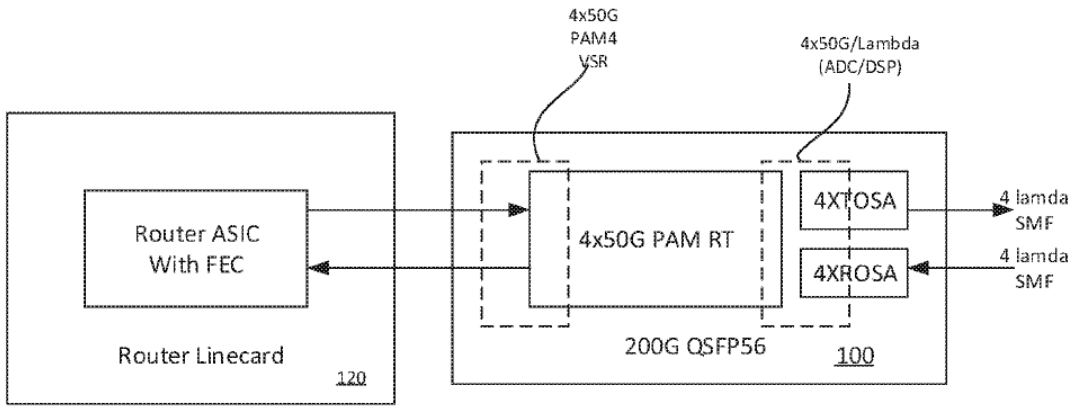
40

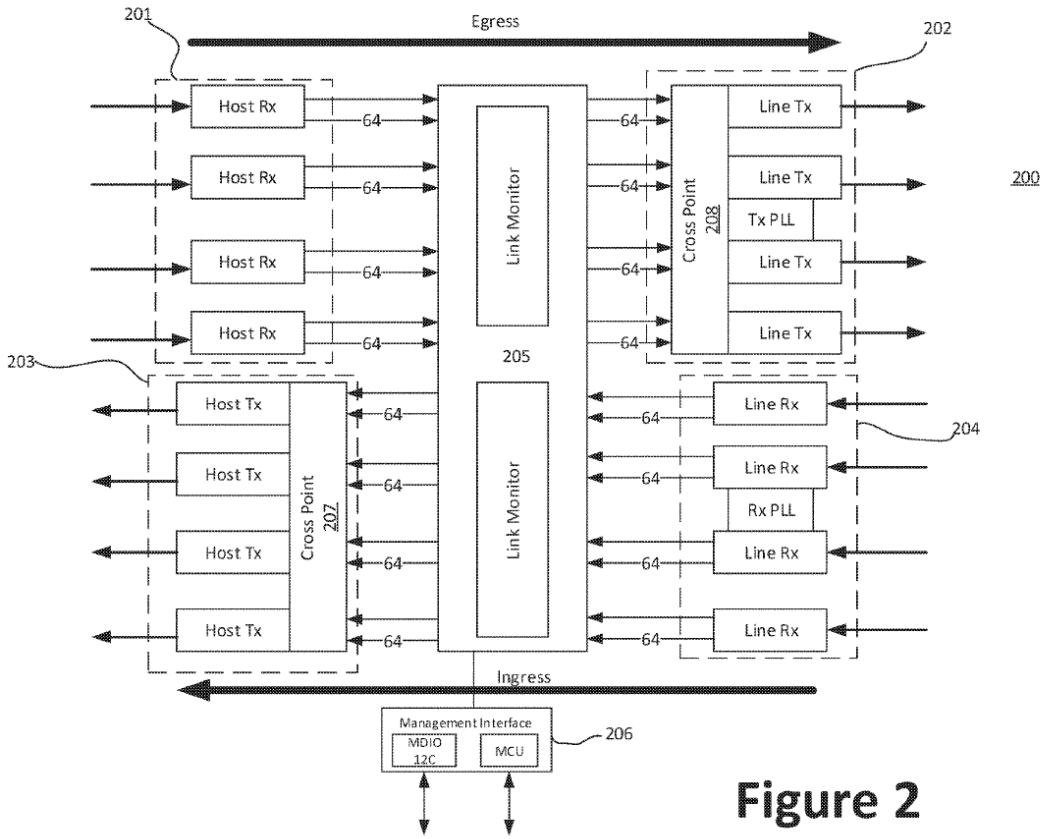
Claim 1	Disclosure
41	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 1	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

Claim 1	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

43

Claim 1	Disclosure
	 **Figure 1** Riani, Figure 1.

Claim 1	Disclosure
45	 Figure 2 Riani, Figure 2.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

47

Claim 1	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

48

Claim 1	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

Claim 1	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

50

Claim 1	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

51

Claim 1	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

Claim 1	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

53

Claim 1	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

54

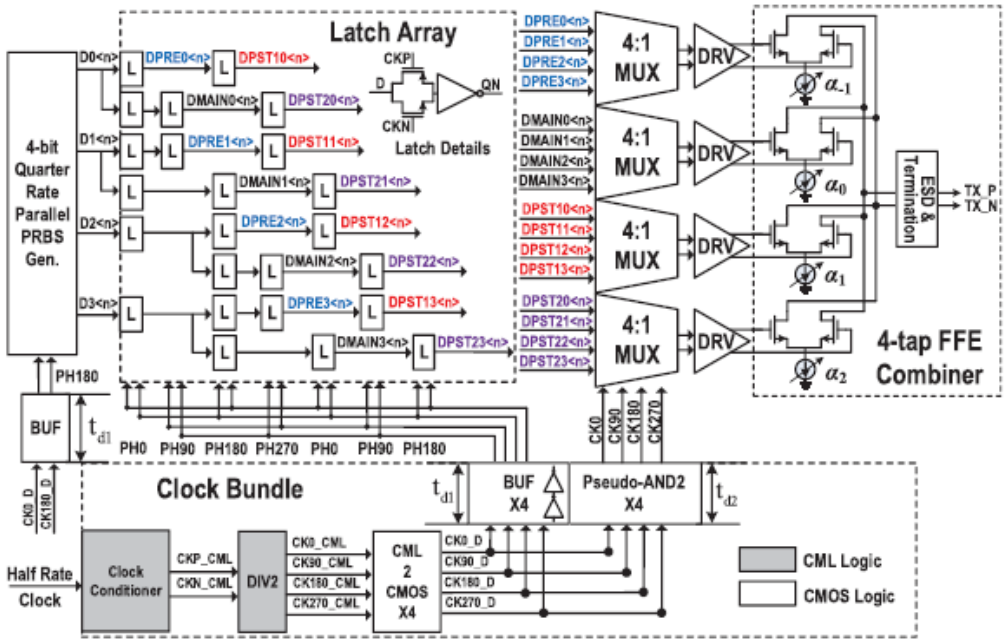
Claim 1	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

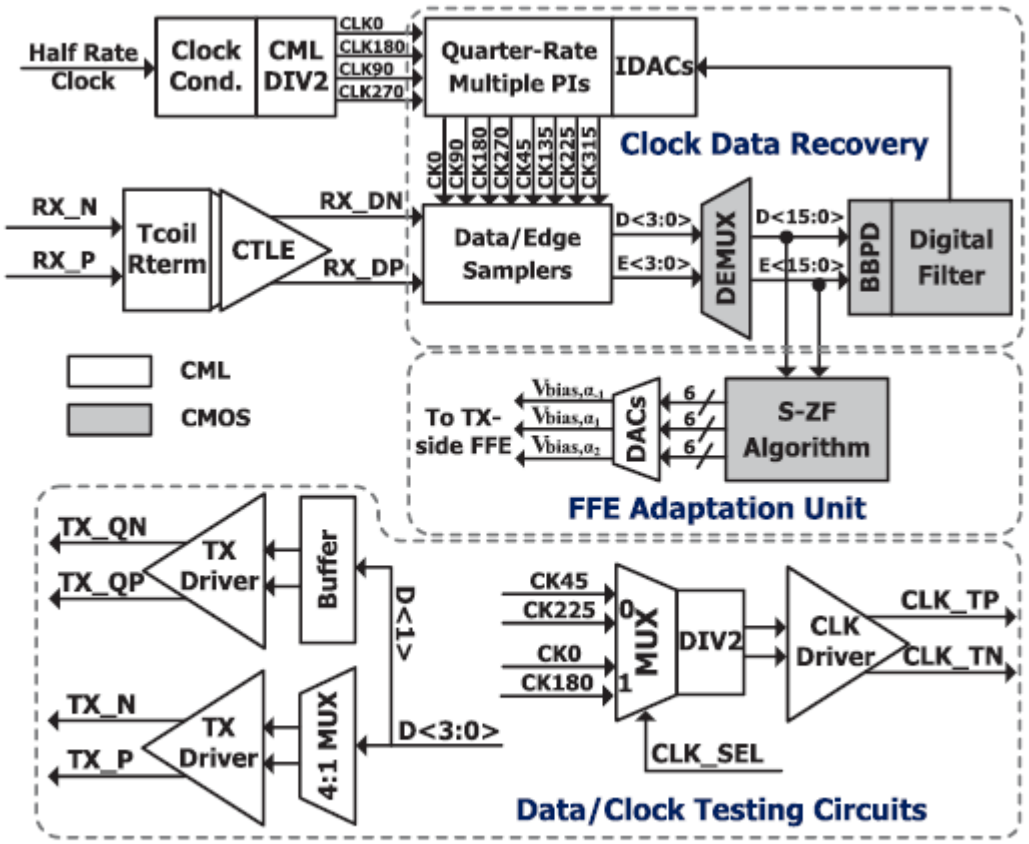
Claim 1	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

56

Claim 1	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

57

Claim 1	Disclosure
	 Fig. 1. Block diagram of the TX chip. Zheng, 2964.

Claim 1	Disclosure
59	 Fig. 8. Block diagram of the RX chip. Zheng, 2966.

Claim 1	Disclosure
60	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

Claim 1	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 1	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

62

Claim 1	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

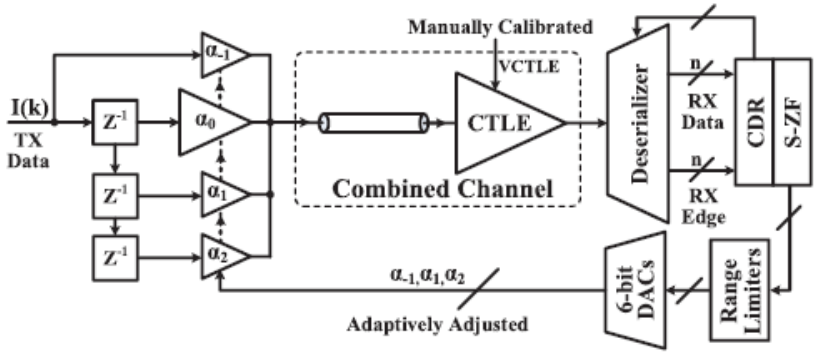


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

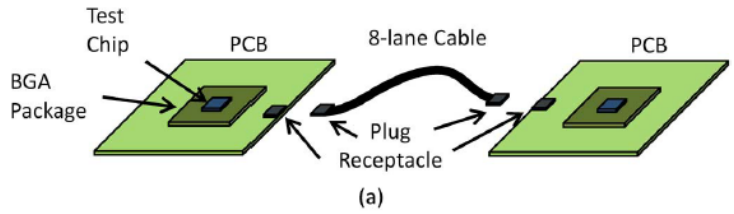
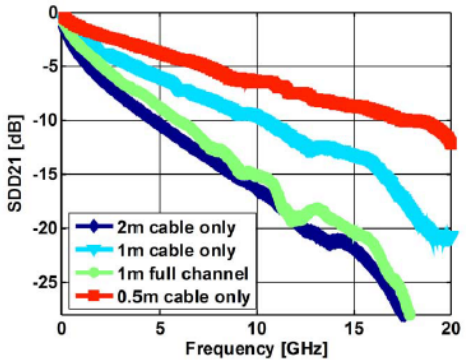
Claim 1	Disclosure
	Zheng, 2970. This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 1, the data transmission system 1 is in accordance with DP1.4, for example, and includes an upstream device 2, a downstream device 4, and an active cable 3 connecting the upstream device 2 and downstream device 4." Kobayashi, 2:41-45.
1[b] each of the first and second connectors being adapted to fit into an Ethernet port of a	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

Claim 1	Disclosure
corresponding host device to receive from that host device an electrical input signal conveying an inbound data stream to the cable and to provide to that host device an electrical output signal conveying an outbound data stream from the cable,	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

65

Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

66

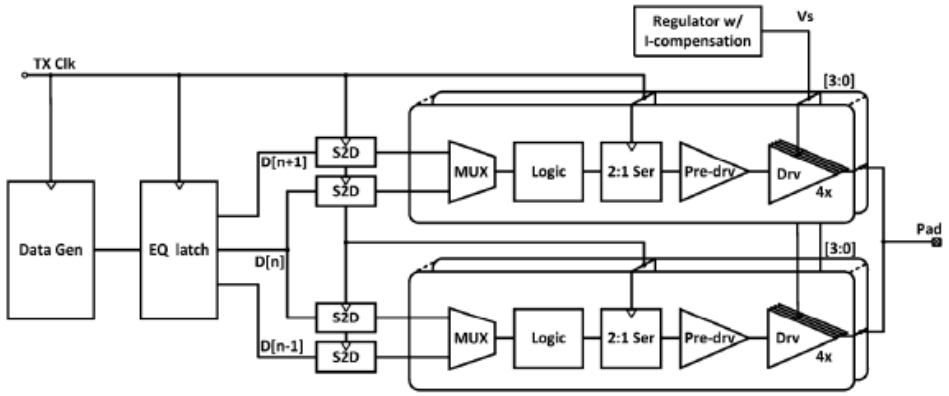
Claim 1	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

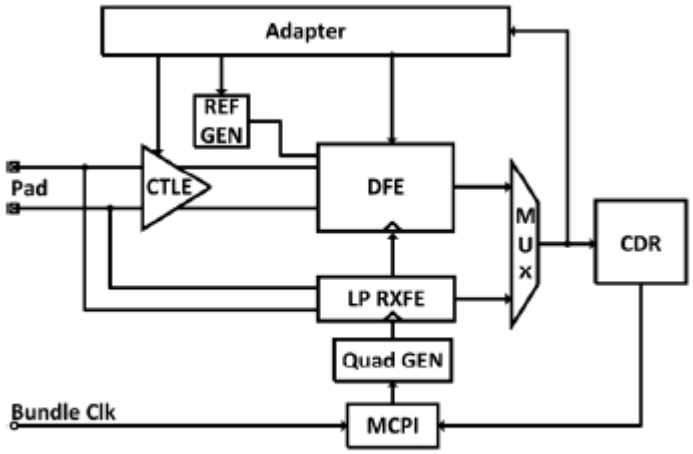
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

68

Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

69

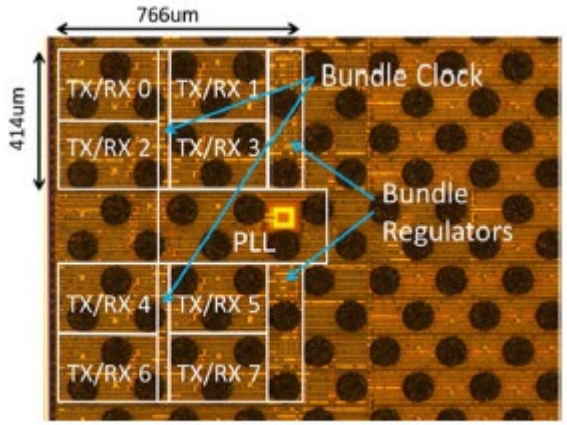
Claim 1	Disclosure
70	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 1	Disclosure
71	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 1	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

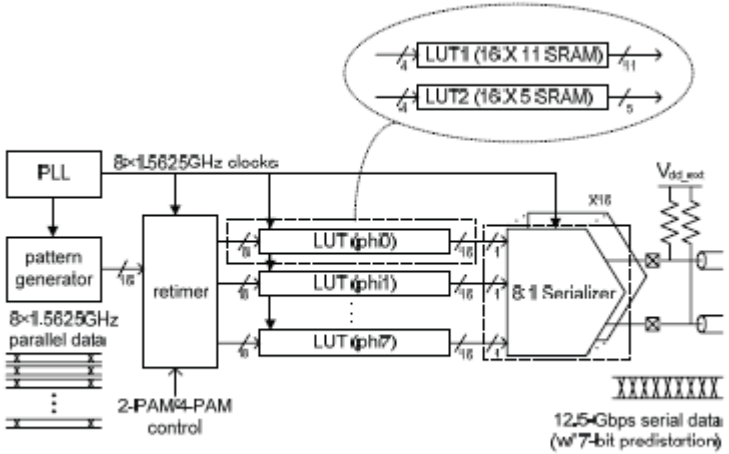
73

Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 1	Disclosure
75	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

76

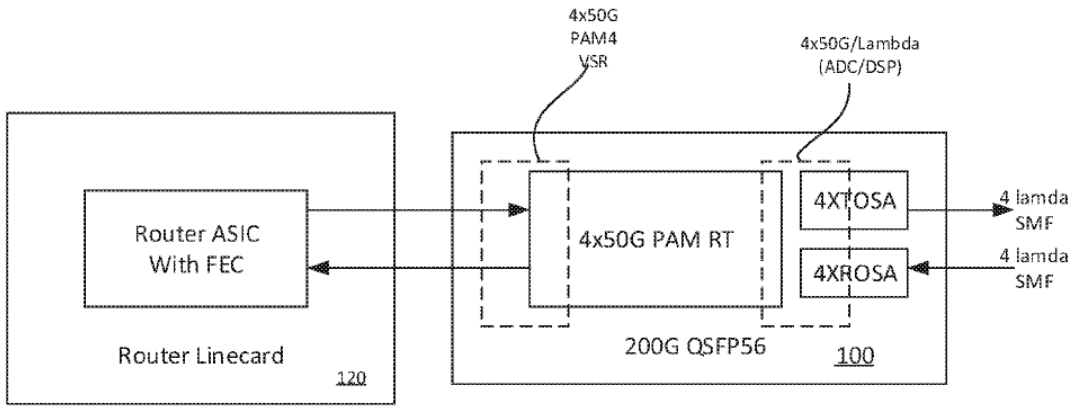
Claim 1	Disclosure
77	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 1	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

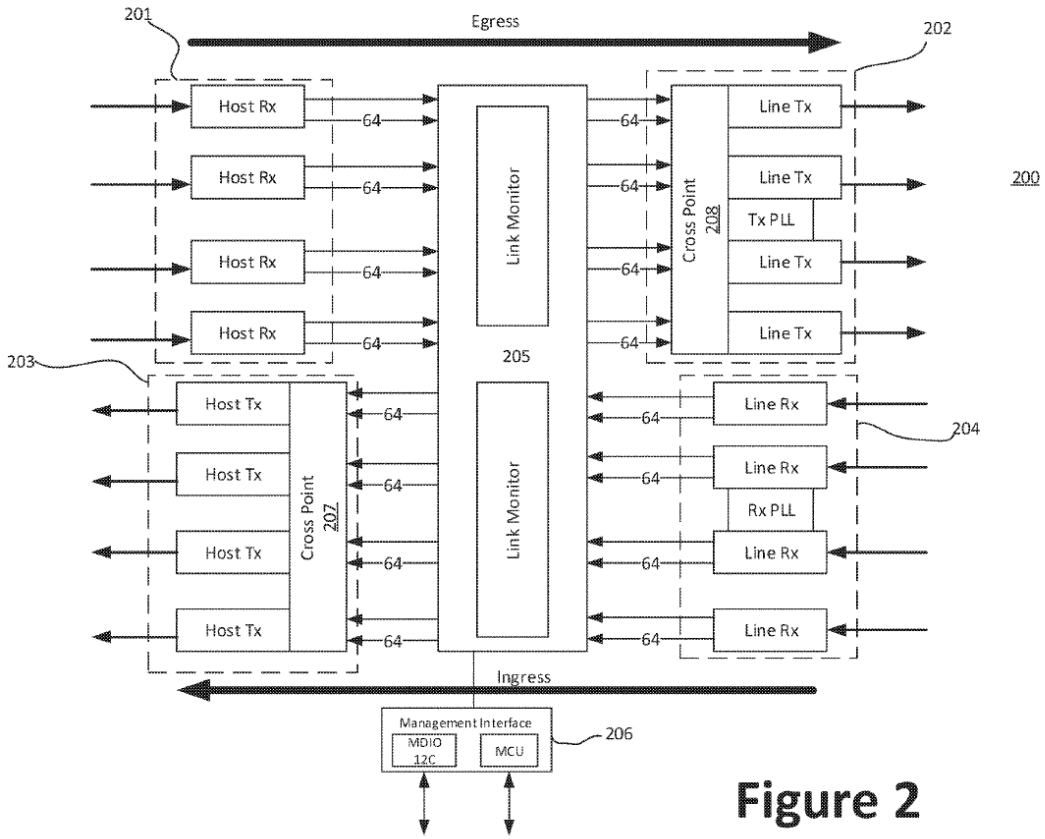
78

Claim 1	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

79

Claim 1	Disclosure
	 **Figure 1** Riani, Figure 1.

80

Claim 1	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 1	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

84

Claim 1	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

85

Claim 1	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

86

Claim 1	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

87

Claim 1	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

88

Claim 1	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

89

Claim 1	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

Claim 1	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector.

Claim 1	Disclosure
	Riani, 11:4-45. It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

92

Claim 1	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

93

Claim 1	Disclosure
94	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

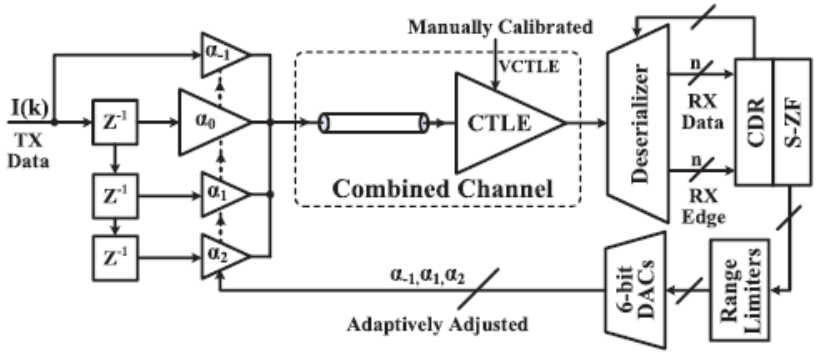
95

Claim 1	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

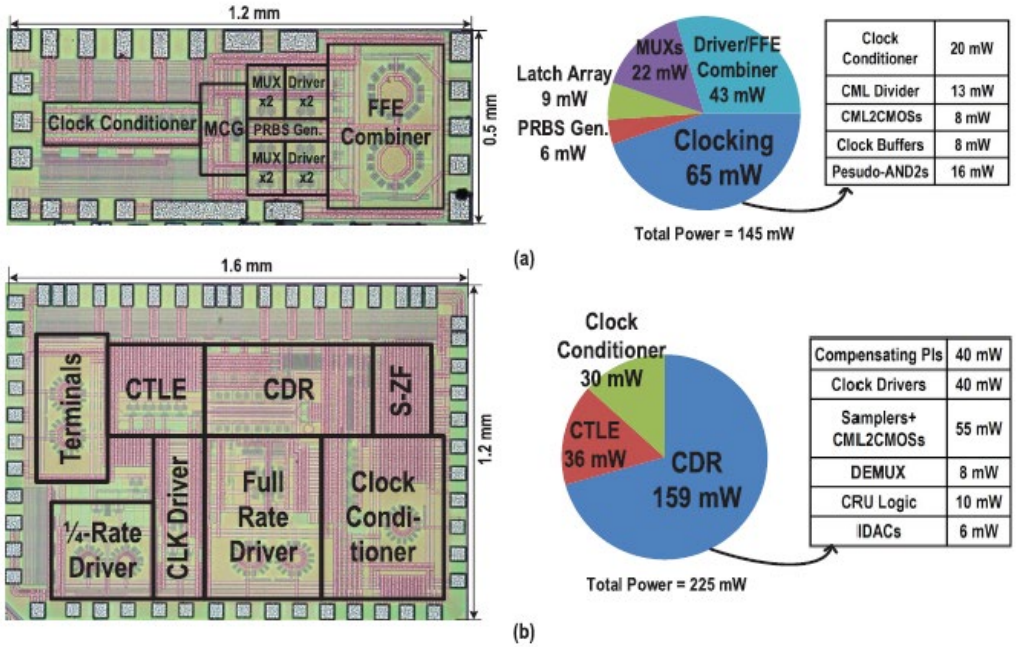
96

Claim 1	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

97

Claim 1	Disclosure
98	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 1	Disclosure
66	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

Claim 1	Disclosure
100	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

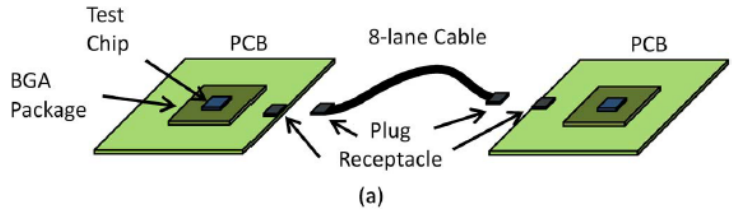
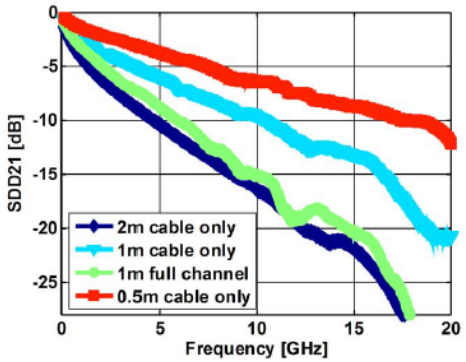
Claim 1	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "In the data transmission system 1 having the above configuration, the source 20 in the upstream device 2 generates a serial bit stream which is a video signal, for example. The source 20 transmits the generated serial bit stream from the downstream facing TX 20a to the retimer 21. The retimer 21 receives the serial bit stream transmitted from the downstream facing TX 20a in the upstream facing TX 21a. The retimer 21 performs a waveform shaping on the received serial bit stream and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 21b to the active cable 3." Kobayshi, 2:60-3:4.
1[c] each of the first and second connectors including a respective	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

101

Claim 1	Disclosure
transceiver that performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream,	**_Abstract_—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

102

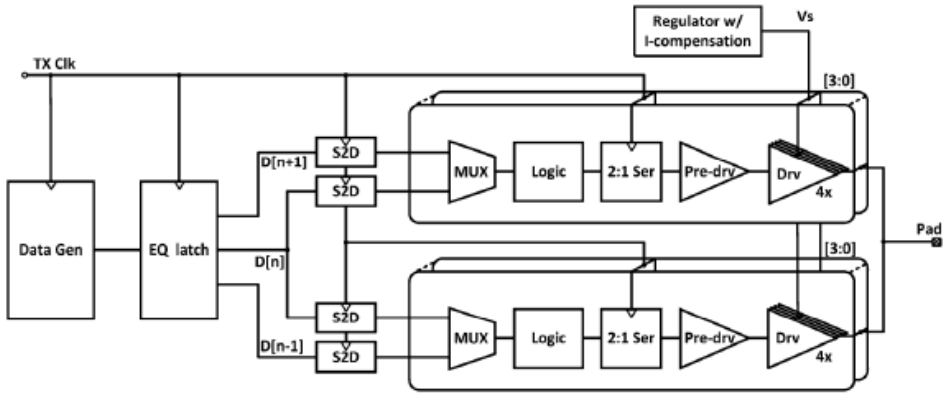
Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

Claim 1	Disclosure
104	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

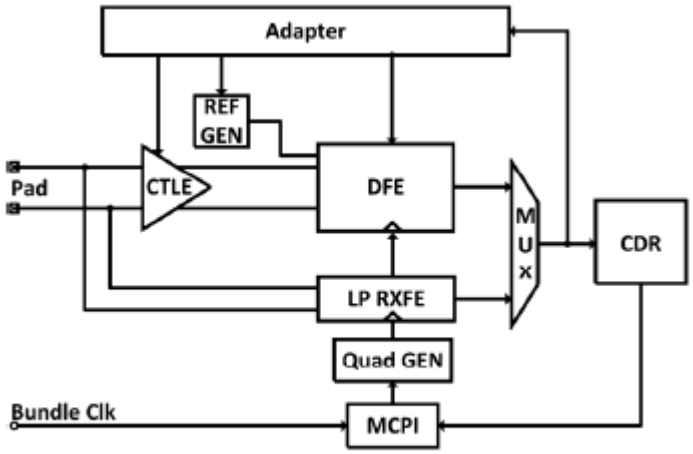
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

106

Claim 1	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

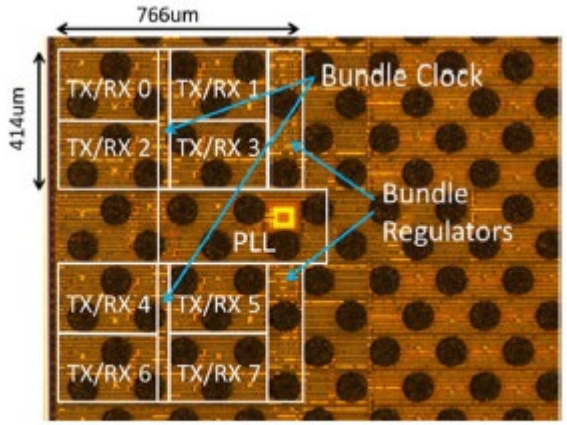
107

Claim 1	Disclosure
108	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 1	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

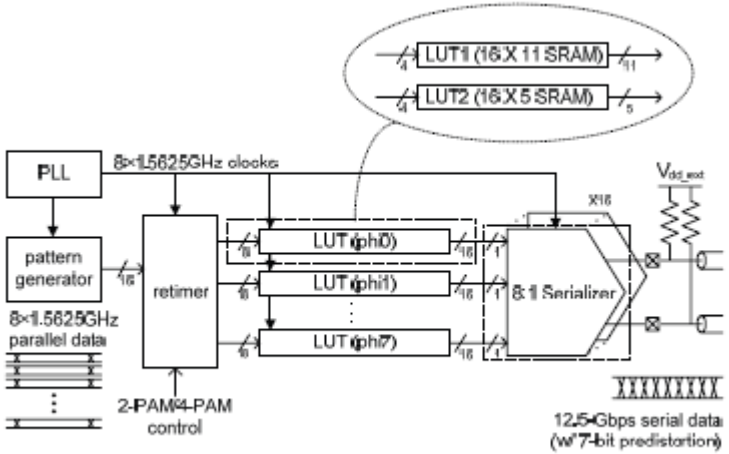
110

Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 1	Disclosure
112	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

113

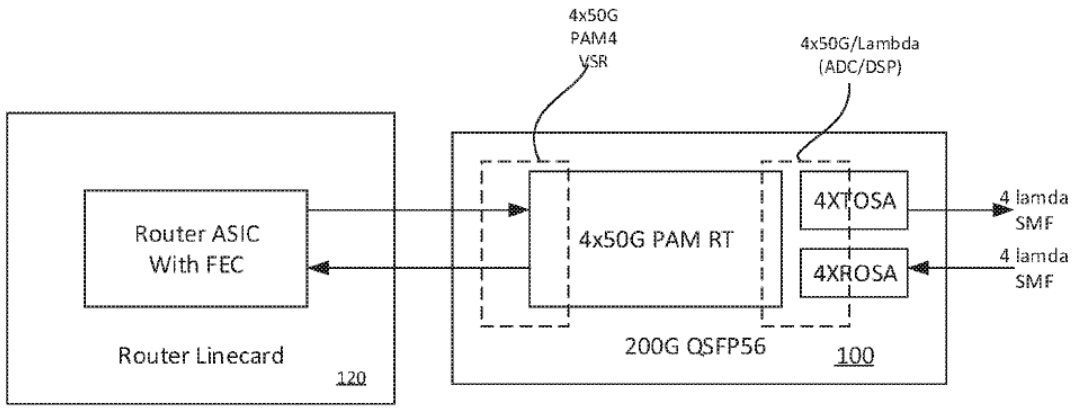
Claim 1	Disclosure
114	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

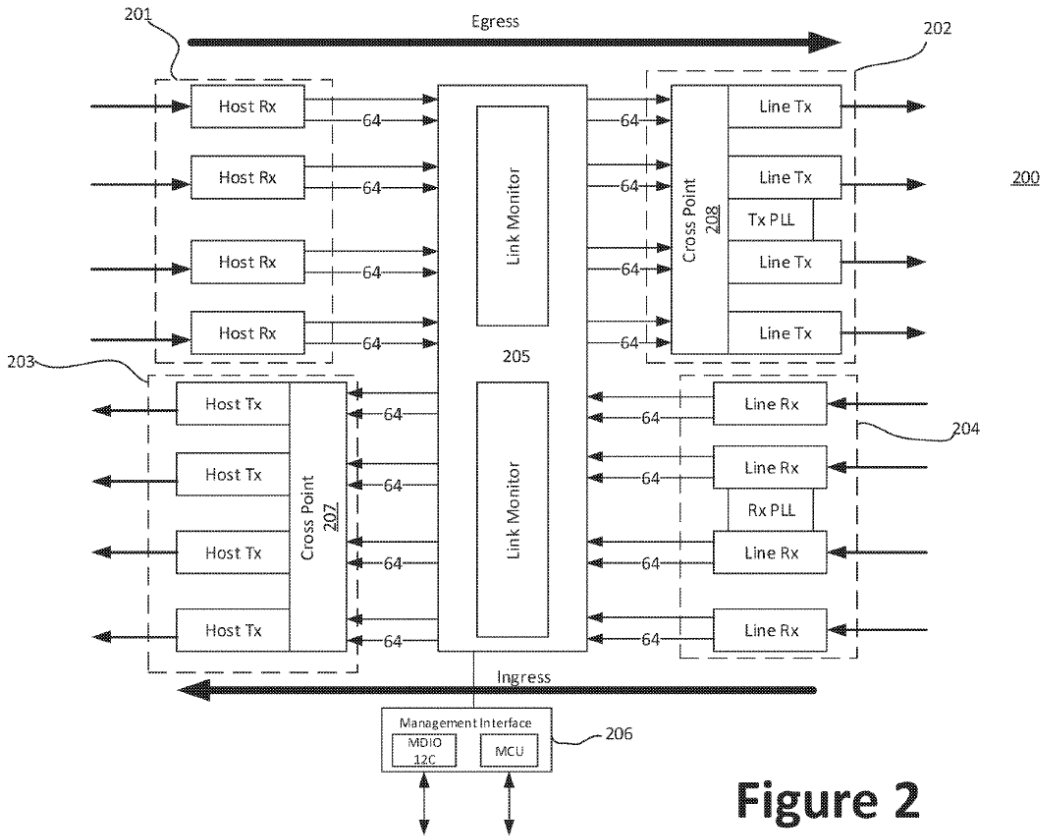
Claim 1	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

115

Claim 1	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

116

Claim 1	Disclosure
117	 **Figure 1** Riani, Figure 1.

Claim 1	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

119

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

120

Claim 1	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

121

Claim 1	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

122

Claim 1	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

123

Claim 1	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

124

Claim 1	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

125

Claim 1	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

126

Claim 1	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

Claim 1	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

Claim 1	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

Claim 1	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

130

Claim 1	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

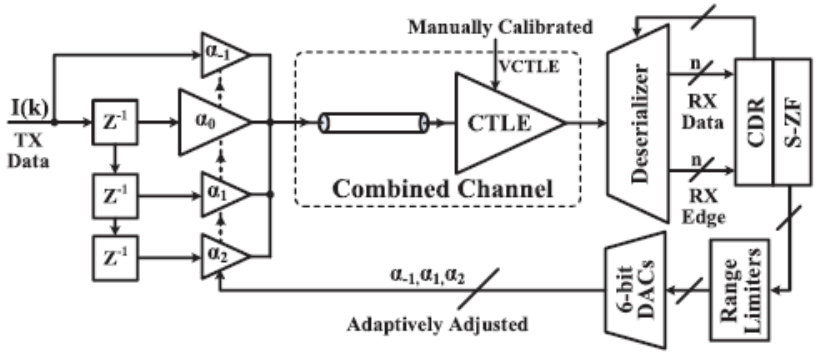
131

Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

132

Claim 1	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 1	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

Claim 1	Disclosure
135	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

136

Claim 1	Disclosure

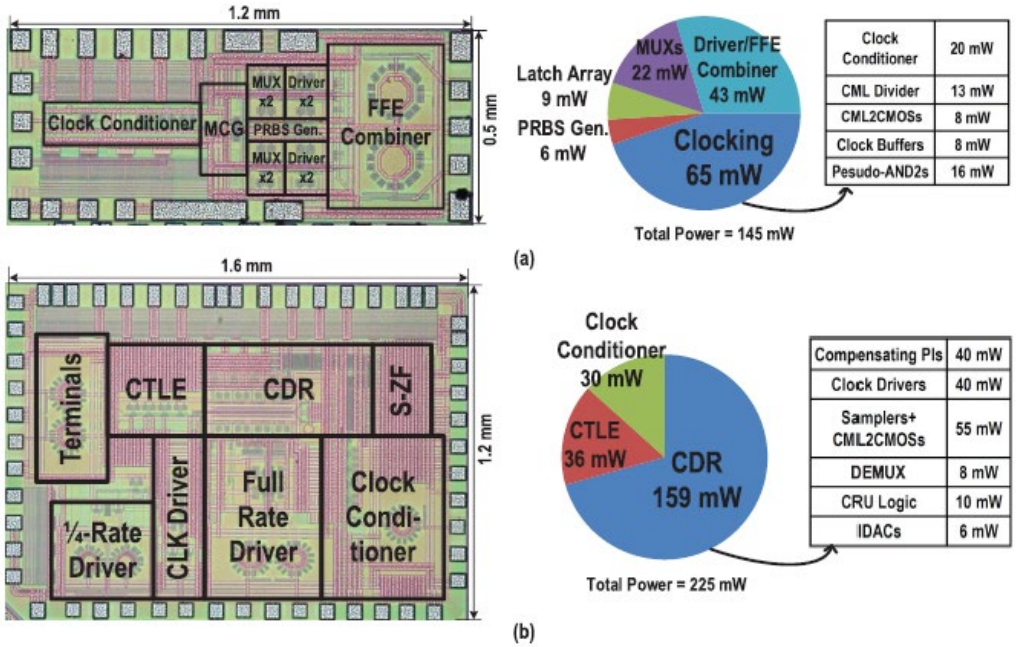
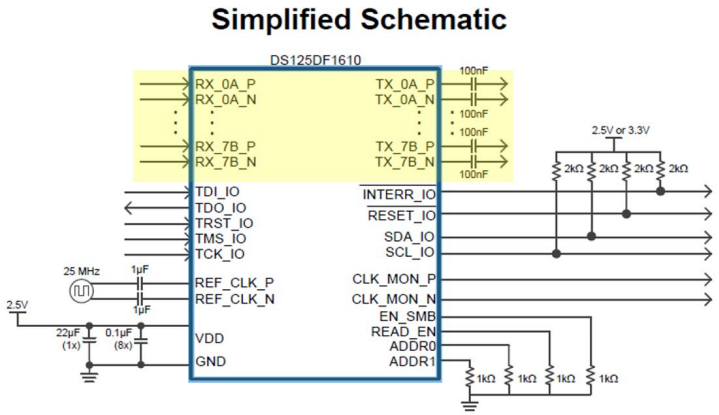
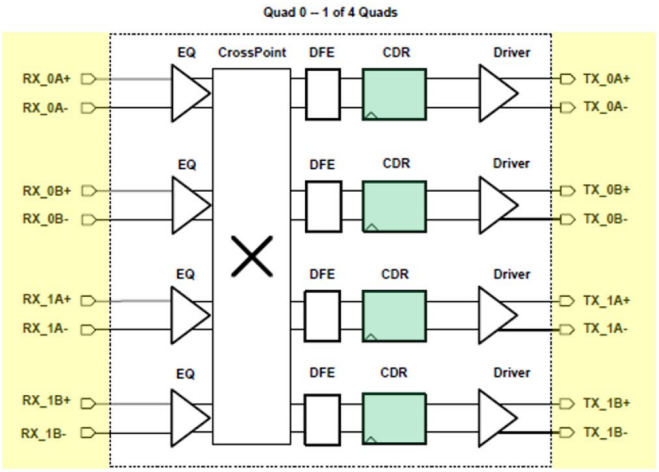


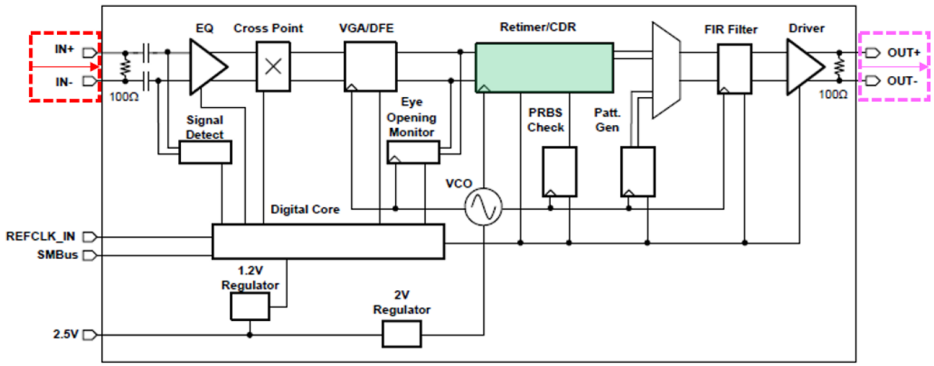
Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

Claim 1	Disclosure
138	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

Claim 1	Disclosure
	Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 circuit is illustrated in the Simplified Schematic and Figure 1 below. The highlighted pins illustrate sixteen channels of data, with each channel having two differential (Positive and Negative) input pins (RX_0A_P/N through RX_7B_P/N) to receive the channel of data and corresponding output pins (TX_0A_P/N through TX_7B_P/N) to transmit the channel of data after processing. *See, e.g.*, DS125DF1610 at §§ 3, 7.2. Thus, DS125DF1610's retimer is a transceiver that, when used in Tang's cable, receives one or more inbound data streams (on RX_0A_P/N through RX_7B_P/N), and transmits one or more transit data streams (via pins TX_0A_P/N through TX_7B_P/N).

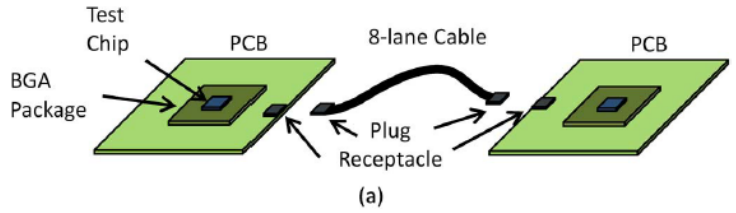
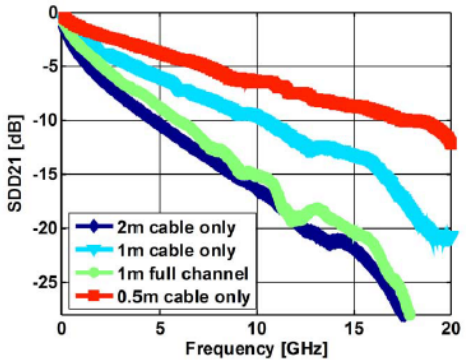
Claim 1	Disclosure
140	  **Figure 1. DS125DF1610 Simplified Cross Point Diagram** Figure 1 above shows four channels and illustrates that DS125DF1610's retimer processes each data channel by, inter alia, performing "clock and data recovery (CDR)." *See id.* at §§ 3, 7.1. A channel's data path is shown in more detail in DS125DF1610 Figure 2 (below), which illustrates the input signal conveying an inbound data stream and the transit signal conveying a transit data stream in Tang's cable

Claim 1	Disclosure
141	when DS125DF1610's retimer implements Tang's signal driver 44.  Figure 2. DS125DF1610 Simplified Data Path Diagram The Retimer/CDR circuit performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream. *See id*. at § 7.3.6. The '252 Patent specification makes clear that the input signal from the host can be pre-processed before CDR is performed on it. An ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data for output as the transit data stream because after extracting the input signal's timing and using that timing to sample and extract the input signal's data in an analog-to-digital CDR process, the retimer generates a new data stream for equalization by the FIR Filter (a digital-to-analog process) and to be driven for output by the differential driver (which produces an analog signal with configurable output voltage levels) to convey the output data stream (transit data stream when outputting over the cable) as an analog signal over the cable conductors, where the new signal is a "re-modulation" of the data. *See id*. at §§ 7.3.6-7.3.9. .

Claim 1	Disclosure
1[d] the respective transceiver for each of the first and second connectors performing clock and data recovery on the respective electrical transit signal to extract and remodulate the transit data stream as the outbound data stream from the cable, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

142

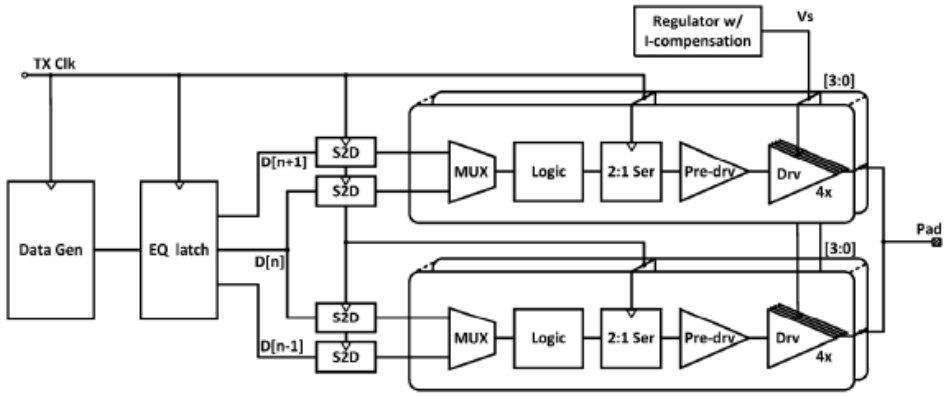
Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

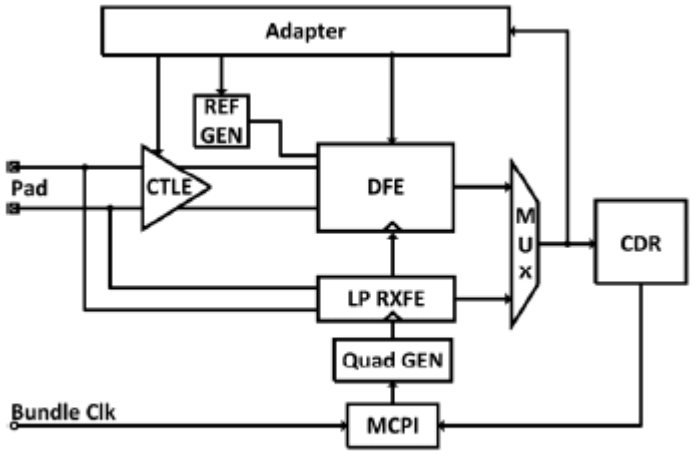
Claim 1	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

145

Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

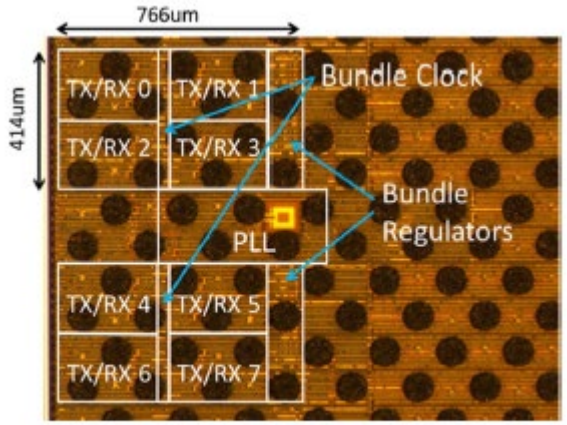
Claim 1	Disclosure
147	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 1	Disclosure
148	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 1	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

150

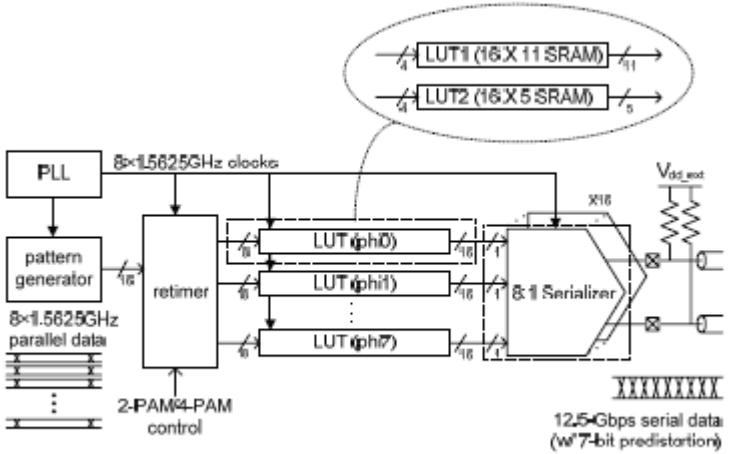
Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 1	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

152

Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

153

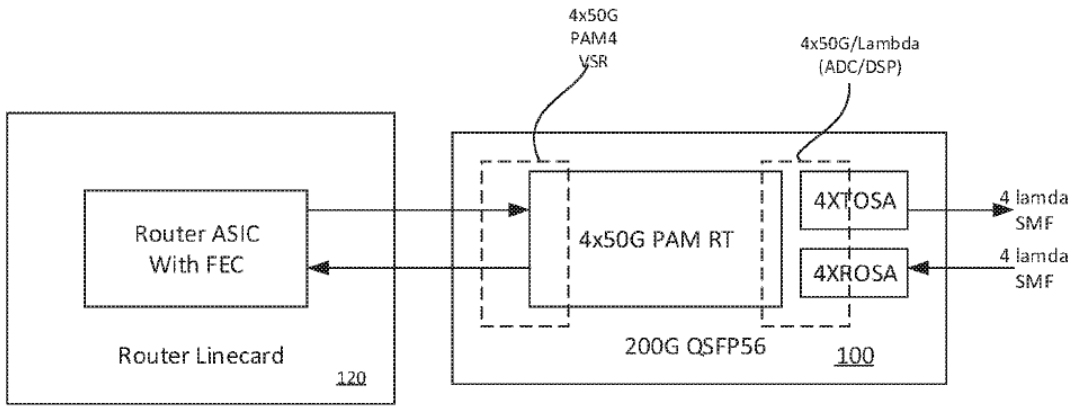
Claim 1	Disclosure
154	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 1	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

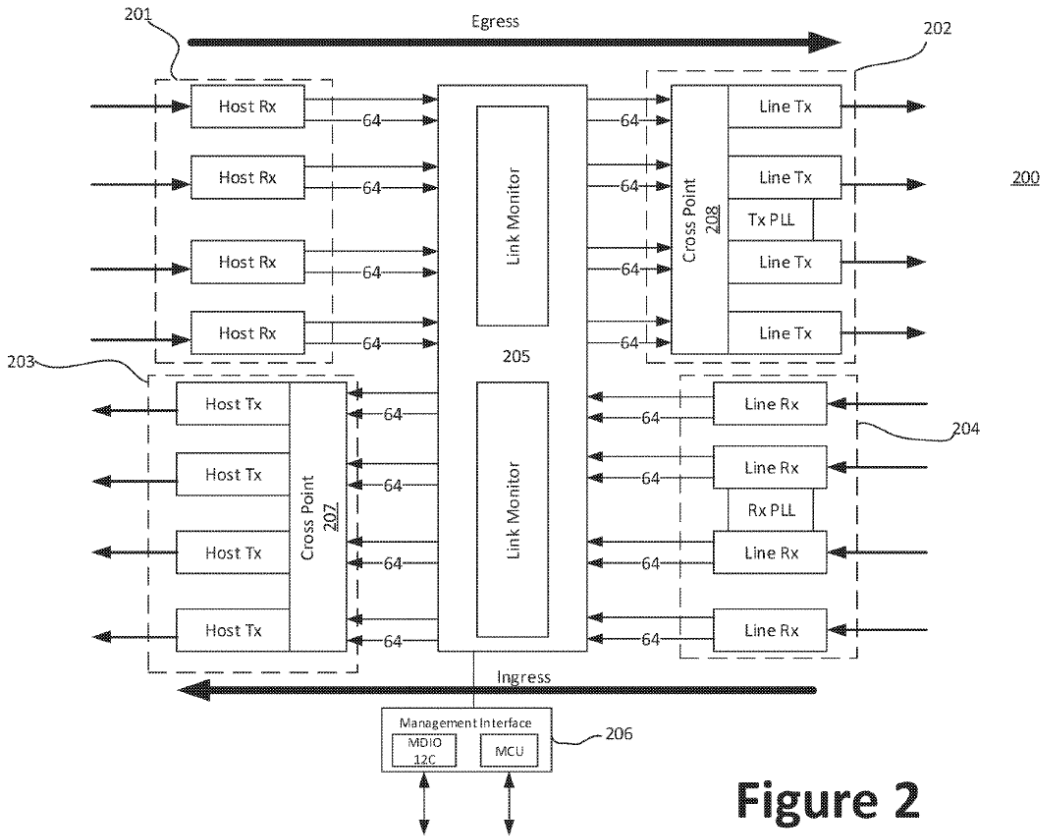
155

Claim 1	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

156

Claim 1	Disclosure
	 **Figure 1** Riani, Figure 1.

157

Claim 1	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

159

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

160

Claim 1	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

161

Claim 1	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

Claim 1	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

163

Claim 1	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

164

Claim 1	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

165

Claim 1	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

166

Claim 1	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

167

Claim 1	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

168

Claim 1	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

169

Claim 1	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

170

Claim 1	Disclosure
171	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

172

Claim 1	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 1	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

174

Claim 1	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

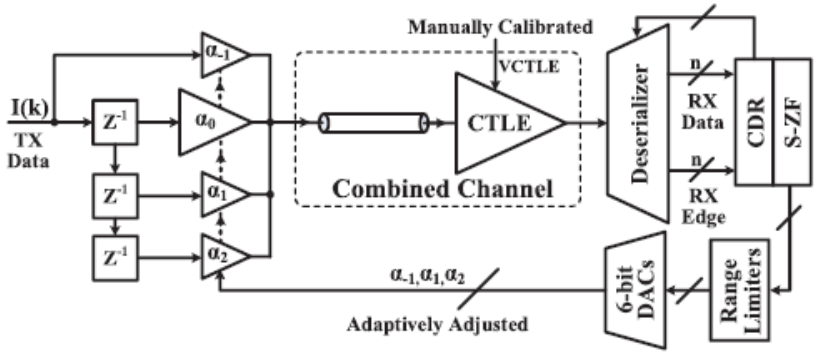


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

176

Claim 1	Disclosure

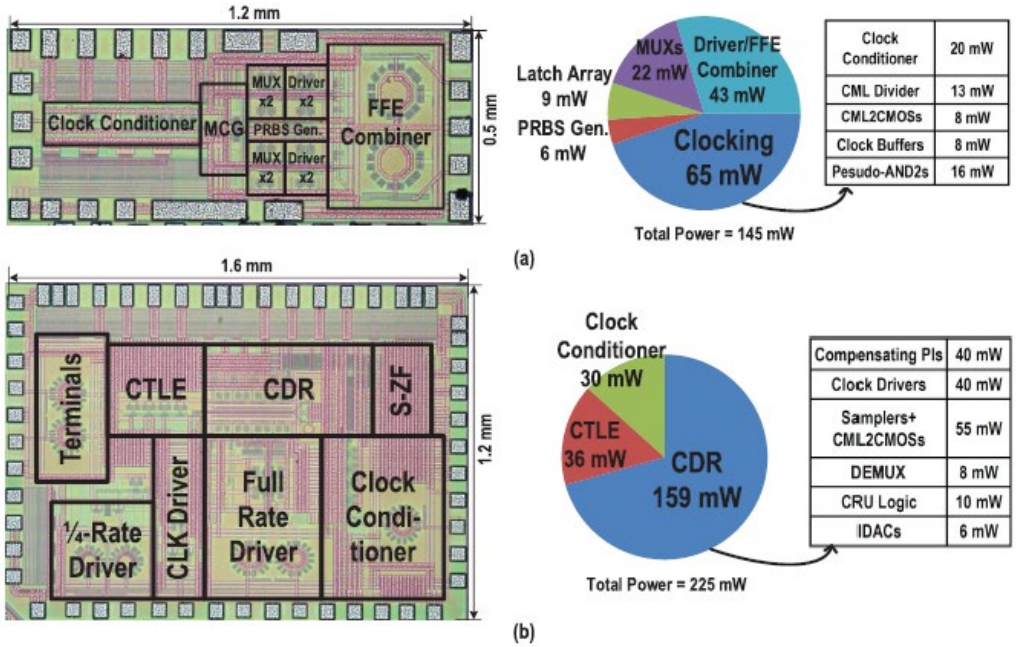


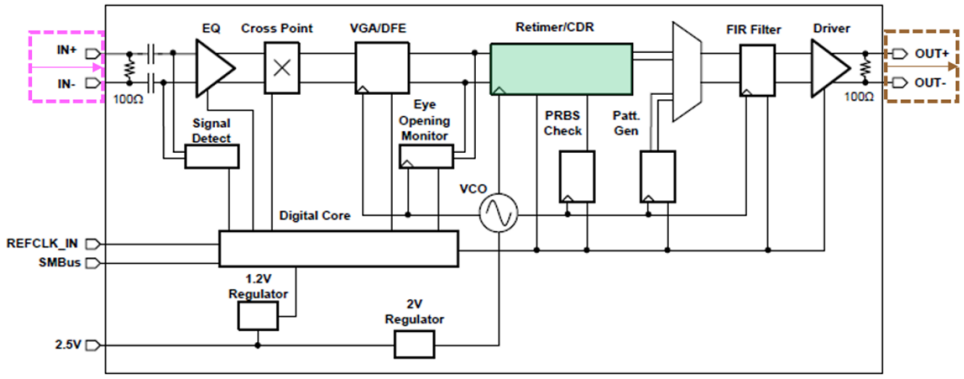
Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

177

Claim 1	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

178

Claim 1	Disclosure
	Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 retimer performs CDR for respective transit signals conveying transit data streams received from electrical conductors to generate output signals conveying outbound data streams that exit cable 10 to respective hosts, as shown in DS125DF1610 Fig. 2 below. For similar reasons to those discussed in claim limitation 1[c], an ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data as the outbound data after extracting the input signal's data and generating a newly-timed signal that is a "re-modulation" of the data.  **Figure 2. DS125DF1610 Simplified Data Path Diagram**

179

Claim 1	Disclosure
	5. The PCIe 3.0 Equalization Adaptation Protocol The PCIe 3.0 Base Specification [10] requires that a compliant Tx be able to produce any of the output equalization (TxEQ) levels shown in Figure 21, whenever requested to do so by the receiving port, during link equalization training. The same specification describes a protocol of TxEQ adaptation, where the host (upstream agent) or the end device (downstream agent) can each request that any of the TxEQ values in Figure 21 be presented to them by the other agent on the link. The intent of adaptation is to allow both agents to adjust the link partner's TxEQ to an optimal value for their receiver, and for the specific channel and operating conditions at hand. Samaan, p. 24.
1[e] the respective transceivers each employing fixed, cable-independent equalization parameters for each of: the remodulation of the transit data stream as the outbound data stream, and the clock and data recovery performed on the electrical input signal.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1.

180

Claim 1	Disclosure
	1.4.4 MgmtInit State The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

181

Claim 1	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

182

Claim 1	Disclosure
183	Table 10- Control field dependency on Explicit Control bit Control field / Explicit Control value Tx Adaptive Input Eq Enable / 1 Tx Adaptive Input Eq Recall / 1 or 0 Tx Input Eq control / 1 Tx CDR control / 1 Rx CDR control / 1 Rx Output Eq control, pre-cursor / 1 Rx Output Eq control, post-cursor / 1 Rx Output Amplitude control / 1 CMIS, 38. **1.5.4.1 Control Set Usage** There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

Claim 1	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

184

Claim 1	Disclosure
185	**1.7 Module Memory Map** This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module. CMIS, 44. **1.5.5 Signal Integrity Control Field Background** The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated. The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11). Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable Control / Tx Adaptive Input Eq Enable value to use this control Tx Input Eq Adaptation Freeze — 1 Tx Input Eq Adaptation Store — 1 Tx Adaptive Input Eq Recall — 1 Tx Input Eq control — 0 The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module. CMIS, 40.

Claim 1	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

Claim 1	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

187

Claim 1	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

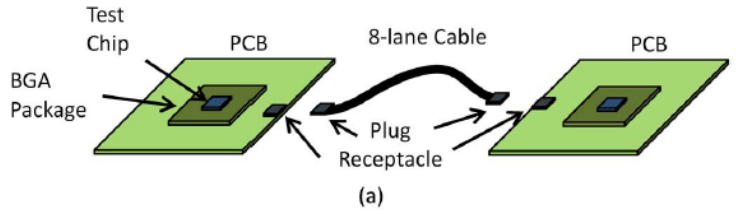
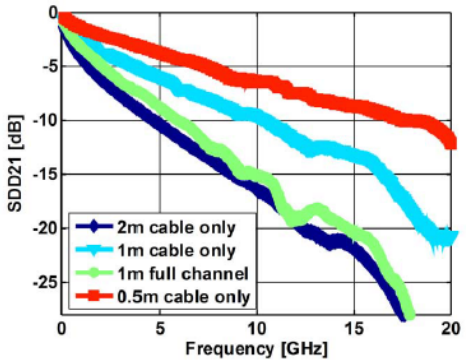
Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

Claim 1	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

189

Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

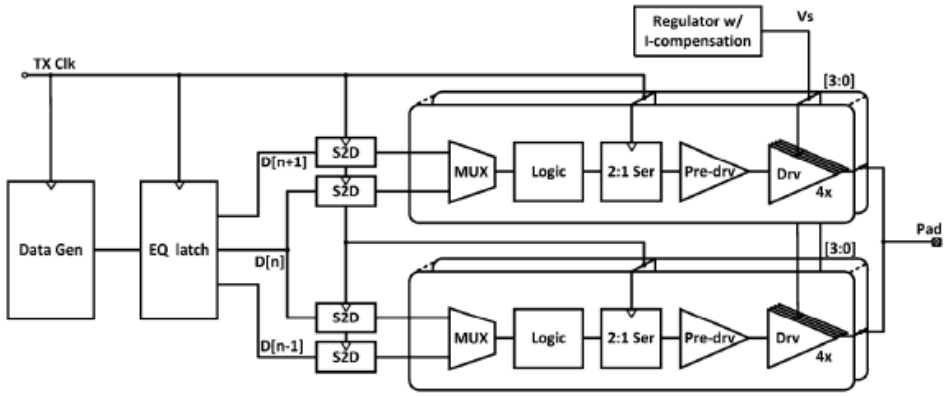
Claim 1	Disclosure
191	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

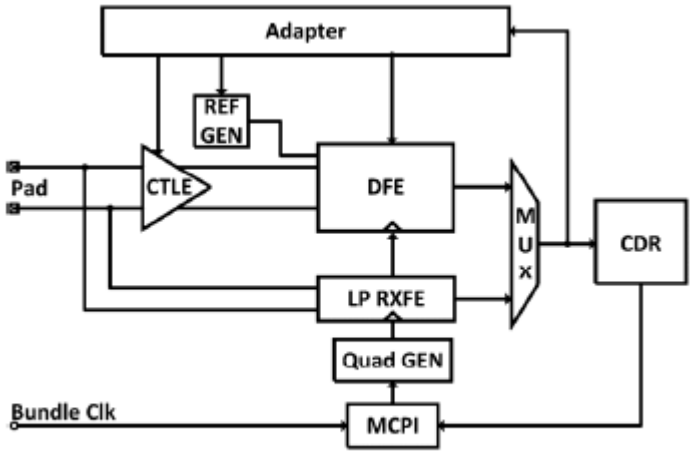
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

192

Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

193

Claim 1	Disclosure
194	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

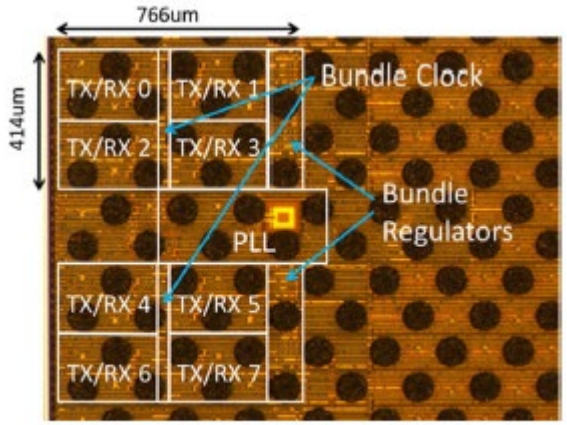
Claim 1	Disclosure
195	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 1	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

196

Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

197

Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

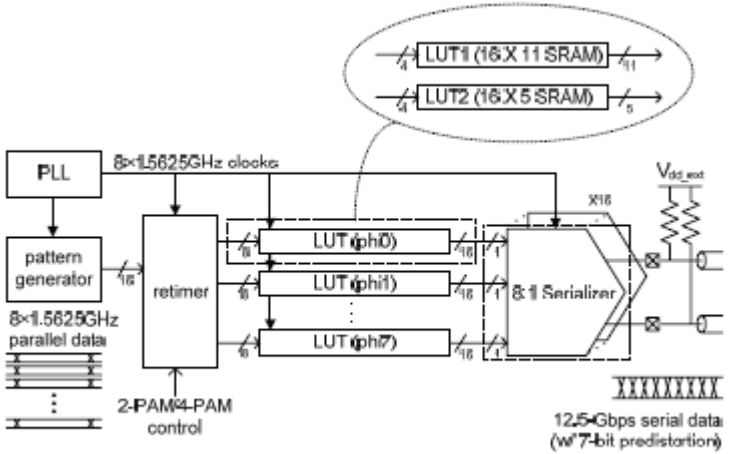
198

Claim 1	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

199

Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

200

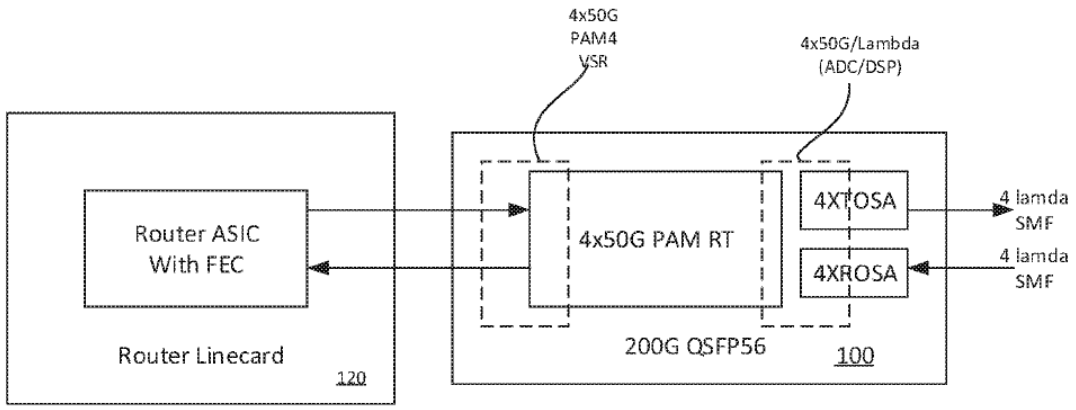
Claim 1	Disclosure
201	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

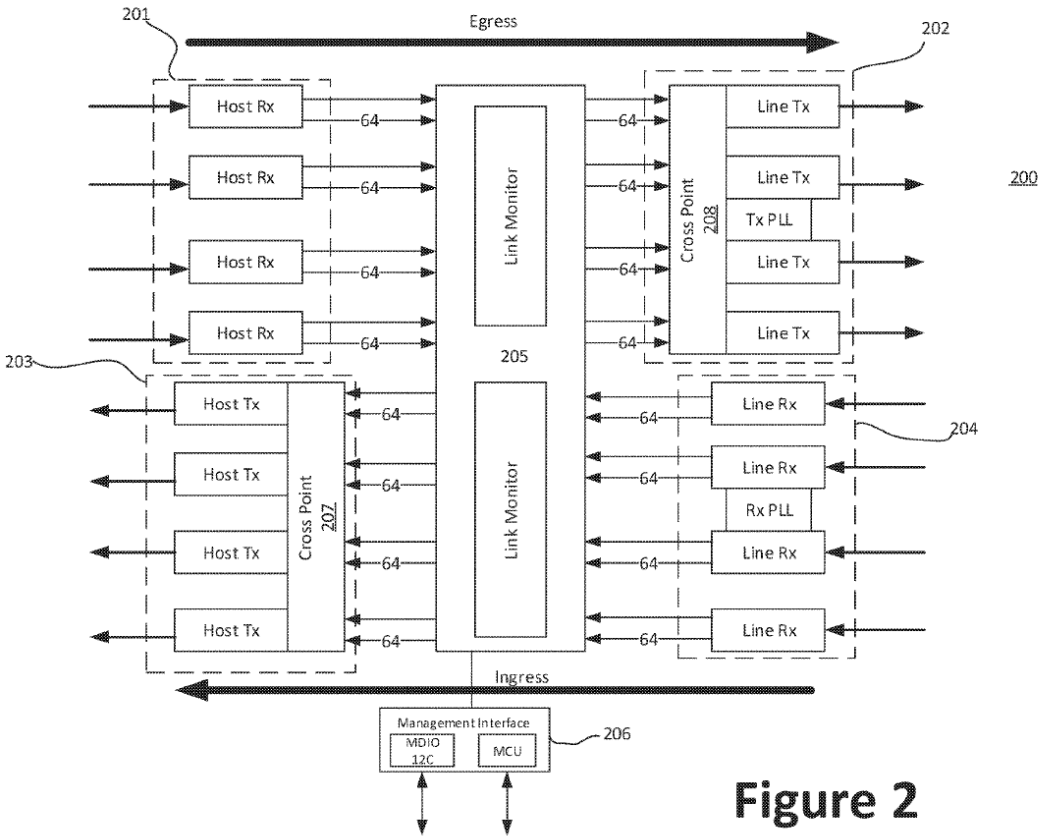
Claim 1	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

202

Claim 1	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

203

Claim 1	Disclosure
204	 **Figure 1** Riani, Figure 1.

Claim 1	Disclosure
205	 Figure 2 Riani, Figure 2.

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

206

Claim 1	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 1	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

208

Claim 1	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

209

Claim 1	Disclosure
210	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

Claim 1	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

211

Claim 1	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

212

Claim 1	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

213

Claim 1	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

214

Claim 1	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

215

Claim 1	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

Claim 1	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

Claim 1	Disclosure
218	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

219

Claim 1	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

220

Claim 1	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

Claim 1	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

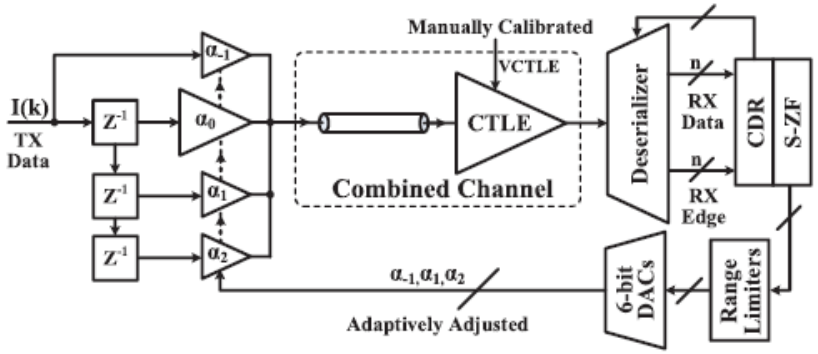


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

222

Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

223

Claim 1	Disclosure

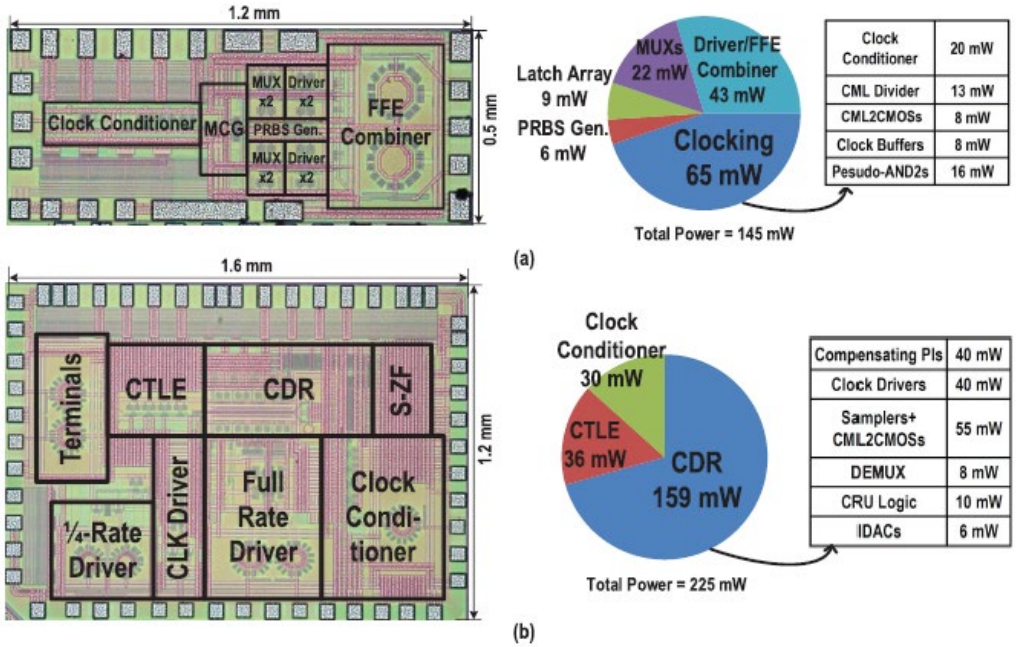


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

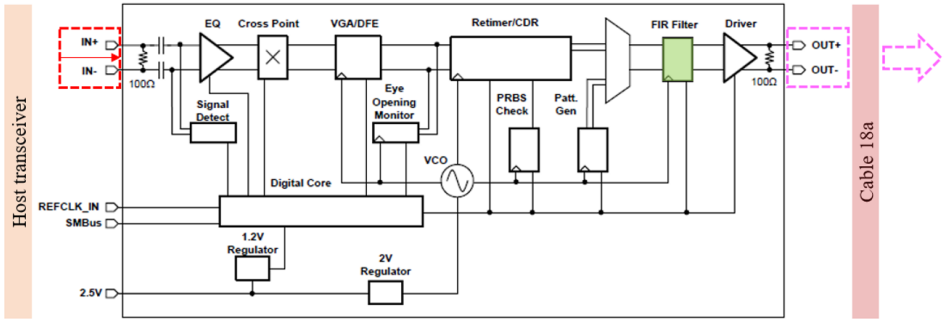
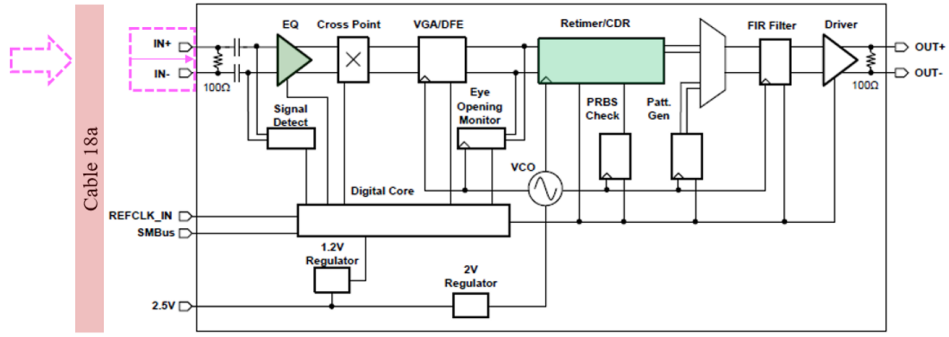
224

Claim 1	Disclosure
225	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The active cable 3 receives the serial bit stream transmitted from the upstream device 2 in the upstream facing RX 30a in one retimer 30. The retimer 30 performs a waveform shaping on the serial bit stream received by the upstream facing RX 30a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 30b to the other retime 31. The retimer 31 receives the serial bit stream transmitted from the retimer 30 in the upstream facing RX 31a. The retimer 31 performs a waveform shaping on the serial bit stream received by the upstream facing RX 31a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 31b to the downstream device 4." Kobayashi, 3:5-19. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX

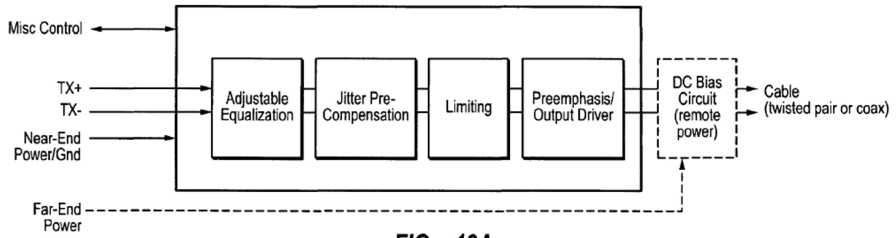
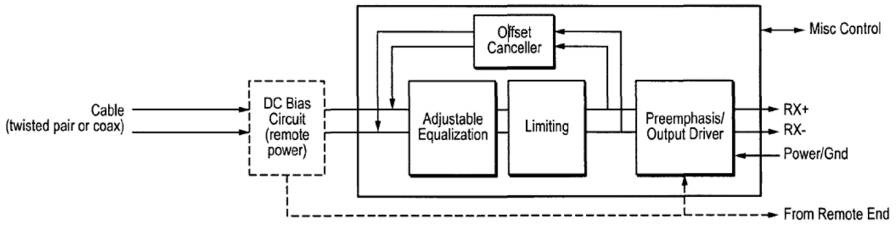
Claim 1	Disclosure
	serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver." Kobayashi, 4:4-9. "The exchange module 202 is configured to communicate capabilities information with a remote link partner including host loss information, RS-FEC ability information, and RS-FEC request information. The exchange module 202 may transmit local host loss information to the link partner and receive remote host loss information from the link partner. In some embodiments, the exchange module 202 may be included in or otherwise established by the autonegotiation module 104 and thus the capabilities information may be exchanged during execution of an autonegotiation protocol. In some embodiments, the exchange module 202 may be included in or otherwise established by the PMD/PMA module 106 and thus the capabilities information may be exchanged during a startup protocol performed by the PMD/PMA module 106." Lusted, paragraph 23. "The autonegotiation module 104 is configured to auto-negotiate line transmission speed, mode of operation, and other communication parameters with a link partner when the communication lane 102 is brought up. Additionally, in some embodiments the autonegotiation module 104 may be configured to exchange capabilities information with the link partner including host insertion loss information, RS-FEC ability information, and RS-FEC request information. The autonegotiation module 104 may encode the capabilities information, for example, in an autonegotiation base page or in a next page. Although illustrated in FIG. 1 as coupled between the PMD/PMA 106 and the communication lanes 102, in other embodiments the autonegotiation module 104 may be included at different positions in the Ethernet port logic 100. For example, in some embodiments the autonegotiation module 104 may be coupled between the PMD/PMA 106 and the FEC 108." Lusted, paragraph 17. "In view of the above transmission characteristics of twisted-pair, current Ethernet chips are usually equipped with preemphasis technique of a transmitting end and equalization technique of a receiving end to compensate the difference in transmission attenuation of physical layer signals on different frequencies on

226

Claim 1	Disclosure
	the same medium during Ethernet baseband transmission. Main functions of the preemphasis technique of the transmitting end and the equalization technique of the receiving end include: amplifying signals on different frequencies differently at the transmitting end to compensate different transmission attenuation in the twisted pair; applying the equalization technique to different frequency components in signals received at the receiving end to compensate the attenuation." Yu, paragraph 34. "An embodiment of the present invention provides an equalizer setting apparatus, which uses a filter circuit 101, an analog-to-digital converter 102, a digital memory 103, a digital comparator 104 and a first digital accumulator 105, the filter circuit 101, the analog-to-digital converter 102, the digital memory 103, the digital comparator 104 and the first digital accumulator 105 are connected in sequence, the output end of the equalizer 201 is connected with the input end of the filter circuit 101, the output end of the first digital accumulator 105 is connected with the set value input end of the equalizer 201, the output end of the analog-to-digital converter 102 is connected with the input end of the digital comparator 104, so that the problems that the setting value of the equalizer is accurately set in the prior art, which is time-consuming and labor-consuming are solved, the method has the advantages of achieving accurate setting of the set value of the equalizer in a single time, fast, simple, convenient, accurate and low-cost manner and saving labor force to a certain extent." Yin, paragraph 26 (English translation). *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See*

Claim 1	Disclosure
	id. at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."  Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram *U.S. Patent No. 7,762,727 to Aronson*

228

Claim 1	Disclosure
	Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id*. at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id*. at 14:16-15:24.  *FIG. 12A*  *FIG. 12B* Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23.

Claim 1	Disclosure
	Aronson discloses the host-side equalization and pre-equalization use fixed cable-independent parameters to address host-side losses. *See, e.g.*, Aronson at 12:10-17, 14:16-23, 15:17-24. "Fixed" parameters reduced complexity and manufacturing costs. See, e.g., U.S. Patent No. 7,239,665 at 1:7-18 (discussing slow convergence with adaptive equalization); U.S. Patent App. Pub. No. 2011/0228821 at [0009]-[0011] (discussing tradeoffs between fixed and adaptive equalization); U.S. Patent No. 6,621,862 at 1:52-55 (large gate count for adaptive equalization can cause "high chip costs"); U.S. Patent App. Pub. No. 2006/0045176 at [0034] ("The adaptive equalizer can be trapped in a failure state (i.e., incorrect state) and not able to recover from the failure state."); U.S. Patent App. Pub. No. 2017/0187463 at [0058] ("[E]qualization techniques… come with different engineering tradeoffs, including trade-offs in power consumption, performance, and/or cost."). **5. The PCIe 3.0 Equalization Adaptation Protocol** The PCIe 3.0 Base Specification [10] requires that a compliant Tx be able to produce any of the output equalization (TxEQ) levels shown in Figure 21, whenever requested to do so by the receiving port, during link equalization training. The same specification describes a protocol of TxEQ adaptation, where the host (upstream agent) or the end device (downstream agent) can each request that any of the TxEQ values in Figure 21 be presented to them by the other agent on the link. The intent of adaptation is to allow both agents to adjust the link partner's TxEQ to an optimal value for their receiver, and for the specific channel and operating conditions at hand. Samaan, p. 24.

230

B. DEPENDENT CLAIM 2

Claim 2	Disclosure
2[a] The active Ethernet cable of claim 1,wherein the respective transceivers each employ cable-dependent equalization parameters for at least one of: the remodulation of the inbound data stream for transit, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

231

Claim 2	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

232

Claim 2	Disclosure
233	Table 10- Control field dependency on Explicit Control bit Control field | Explicit Control value Tx Adaptive Input Eq Enable | 1 Tx Adaptive Input Eq Recall | 1 or 0 Tx Input Eq control | 1 Tx CDR control | 1 Rx CDR control | 1 Rx Output Eq control, pre-cursor | 1 Rx Output Eq control, post-cursor | 1 Rx Output Amplitude control | 1 CMIS, 38. **1.5.4.1 Control Set Usage** There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

Claim 2	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

234

Claim 2	Disclosure
235	**1.7 Module Memory Map** This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module. CMIS, 44. **1.5.5 Signal Integrity Control Field Background** The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated. The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11). Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable Control / Tx Adaptive Input Eq Enable value to use this control Tx Input Eq Adaptation Freeze — 1 Tx Input Eq Adaptation Store — 1 Tx Adaptive Input Eq Recall — 1 Tx Input Eq control — 0 The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module. CMIS, 40.

Claim 2	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

Claim 2	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

237

Claim 2	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

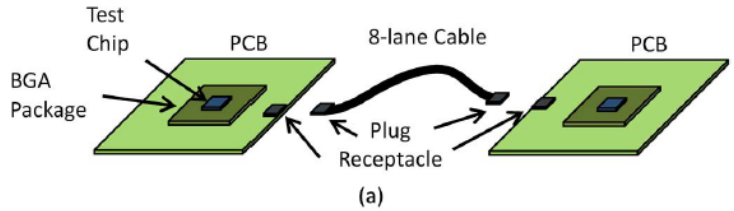
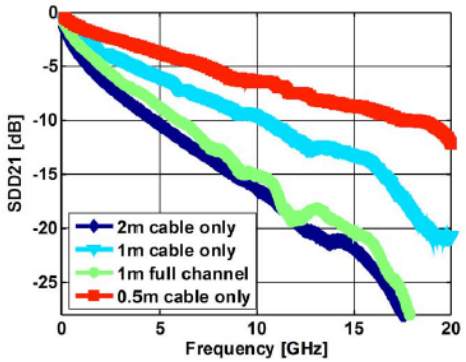
CMIS, 71.

Claim 2	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

239

Claim 2	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

240

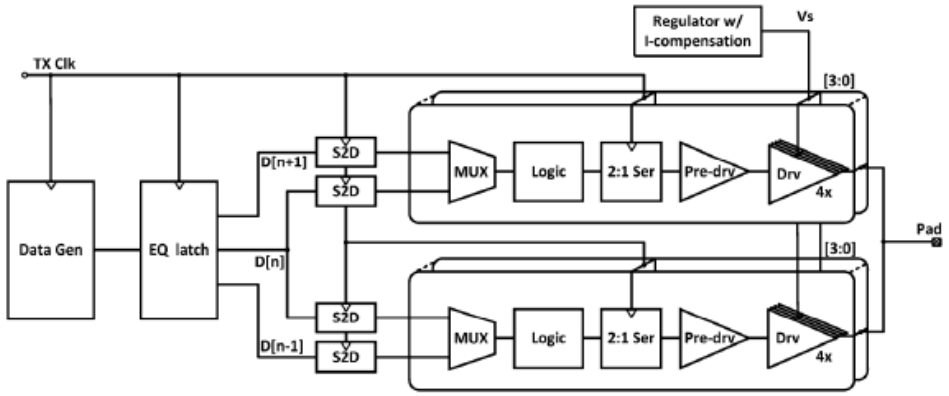
Claim 2	Disclosure
241	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

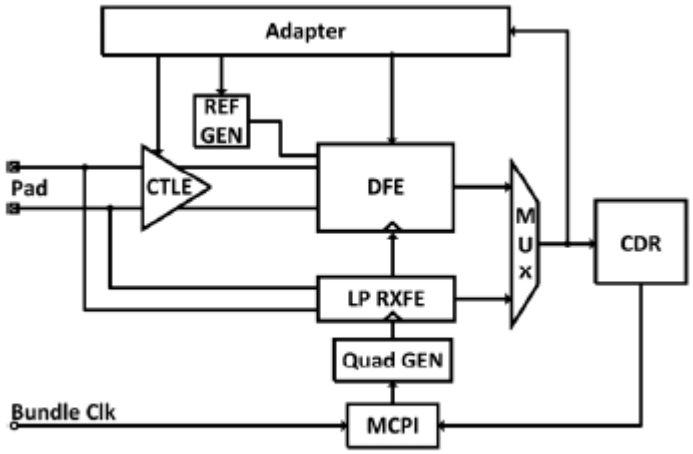
Claim 2	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

242

Claim 2	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

243

Claim 2	Disclosure
244	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

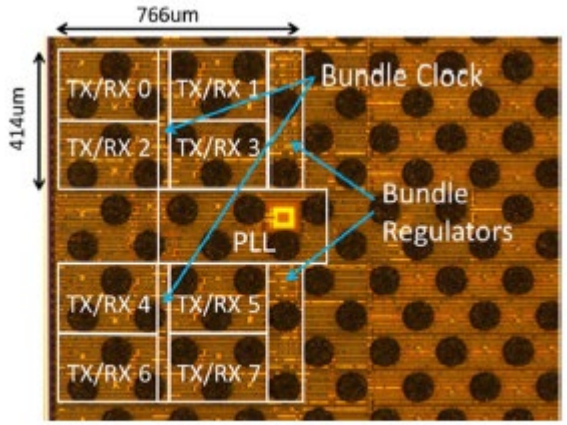
Claim 2	Disclosure
245	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 2	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

246

Claim 2	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

247

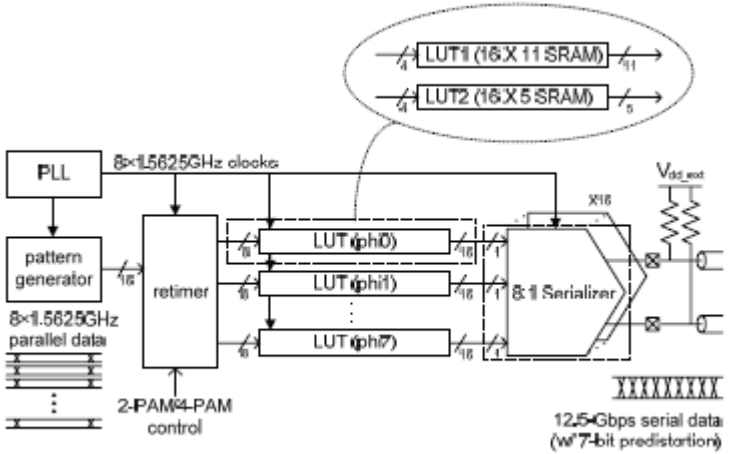
Claim 2	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 2	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

249

Claim 2	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

250

Claim 2	Disclosure
251	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 2	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

Claim 2	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

253

Claim 2	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

254

Claim 2	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

255

Claim 2	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 2	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

257

Claim 2	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

258

Claim 2	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

Claim 2	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

260

Claim 2	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 2	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

262

Claim 2	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

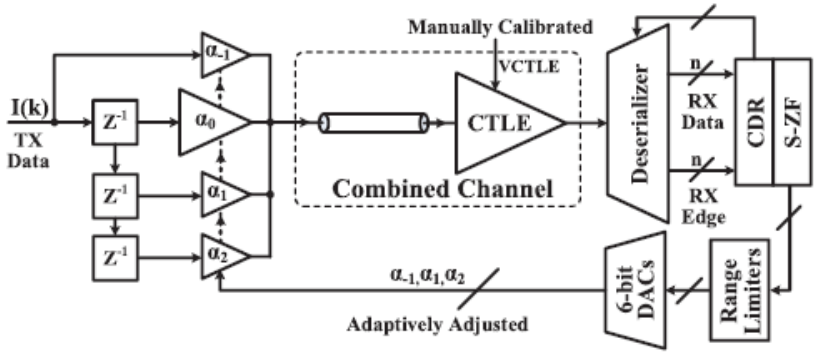


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

263

Claim 2	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

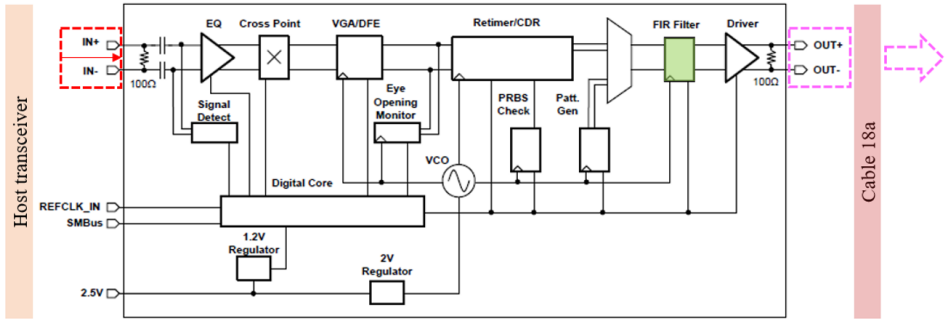
264

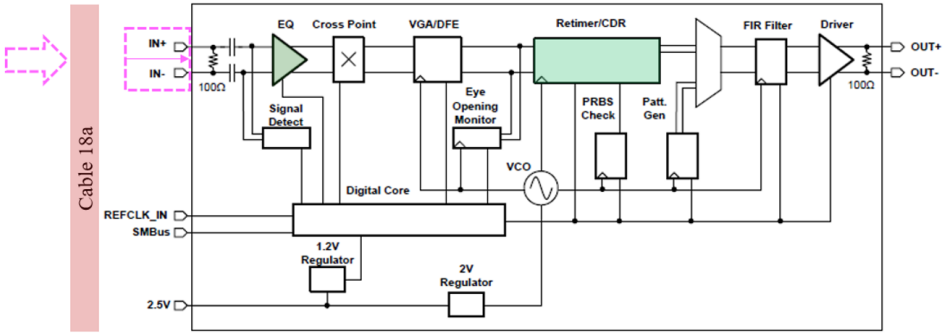
Claim 2	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The link activation module 204 is configured to determine a total insertion loss for the connection with the remote host based on the local host loss, the remote host loss, and cable loss. The link activation module 204 compares the total insertion loss to a FEC limit and to a specification limit. The link activation module 204 may bring up the communication link with the link partner without FEC enabled if the total insertion loss is below the specification limit and below the FEC limit. Additionally, the link activation module 204 may bring up the communication link with the link partner with FEC enabled if the total insertion loss is below the specification limit but not below the FEC limit. Further, the link activation module 204 may not bring up the communication link with the link partner if the total insertion loss is not below the specification limit (e.g., the link activation module 204 may keep the communication link down or bring the communication link down)." Lusted, paragraph 24.

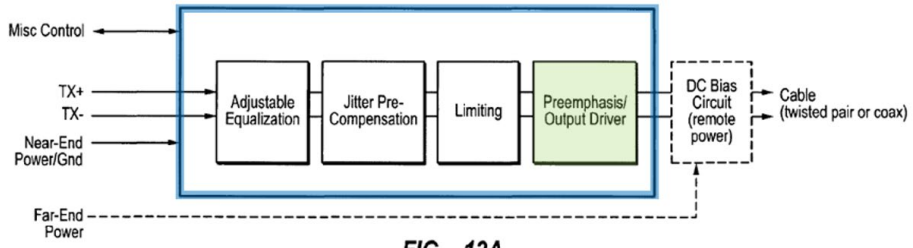
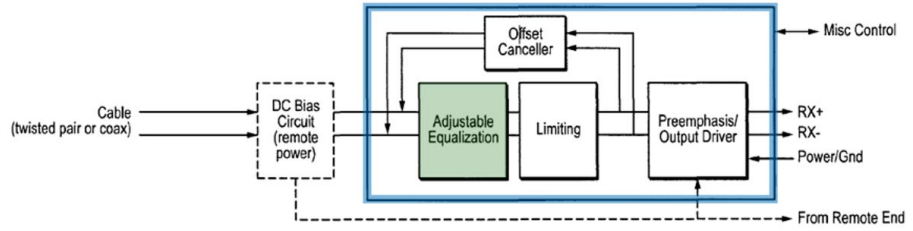
265

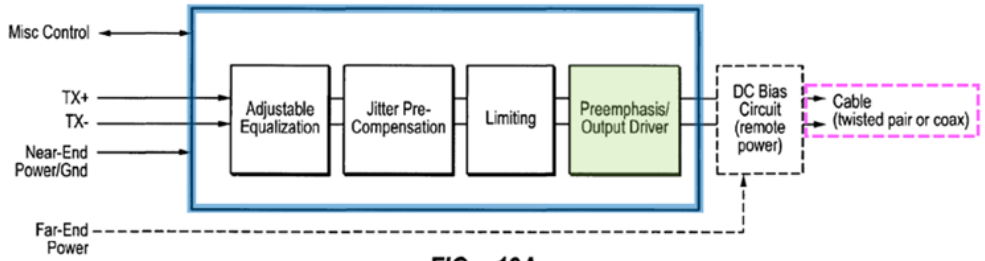
Claim 2	Disclosure
	"The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. "FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current

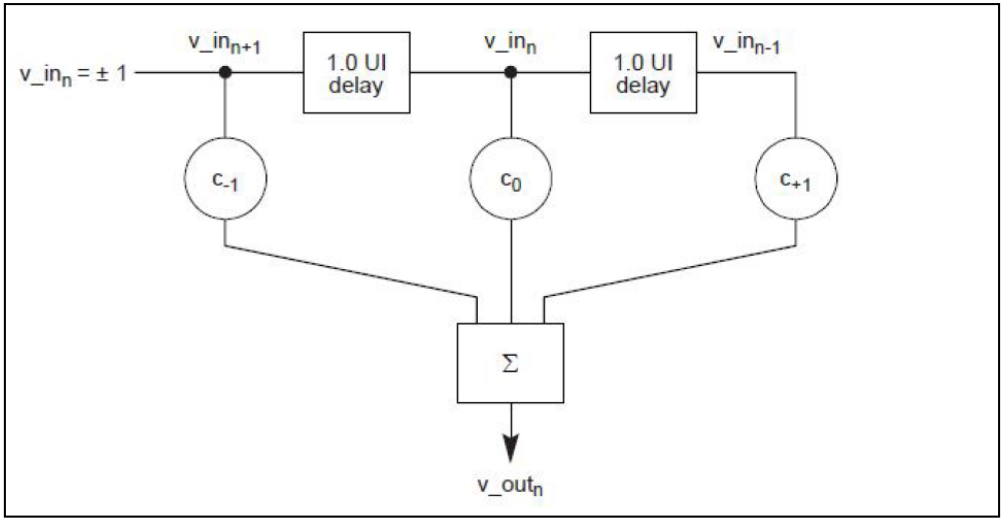
266

Claim 2	Disclosure
267	voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."  Figure 2. DS125DF1610 Simplified Data Path Diagram

Claim 2	Disclosure
268	 Figure 2. DS125DF1610 Simplified Data Path Diagram Aronson discloses cable-side equalization and pre-equalization that use cable-dependent parameters to address cable-side losses. *See, e.g.*, Aronson, 14:33-43, 15:1-8. Aronson's transmitter IC's line-side pre-emphasis/output driver overcomes "loss on the cable" using cable-dependent parameters that "match the particular length [and/or] the characteristics of the copper cabling." *See id*. at 14:33-43, Fig. 12A (below). Similarly, Aronson's receiver IC equalizes received line-side signals to "compensat[e] for the cable" losses using cable-dependent parameters based on "cable length and characteristics." *See id*. at 15:1-8, Fig. 12B (below).

Claim 2	Disclosure
	 FIG. 12A  FIG. 12B Using "cable dependent" parameters for equalization would have been the conventional and obvious way to address losses over the cable. *See, e.g.*, Aronson at 14:33-43 (pre-emphasis/output driver performs pre-equalization to overcome "loss on the cable" using cable-dependent parameters that "match the particular length [and/or] the characteristics of the copper cabling"), Fig. 12A (below).

Claim 2	Disclosure
	 *FIG. 12A*

Claim 2	Disclosure
	 **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.
2[b] the clock and data recovery performed on the electrical transit signal.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: `Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers.` CMIS, 1.

Claim 2	Disclosure
	1.4.4 MgmtInit State The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

272

Claim 2	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

273

Claim 2	Disclosure
	Table 10- Control field dependency on Explicit Control bit Control field: Explicit Control value Tx Adaptive Input Eq Enable: 1 Tx Adaptive Input Eq Recall: 1 or 0 Tx Input Eq control: 1 Tx CDR control: 1 Rx CDR control: 1 Rx Output Eq control, pre-cursor: 1 Rx Output Eq control, post-cursor: 1 Rx Output Amplitude control: 1 CMIS, 38. **1.5.4.1 Control Set Usage** There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

274

Claim 2	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

Claim 2	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

276

Claim 2	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

277

Claim 2	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

278

Claim 2	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	RQD
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

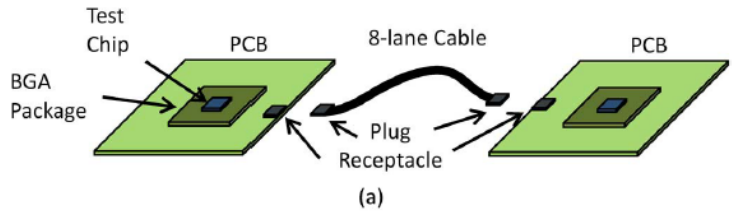
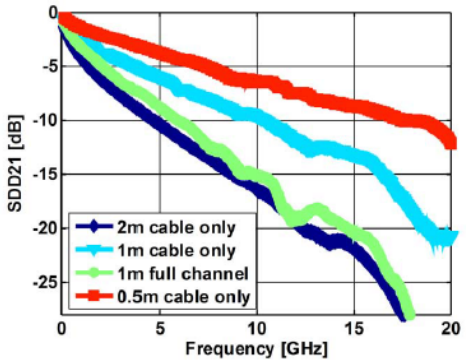
279

Claim 2	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

280

Claim 2	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

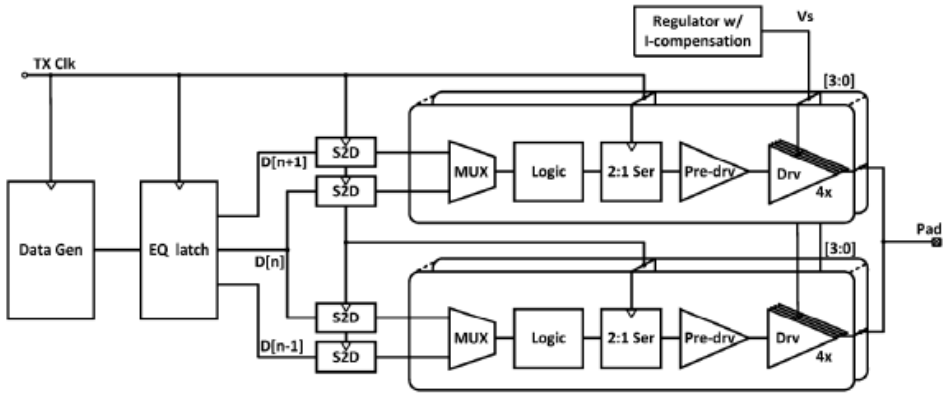
281

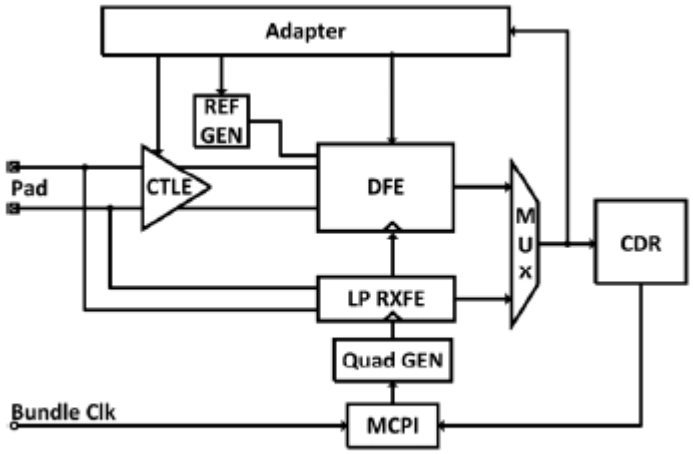
Claim 2	Disclosure
282	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 2	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

283

Claim 2	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

Claim 2	Disclosure
285	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 2	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

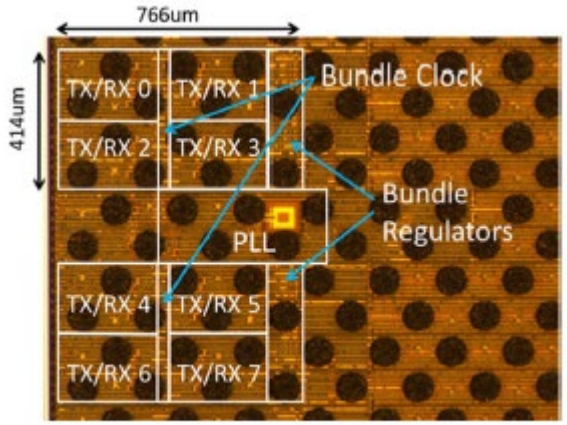
286

Claim 2	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

287

Claim 2	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

288

Claim 2	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

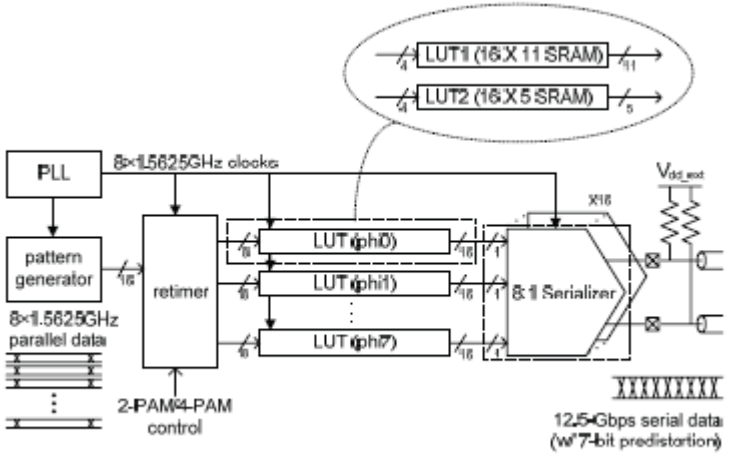
289

Claim 2	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

290

Claim 2	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

291

Claim 2	Disclosure
292	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 2	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

293

Claim 2	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

294

Claim 2	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

Claim 2	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

296

Claim 2	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 2	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

298

Claim 2	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

299

Claim 2	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

300

Claim 2	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

301

Claim 2	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

302

Claim 2	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

303

Claim 2	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

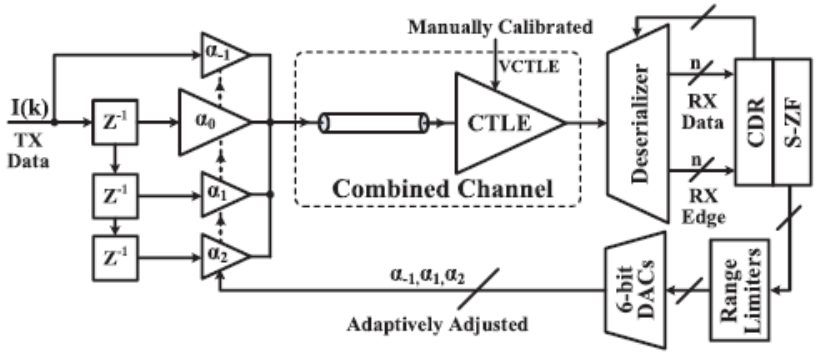


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

304

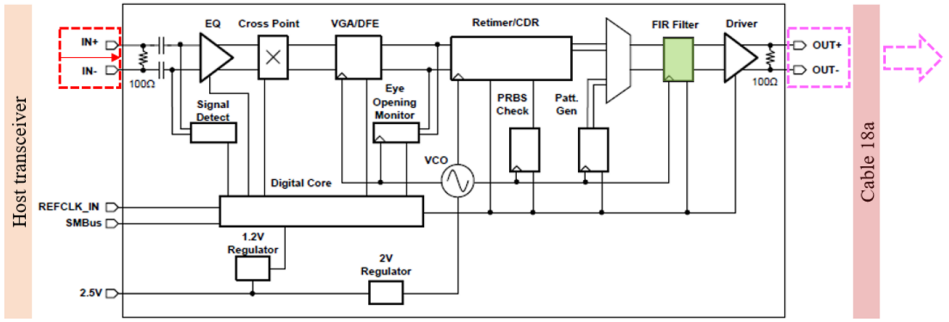
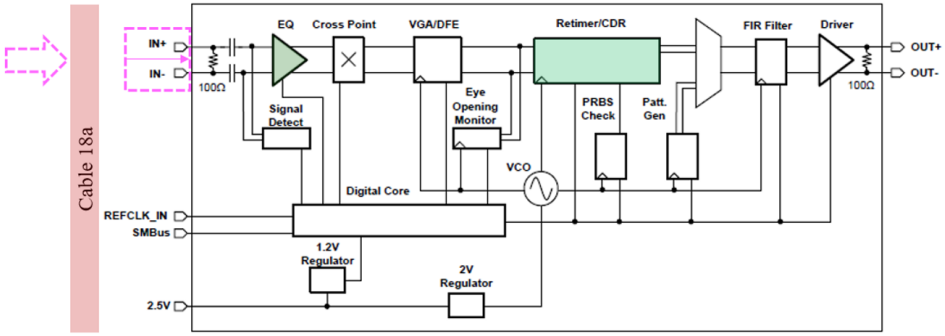
Claim 2	Disclosure
305	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

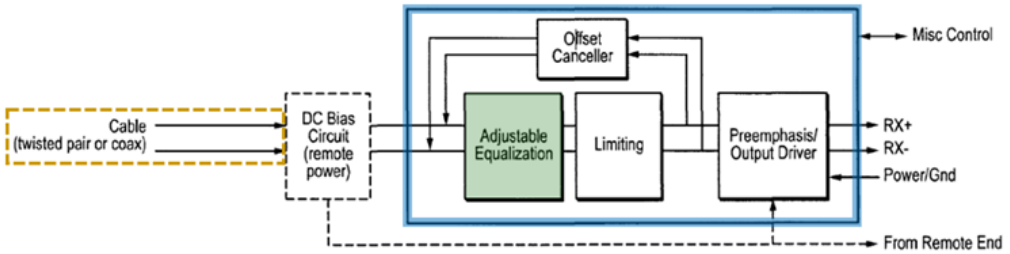
Claim 2	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20.

306

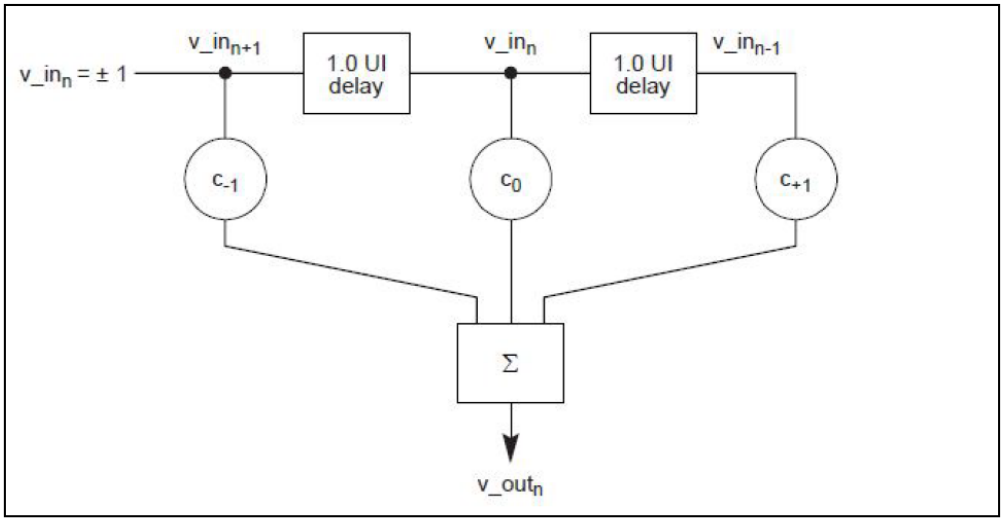
Claim 2	Disclosure
	"FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."

307

Claim 2	Disclosure
308	 Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram Using "cable dependent" parameters for equalization would have been the conventional and obvious way to address losses over the cable. Aronson at 15:1-8 (equalization "compensate[s] for the cable" losses using cable dependent parameters based on "cable length and characteristics"), Fig. 12B (below).

Claim 2	Disclosure
	 *FIG. 12B*

309

Claim 2	Disclosure
310	 **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.

C. DEPENDENT CLAIM 3

Claim 3	Disclosure
3. The active Ethernet cable of claim 2, wherein said cable-dependent equalization parameters adapt during usage of the Ethernet cable.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

311

Claim 3	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

312

Claim 3	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

313

Claim 3	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

Claim 3	Disclosure
315	**1.7 Module Memory Map** This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module. CMIS, 44. **1.5.5 Signal Integrity Control Field Background** The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated. The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11). Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable Control / Tx Adaptive Input Eq Enable value to use this control Tx Input Eq Adaptation Freeze — 1 Tx Input Eq Adaptation Store — 1 Tx Adaptive Input Eq Recall — 1 Tx Input Eq control — 0 The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module. CMIS, 40.

Claim 3	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

316

Claim 3	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

317

Claim 3	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

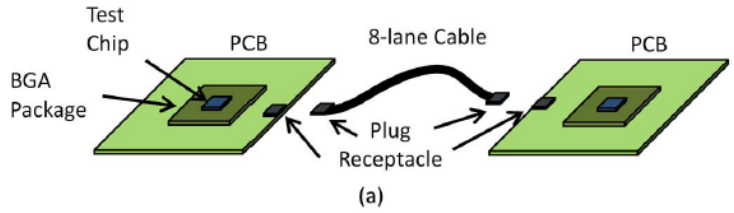
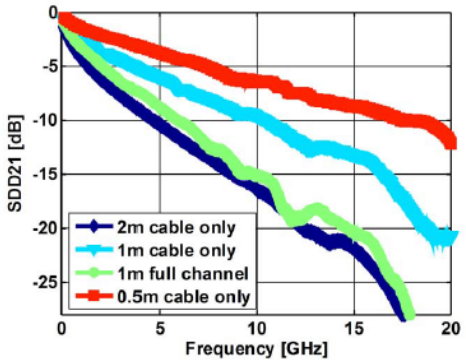
CMIS, 71.

Claim 3	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

319

Claim 3	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

320

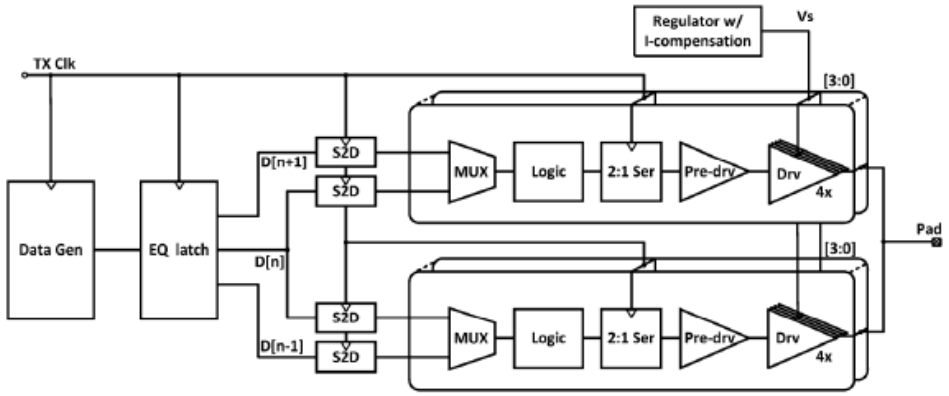
Claim 3	Disclosure
321	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

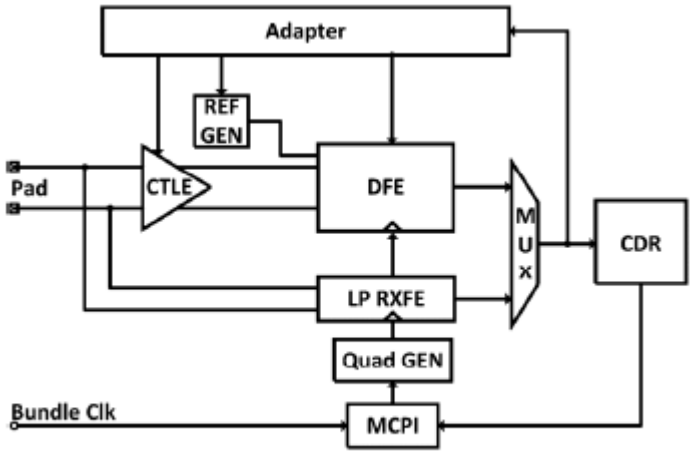
Claim 3	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

322

Claim 3	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

323

Claim 3	Disclosure
324	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

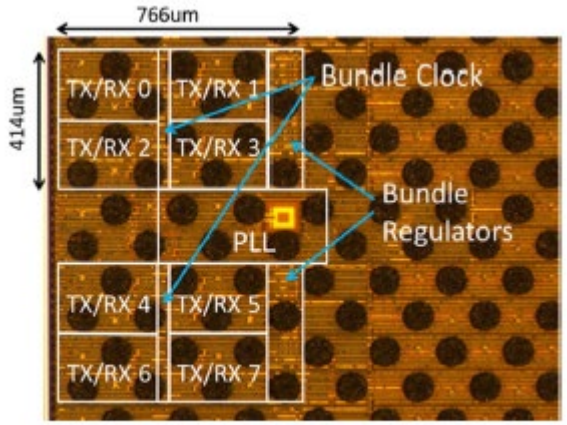
Claim 3	Disclosure
325	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 3	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

326

Claim 3	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

327

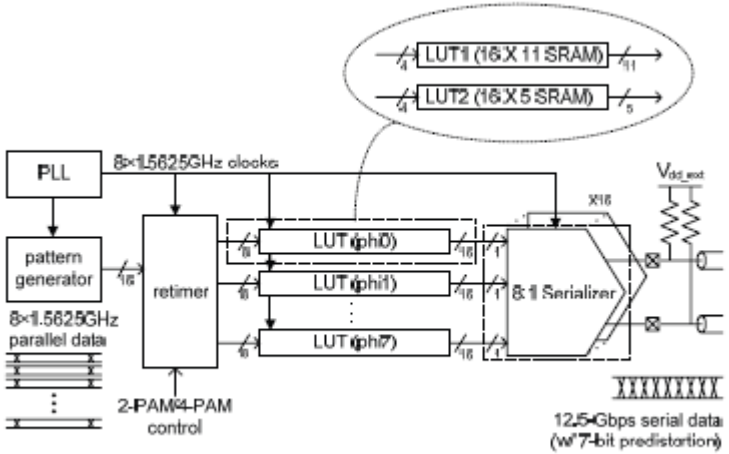
Claim 3	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 3	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

329

Claim 3	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

330

Claim 3	Disclosure
331	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 3	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

Claim 3	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

333

Claim 3	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

334

Claim 3	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

335

Claim 3	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 3	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

337

Claim 3	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

338

Claim 3	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

339

Claim 3	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

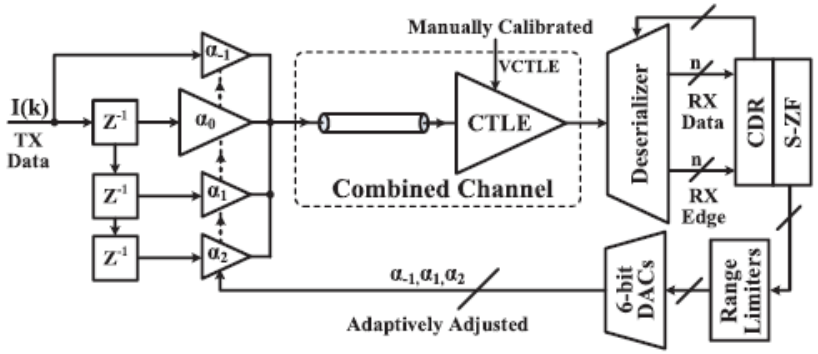
340

Claim 3	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

341

Claim 3	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

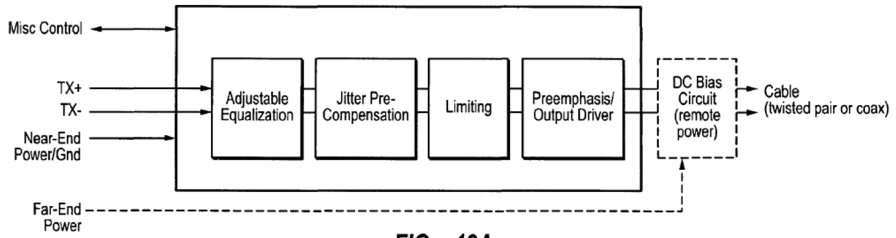
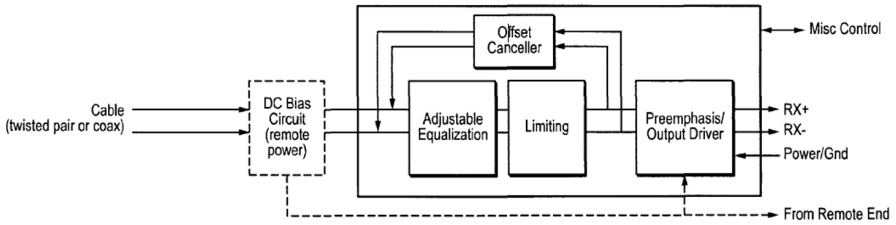
342

Claim 3	Disclosure
343	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

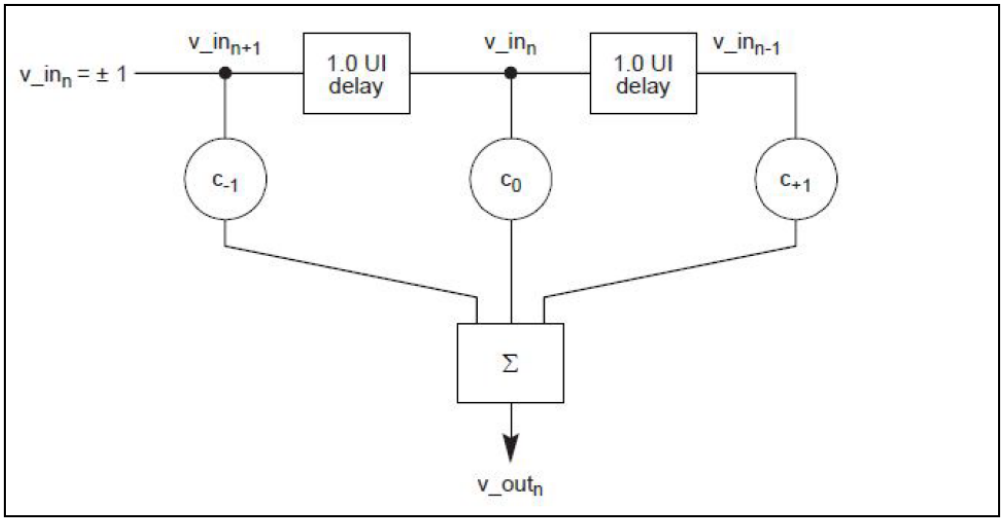
Claim 3	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

344

Claim 3	Disclosure
345	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. *U.S. Patent No. 7,762,727 to Aronson*

Claim 3	Disclosure
	Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id*. at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id*. at 14:16-15:24.  *FIG. 12A*  *FIG. 12B* Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23. Aronson discloses that cable-dependent equalization parameters used for cable-side equalization and pre-

346

Claim 3	Disclosure
347	equalization, respectively (*see, e.g.*, claim 2), are adapted during usage of the Ethernet cable. *See, e.g.*, Aronson at 15:1-8, Figs. 12A-12B.  **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.

D. DEPENDENT CLAIM 4

Claim 4	Disclosure
4[a] The active Ethernet cable of claim 2, wherein said cable-dependent equalization parameters are fixed during normal usage of the Ethernet cable, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

348

Claim 4	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

349

Claim 4	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

350

Claim 4	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

351

Claim 4	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

352

Claim 4	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

353

Claim 4	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

354

Claim 4	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

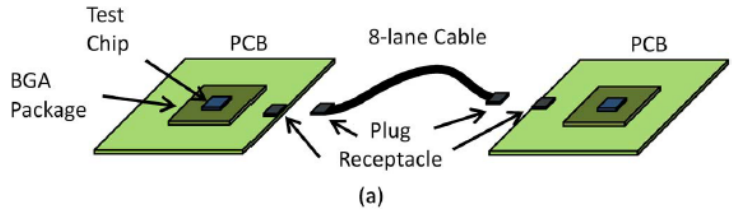
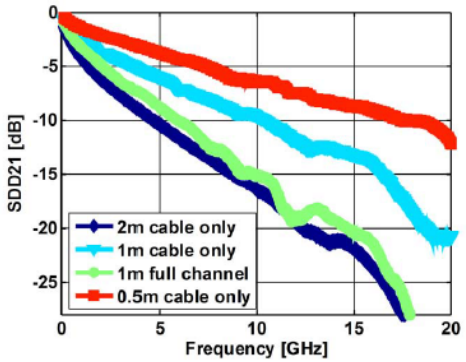
355

Claim 4	Disclosure
	Abstract—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

356

Claim 4	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

357

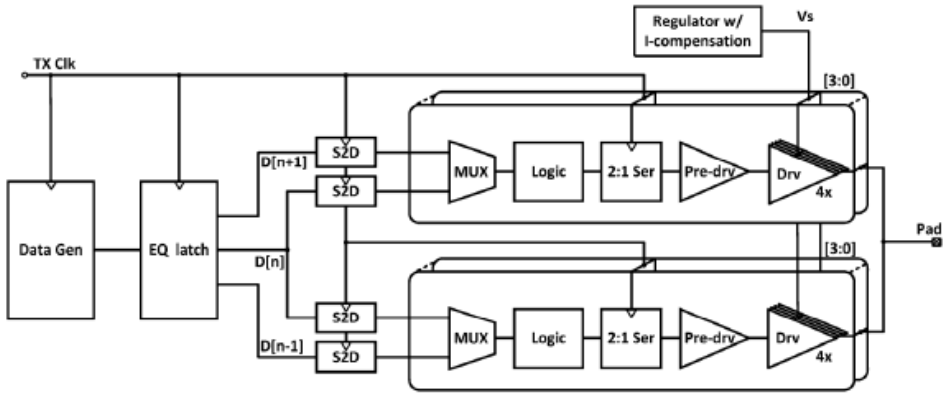
Claim 4	Disclosure
358	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

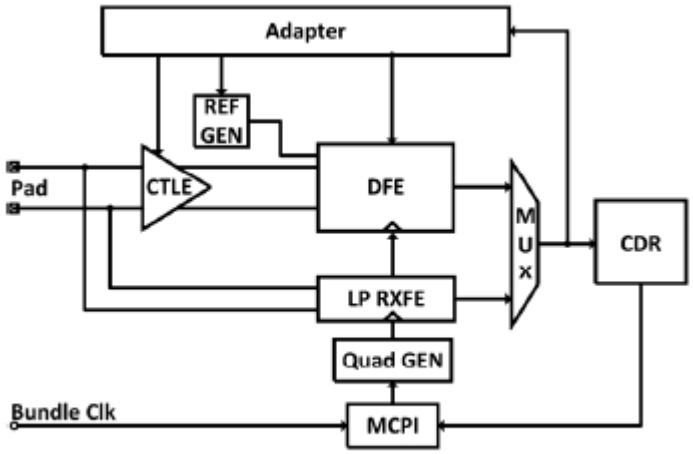
Claim 4	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

359

Claim 4	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

360

Claim 4	Disclosure
361	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

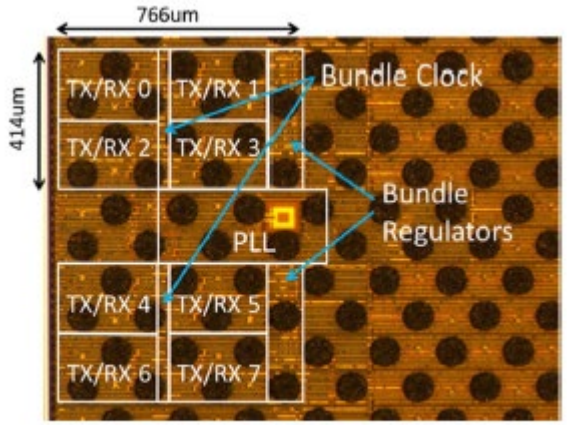
Claim 4	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 4	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

363

Claim 4	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

364

Claim 4	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

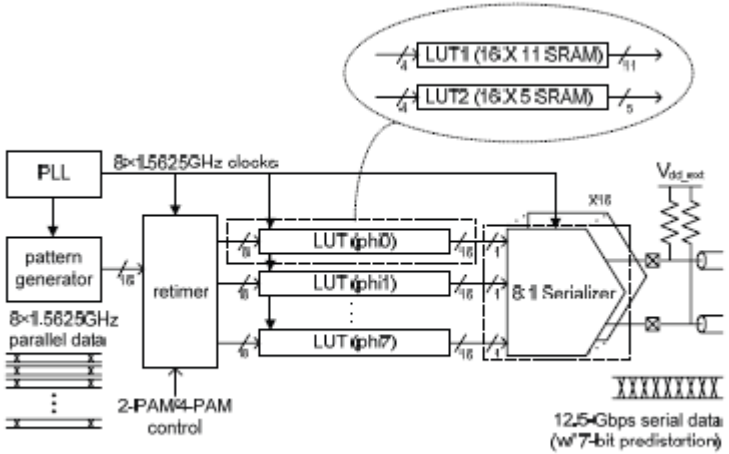
365

Claim 4	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

366

Claim 4	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

367

Claim 4	Disclosure
368	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 4	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

369

Claim 4	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

370

Claim 4	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

Claim 4	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

Claim 4	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 4	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

374

Claim 4	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

375

Claim 4	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

376

Claim 4	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

377

Claim 4	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 4	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

379

Claim 4	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

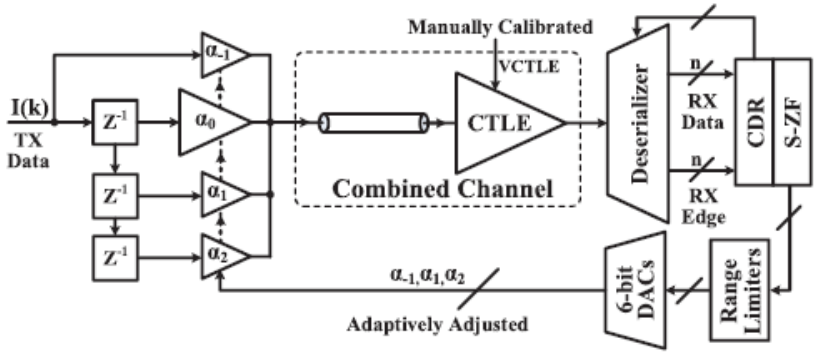


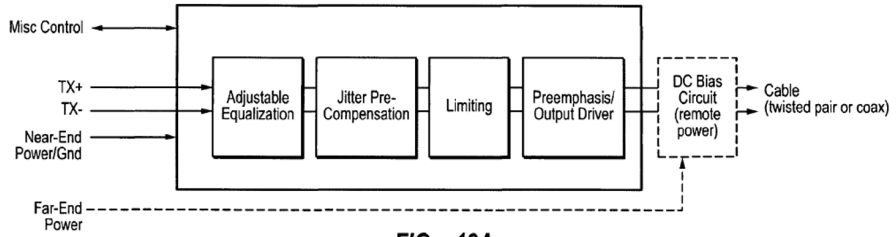
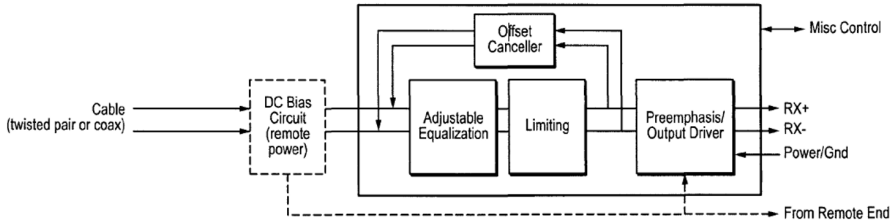
Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

380

Claim 4	Disclosure
381	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

Claim 4	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. *U.S. Patent No. 7,762,727 to Aronson* Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id*. at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id*. at 14:16-15:24.

382

Claim 4	Disclosure
	 FIG. 12A  FIG. 12B Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23. Aronson discloses that cable-dependent equalization parameters used for cable-side equalization and pre-equalization, respectively (*see, e.g.*, claim 2), are fixed during normal usage of the Ethernet cable. *See, e.g.*, Aronson at 14:35-39, 15:4-8, 15:21-23 (describing factory-set cable-dependent equalization parameters).

Claim 4	Disclosure
4[b] wherein said cable-dependent equalization parameters are determined during manufacturing-testing of the Ethernet cable.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

384

Claim 4	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

385

Claim 4	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

386

Claim 4	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

387

Claim 4	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

388

Claim 4	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

389

Claim 4	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

390

Claim 4	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

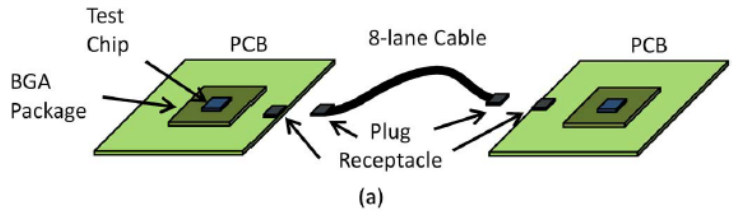
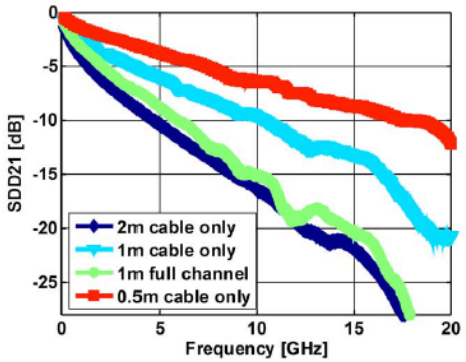
391

Claim 4	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

392

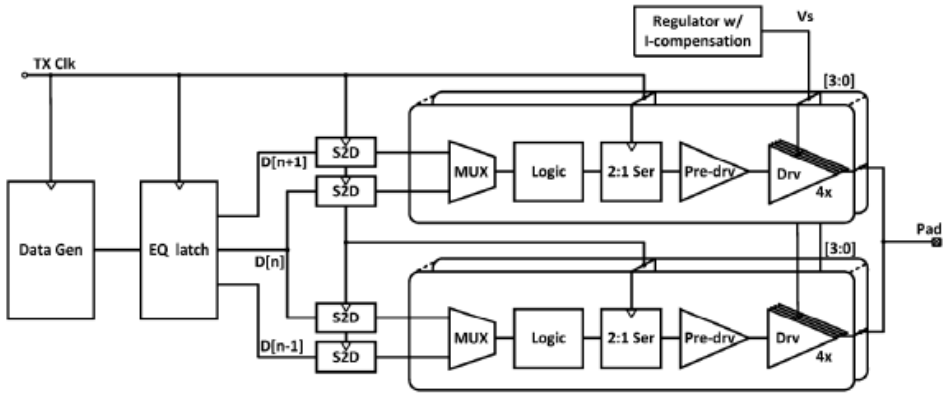
Claim 4	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

393

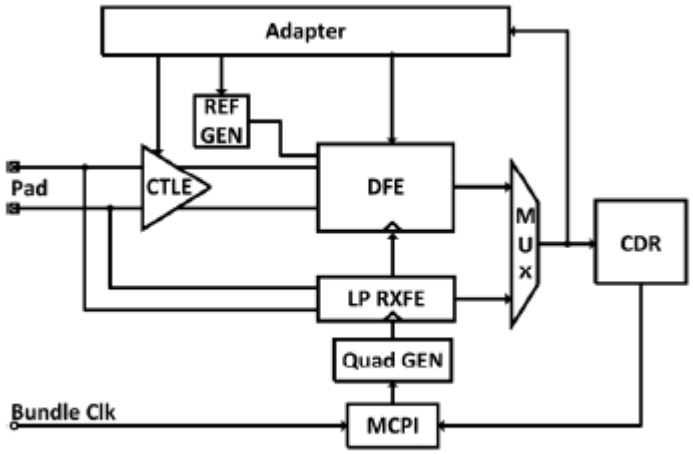
Claim 4	Disclosure
394	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 4	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 4	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

Claim 4	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

397

Claim 4	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

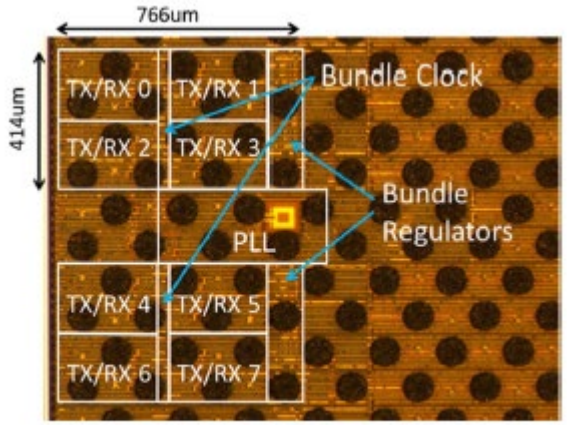
398

Claim 4	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

399

Claim 4	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

400

Claim 4	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

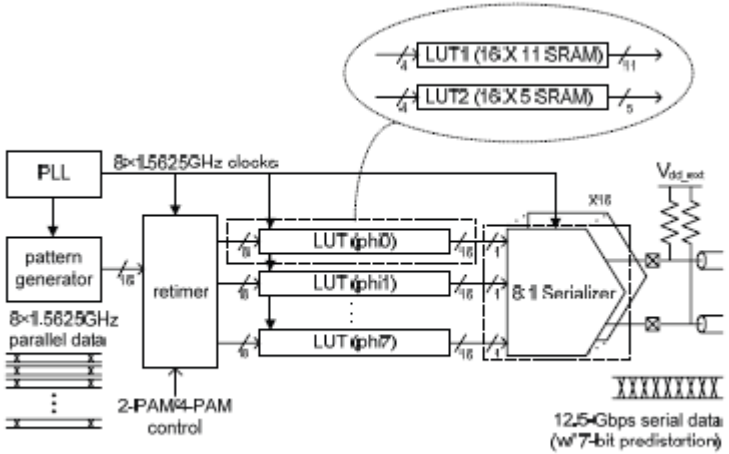
401

Claim 4	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

402

Claim 4	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

403

Claim 4	Disclosure
404	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 4	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

405

Claim 4	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

406

Claim 4	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

Claim 4	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

Claim 4	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

409

Claim 4	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

410

Claim 4	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

411

Claim 4	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/ LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

412

Claim 4	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

413

Claim 4	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

414

Claim 4	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

415

Claim 4	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

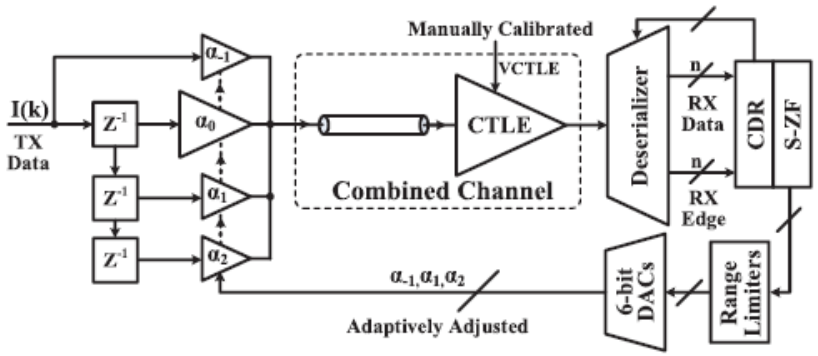
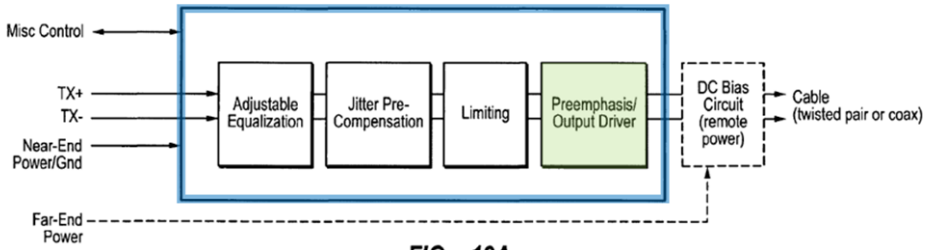
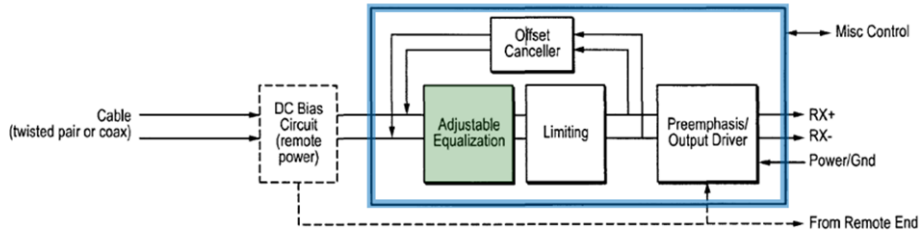


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 4	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

417

Claim 4	Disclosure
418	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55.

Claim 4	Disclosure
	U.S. Patent No. 7,762,727 to Aronson One obvious way to implement the equalization coefficients used by a FIR filter and an equalizer is during manufacturing-testing. *See, e.g.*, Aronson at 14:35-39 (Fig. 12A's pre-equalization parameters for transit over the line-side are "adjusted at factory setup"), 15:1-8 (Fig. 12B's equalization may be "fixed" or "adjustable but factory set" "at the time of the cable manufacture").  *FIG. 12A*  *FIG. 12B* Aronson discloses that cable-dependent equalization parameters used for cable-side equalization and pre-equalization, respectively (*see, e.g.*, claim 2), are fixed during normal usage of the Ethernet cable. *See, e.g.*, Aronson at 14:35-39, 15:4-8, 15:21-23 (describing factory-set cable-dependent equalization parameters).

419

E. DEPENDENT CLAIM 5

Claim 5	Disclosure
5. The active Ethernet cable of claim 4, wherein the inbound data stream and the outbound data stream each have a per-lane symbol rate in excess of 50 GBd.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** ***Index Terms*—Active inductor CTLE, bidirectional link, collaborative CDR, conditional phase detection, decision-feedback equalizer (DFE), phase error decimation, regulated CMOS clocking, source-series terminated (SST) driver.** Musah, 3079.

420

Claim 5	Disclosure
	II. LINK ARCHITECTURE The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

421

Claim 5	Disclosure
	IV. MEASUREMENTS Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. The measured TX output is shown in Fig. 16 at 32 Gb/s operation over a short channel. Fig. 17 shows the on-die 10^{-10} BER eye at 32 Gb/s over a 0.5 m cable using PRBS10 data on a single lane. Amortizing the PLL and clock distribution power over 4 lanes, the link power consumption is 205 mW at 32 Gb/s. Transceiver power efficiency given for multiple data rates from 4 to 32 Gb/s is shown in Fig. 18. Excluding the PLL power, the transceiver power efficiency ranges from 3.4 to 5.7 pJ/b, while equalizing 24 dB and 16 dB of Nyquist rate channel loss at 16 Gb/s and 32 Gb/s, respectively. Due to circuit BW limitations, the equalized channel loss tolerance is 24 dB up to 27 Gb/s and reduces to 16 dB at the highest data rate. When the HP RX is disabled and the LP RX is enabled, the measured power efficiencies are 1.0 pJ/b at 4 Gb/s, 1.5 pJ/b at 8 Gb/s and 2.2 pJ/b at 12 Gb/s with 4.5 dB, 8 dB and 10 dB of channel loss, respectively. A breakdown of the link power is shown in Table I, with PLL power amortized over 4 lanes. Musah, 3086.

422

Claim 5	Disclosure

TABLE II
COMPARISON TABLE

	[1]	[2]	[3]	[4]	This Work	
CMOS technology	32nm	32nm	28nm	40nm	22nm	
Supply voltage [V]	1.2/1.05	1.00	1.5/1.05	0.90	0.90	1.07
Data rate [Gb/s]	28	32	28	28	16	32
Channel loss at Nyquist [dB]	35	30	34	20	24	16
TX FFE	4-tap	4-tap	3-tap	3-tap	3-tap	3-tap
RX DFE	15-tap	15-tap	14-tap	N/A	6-tap	6-tap
RX peaking amplifier	2-stage	2-stage	1-stage	1-stage	1-stage	1-stage
TX output swing [$V_{pp\text{-}diff}$]	0.950			0.6	0.515	0.630
TX output jitter [ps_{rms}]	0.45		0.25	0.2		<0.54
RX input sensitivity [mV]	15		15	27		35
Area density [mm^2/Gb/s]	0.0289	0.0219	0.0298	0.0506	0.0049	0.0025
Transceiver power efficiency per lane [pJ/b]	24.8	21	20	6.96	4.6	6.4

Musah, 3087.

V. CONCLUSION

We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane.

Musah, 3088.

423

Claim 5	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G

424

Claim 5	Disclosure
	data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below. Riani, 4:43-5:5. ***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER < 10^{-12} over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

425

Claim 5	Disclosure
	A. Transmitter Chip Measurement The TX output is measured after a channel consisting of a doubled bonding wire, a 4-cm PCB trace, and a 0.5-m connection cable. Fig. 23(a) shows the over-equalized eye-diagram at 5 Gb/s, where the four sub-levels are contributed by the four FFE taps. Fig. 23(b) and (c) gives the output eye-diagrams at 40 Gb/s before and after applying the four-tap FFE. Obviously, the FFE can significantly improve the eye opening. The eye height and eye width are optimized from 140 mV and 0.45 UI to 180 mV and 0.68 UI, respectively. Meanwhile, the thickness of the eyelid is dramatically reduced from around 330 to 140 mV. Fig. 23(d) shows the properly compensated eye-diagram at the maximum operation speed of 50 Gb/s. Its eye height and eye width are 50 mV and 0.38 UI. Clearly, a wide operation range from 5 to 50 Gb/s is achieved, which is mainly attributed to the multi-MUX-based FFE implementation. Fig. 24 further shows the TX output with four separate eyes. It can be seen that the horizontal eye widths for both fixed clock and PRBS patterns are almost identical, thus proving that the four sampling phases are properly aligned. Zheng, 2972-2973. *B. Receiver Chip Measurement* The RX standalone measurement results are presented in this part. Fig. 25(a) shows the eye-diagram of the 40-Gb/s input data generated by an Anritsu MP1812A through combining four 10-Gb/s PRBS7 sequences, where the single-end eye height and eye width are around 360 mV and 0.71 UI. Fig. 25(b) shows the eye-diagram of the 10-Gb/s recovered

426

Claim 5	Disclosure
	data with a total jitter of 12.73 ps. The eye-diagrams of the recovered clocks (divided by 2) for the edge sampling and data sampling are shown in Fig. 25(c) and (d), which reveal that the introduced LPFs can optimize the total jitter from 11.48 to 7.66 ps. To demonstrate the effect of the LPFs with adaptively adjusting bandwidth, the jitter transfer (JTRAN) and JTOL curves are measured using a Tektronix BSA286C with a CDR block. The input peak-to-peak swing is tuned to 800 mV and the control voltage of the CTLE is manually set to 710 mV. The JTRAN curves in Fig. 26 illustrate that the bandwidth of the data-sampling path depending on the LPFs is 4 MHz, which is much smaller than 18 MHz for the edge-sampling path determined by the loop parameters. The measured JTOL in Fig. 26 indicates that the embedded LPFs result in a significant dip attenuation around the corner frequency and improve the JTOL amplitudes apparently at high jitter frequencies. Meanwhile, the adaptively adjusting bandwidth of the LPFs makes them exhibit little effect on the phase-tracking slew rate at low jitter frequencies. Additionally, the corner frequency of the JTOL is about 20 MHz, which is much larger than the JTRAN bandwidth of 4 MHz. Zheng, 2973-2974.

Claim 5	Disclosure

Transmitter Chip				
	[4]	[16]	[20]	This work
Technology (nm)	65	14	65	65
Supply (V)	1.2	N/A	1.2	1.2
Data Rate (Gb/s)	50-64	16-40	38.8-42	5-50
Area (mm^2)	1.2 × 1.0	0.215 × 0.13	0.9 × 0.7	1.2 × 0.5
TX Equalization	4-tap FFE	4-tap FFE	5-tap FFE	4-tap FFE
1UI-Delay Gen.	LC-Delay	Multi-MUX	CML-Delay	Multi-MUX
MUX Type	4:1	4:1	N/A	4:1
Data Jitter RJ (ps_{rms})	N/A	0.33@28Gb/s 0.51@40Gb/s	N/A	0.23@40Gb/s 0.18@50Gb/s
Data Jitter (ps) TJ (BER=1e-12)	N/A	10.72@28Gb/s 12.89@40Gb/s	15.1@40Gb/s	9.90@40Gb/s 10.58@50Gb/s
Power (mW)	199*	518	135**	145*
Energy Efficiency (pJ/bit)	3.1	12.9	3.4	3.6

*Exclude PLL power consumption, ** Standalone FFE without data serializer.

Zheng, 2974.

"FIG. 5 illustrates an EYE waveform of a signal obtained after passing a 56 Gbps PAM-4 signal through a raised cosine filter with a bandwidth of 22 GHz and a roll-off of 0.5. In FIG. 5, when the signal has a value higher than a determination threshold 501, the signal is determined to be +3, when the signal has a value between the determination threshold 501 and a determination threshold 502, the signal is determined to be +1, when the signal has a value between the determination threshold 502 and a determination threshold 503, the signal is determined to be −1, and when the signal has a value lower than the determination threshold 503, the signal is determined to be −3. Ranges above the determination thresholds 501, 502, and 503 through which the signal does not pass is referred to as the EYE widths. As seen in FIG. 5, the EYE width is defined by the upper EYE width 505 and the lower the EYE width 506, and the EYE widths 505 and 506 are defined by the data transition time, −3→+3 and +3→−3. The EYE width is narrowed by approximately 15 ps even under the condition that the bandwidth is limited by a frequency higher than the

428

Claim 5	Disclosure
	Nyquist frequency and there is almost no influence of ISI. In this manner, the EYE width of the PAM multilevel signal is narrowed relative to that of the NRZ signal due to data pattern dependence." Norimatsu, 4:27-50.

F. INDEPENDENT CLAIM 6

Claim 6	Disclosure
6[pre] A communication method that comprises, in a network cable having conductor pairs electrically connecting a first connector to a second connector	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: **1.1 Introduction** This specification defines the Common Management Interface Specification (CMIS) for QSFP-DD, OSFP and COBO in order to enable flexible use of modules and cable assemblies (hereafter referred to as modules unless cable assemblies are specifically mentioned) by the user. It shares some similarities with the management interfaces used in other form factors such as QSFP, SFP, and CFP. This specification supports a module with 8 host electrical lanes. Support for modules with more host lanes is possible (in increments of up to 8 lanes) by a bank selection feature (which allows several instances of an adressable page to exist in the paged register map). Due to this new mechanism this specification is not backwards compatible with SFF-8636. Some host ports may support both 4-channel and 8-channel devices; when a 4-channel device is inserted into such a port the appropriate memory map (e.g. SFF-8636) must be used. This case is outside the scope of this document. The host-to-module interactions needed to initialize active modules are presented in two state machines. A module-level state machine defines the initialization of the management interface as well as control of the module-wide power mode. A data path-specific state machine defines the host and module interactions and behaviors needed for power up and initialization of the data path(s) in the module. The CMIS is intended to cover a wide range of possible module functionalities and applications, ranging from cable assemblies to coherent DWDM modules. The CMIS is designed with the intent not to burden the simple cases but to provide a common set of management interactions and a common register map. In some modules, multiplexing of the high speed signals between the host electrical interface and the module media interface may occur in the module. In this specification, all references to lane numbers are based on the electrical connector interface lanes, unless otherwise indicated. In cases where a status or control aspect is applicable only to lanes after multiplexing or demultiplexing has occurred, the status or control is intended to apply to all lanes in the data path, unless otherwise indicated. Symmetry is expected between the Tx and the Rx hardware structure; for example Tx lanes that are multiplexed in the Tx are demultiplexed in the Rx. Each Application selected by the host is applied to the same Tx and Rx host electrical lanes. CMIS, 5-6.

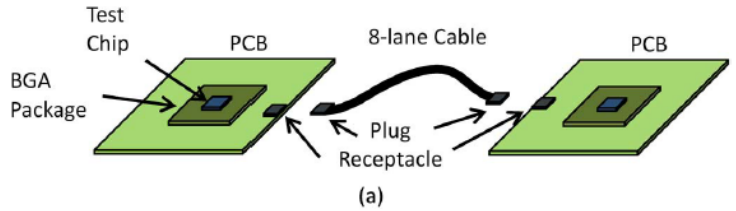
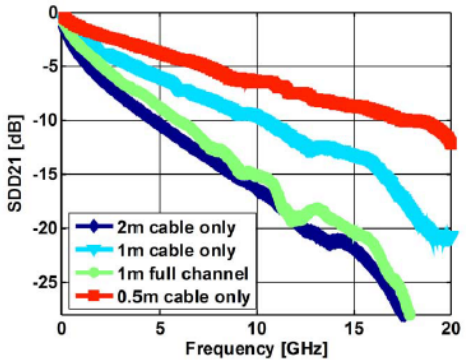
430

Claim 6	Disclosure
	Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. Table 1 and Table 2 provide a summary of the high-level behaviors and properties of each module state for passive copper cable assembly and active module implementations, respectively. Implementers of passive copper cable assemblies should note that this specification was designed to allow for flat-memory EEPROM implementations for that cable type. Active modules and active cable assemblies shall use the paged memory model. For all module types, refer to sections 1.4.3-1.4.9 for detailed requirements for each state. CMIS, 18.

Claim 6	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

433

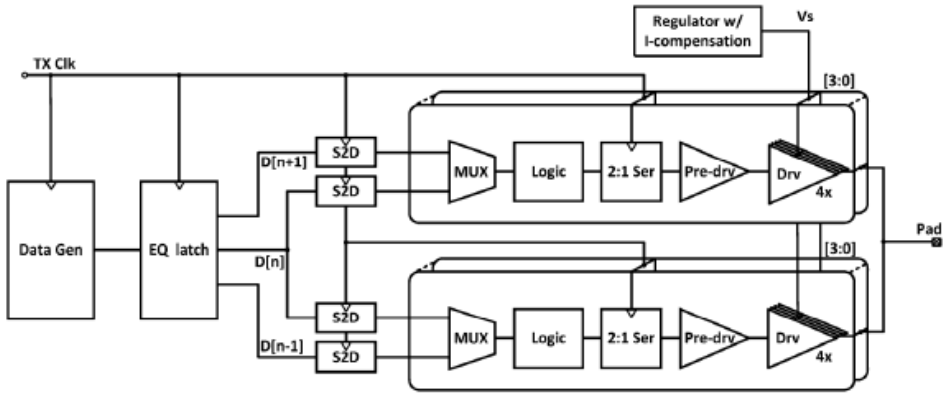
Claim 6	Disclosure
434	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

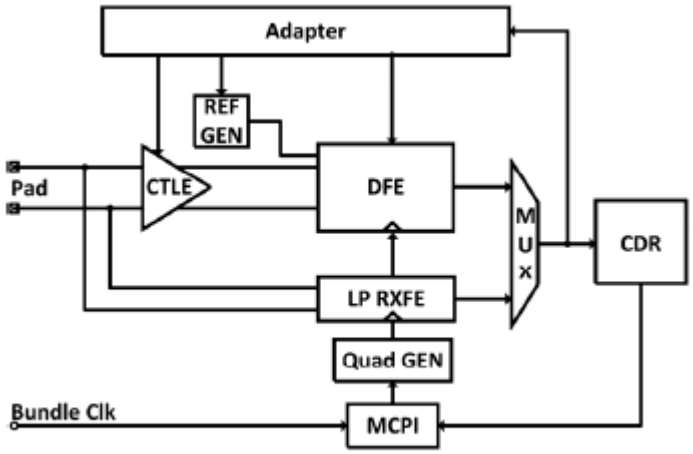
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

435

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

436

Claim 6	Disclosure
437	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

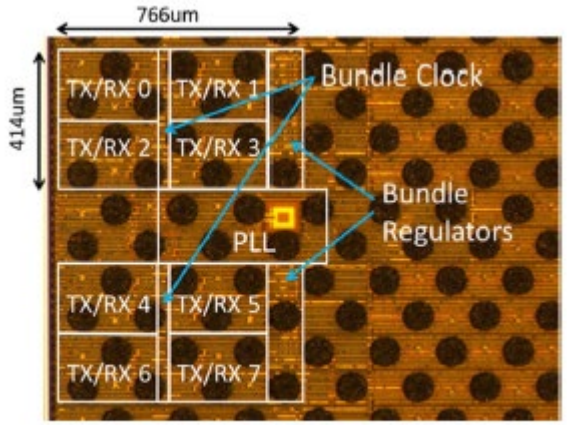
Claim 6	Disclosure
438	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

439

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

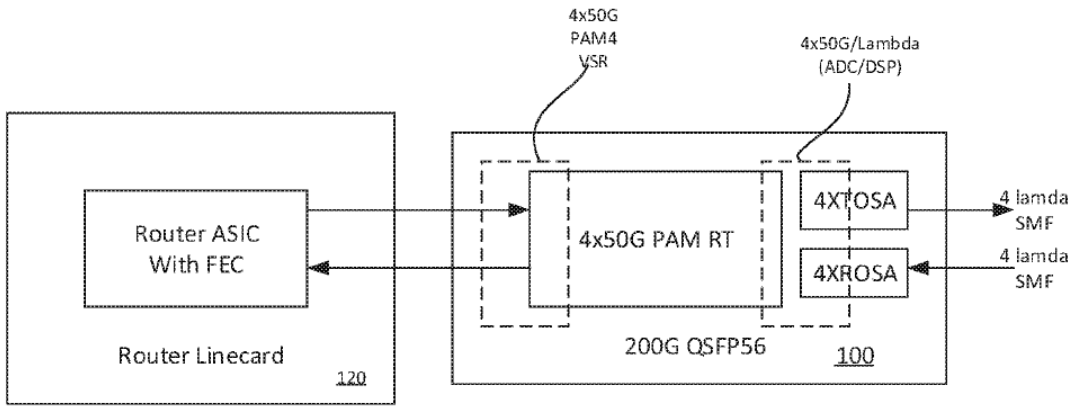
440

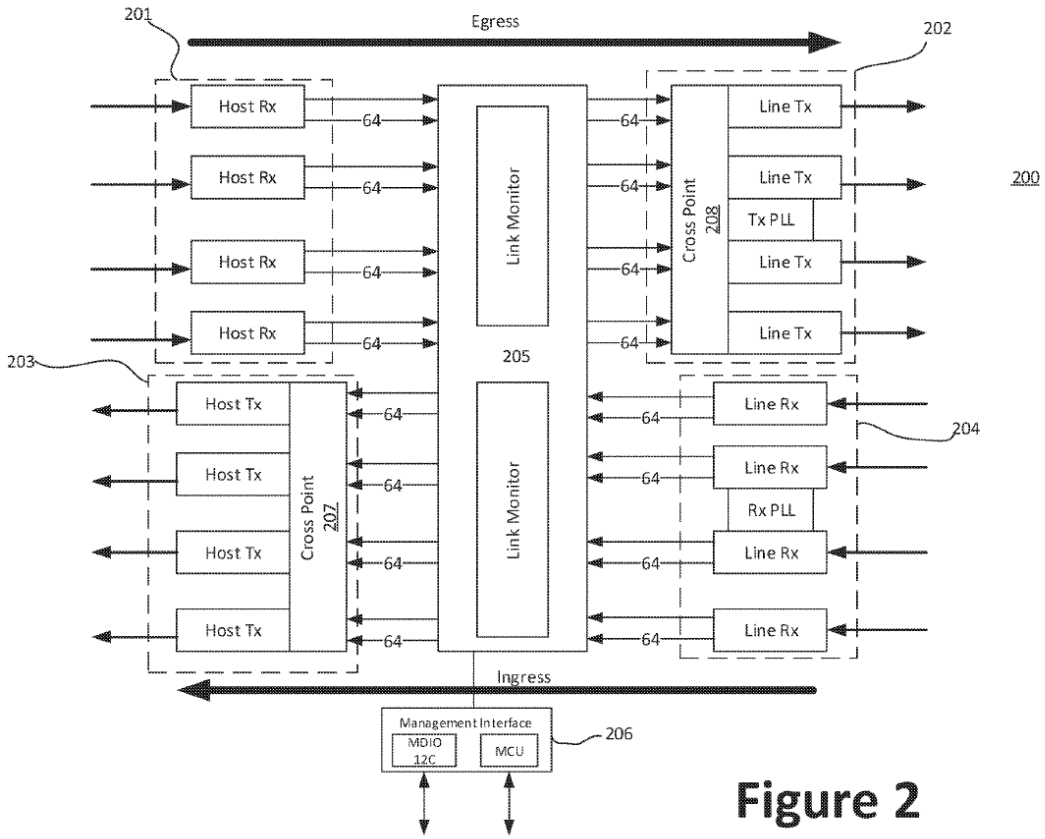
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

441

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free $(BER < 10^{-12})$ transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088.

442

Claim 6	Disclosure
443	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

445

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

447

Claim 6	Disclosure
	Riani, 4:43-5:5. ***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

Claim 6	Disclosure
	THE exponential growth of cloud computing, social networking, and multimedia sharing has led to an explosive bandwidth demand on data communication in both telecommunication equipment and inter/intra data center [1], [2]. To accommodate to this requirement, the data rate of the wireline serializer/deserializer (SerDes) transceiver has been continuously increased [3]–[5]. Currently, 25–28 Gb/s serial links approved by InfiniBand EDR, 32GFC, and CEI-28G have stepped into the period of industrial deployment [1], [3], [6]. The 38–64 Gb/s transceivers, which will play a key role in the next-generation data rate supported by Ethernet 400GbE, InfiniBand HDR, and CEI-56G, have attracted increasing attentions in the industry and the academia [2], [4], [5], [7]–[11]. The main challenges in designing such high-speed transceivers originate from the ever decreased unit interval (UI) period, which not only poses high bandwidth requests on the blocks located at the critical path, but also makes the link timing budget extremely tight. Moreover, advanced processes cannot completely solve these problems, since the parasitic capacitances/resistances at the high-speed outputs usually do not scale well with the technology due to the bonding and/or electro-static discharge (ESD) protection requirements. Zheng, 2963.

449

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

450

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

451

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

453

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

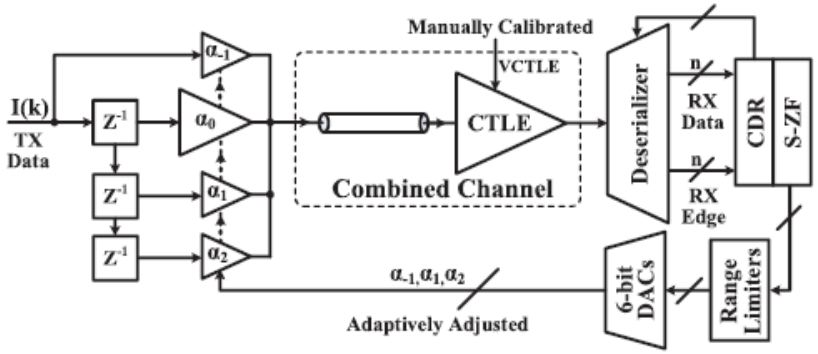


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

454

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

455

Claim 6	Disclosure

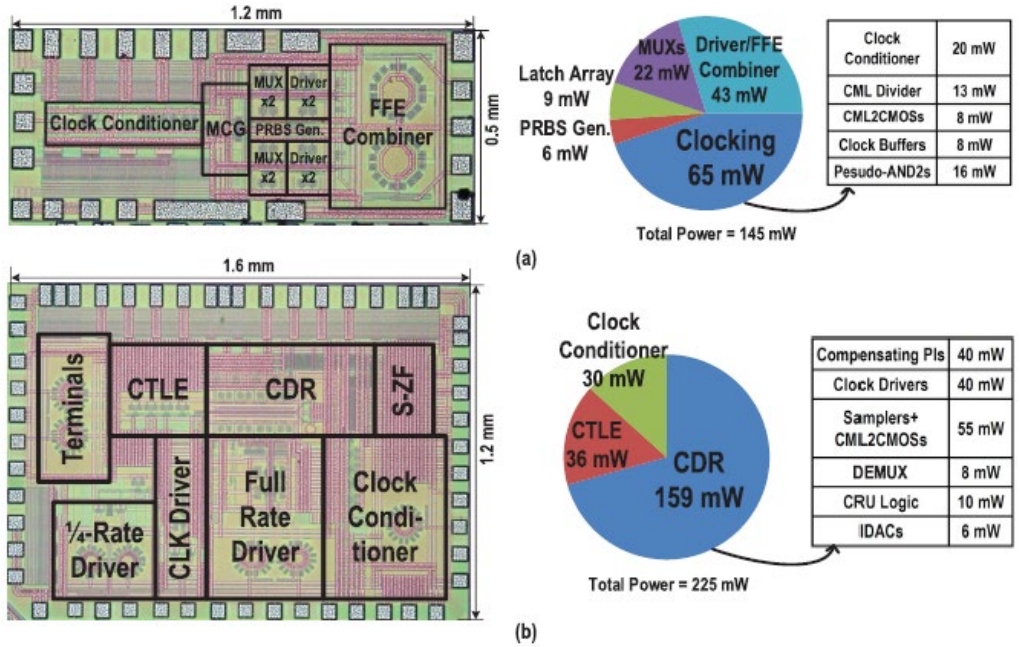


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

456

Claim 6	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 1, the data transmission system 1 is in accordance with DP1.4, for example, and includes an upstream device 2, a downstream device 4, and an active cable 3 connecting the upstream device 2 and downstream device 4." Kobayashi, 2:41-45.
6[a] receiving with the first connector a first electrical input signal conveying a first inbound data stream from a first host device;	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

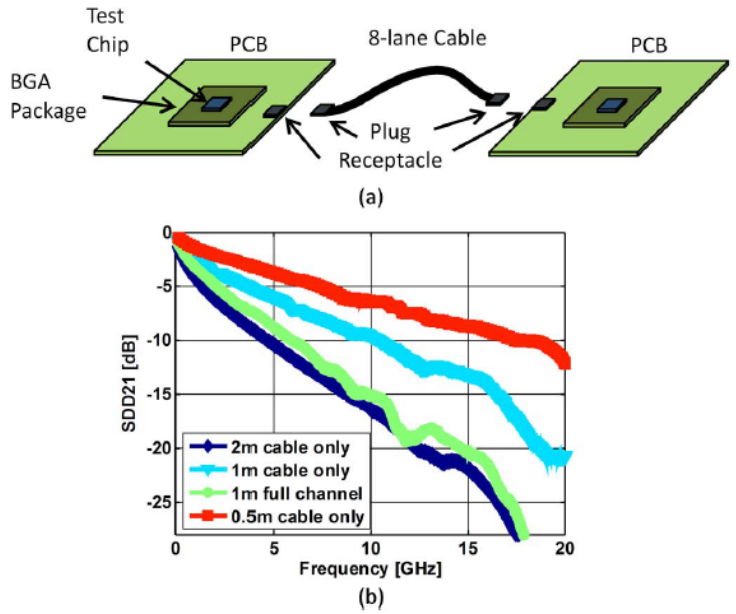
457

Claim 6	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

458

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

459

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

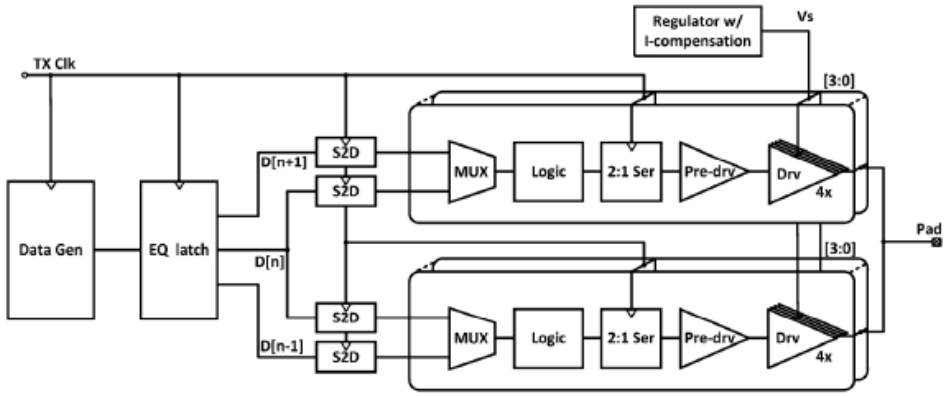
460

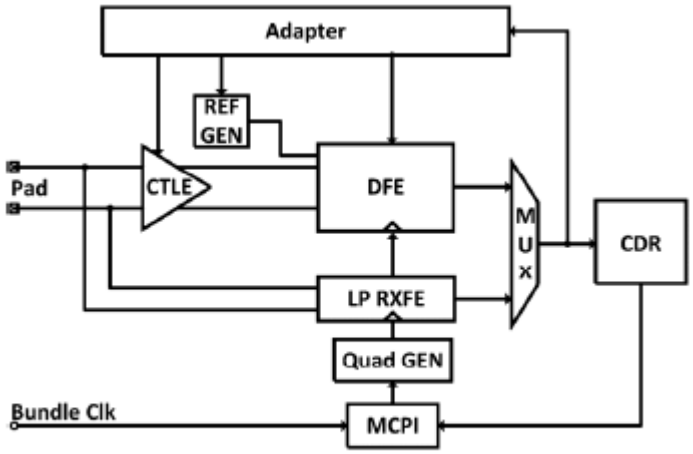
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

461

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

462

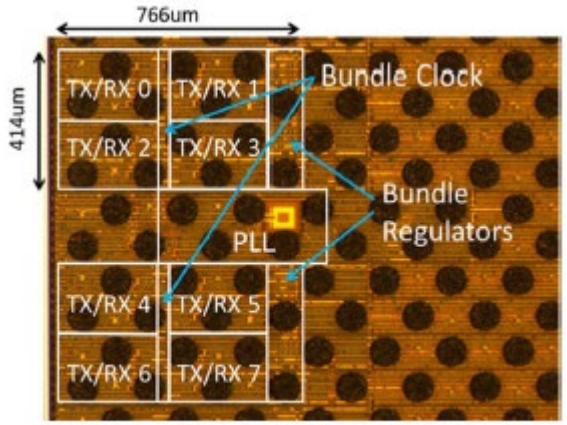
Claim 6	Disclosure
463	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

466

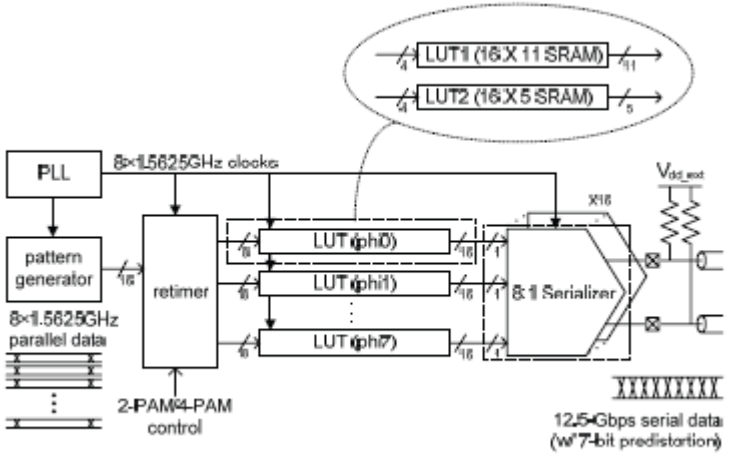
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

468

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

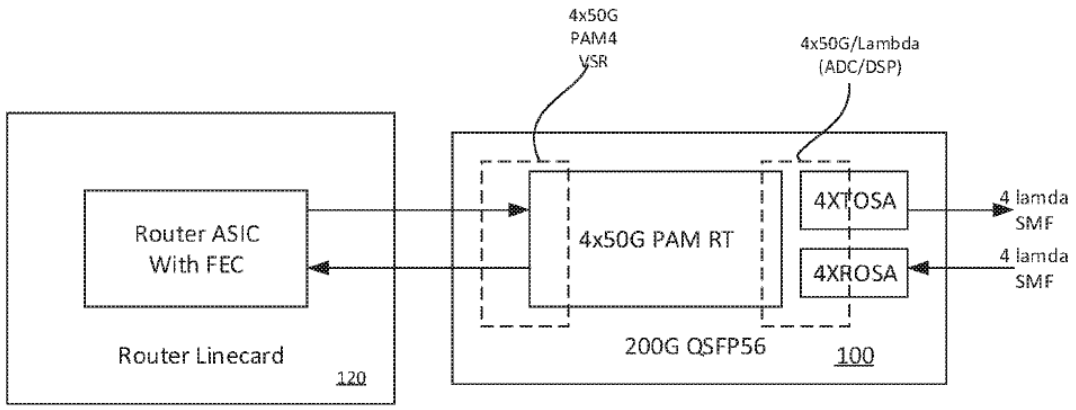
469

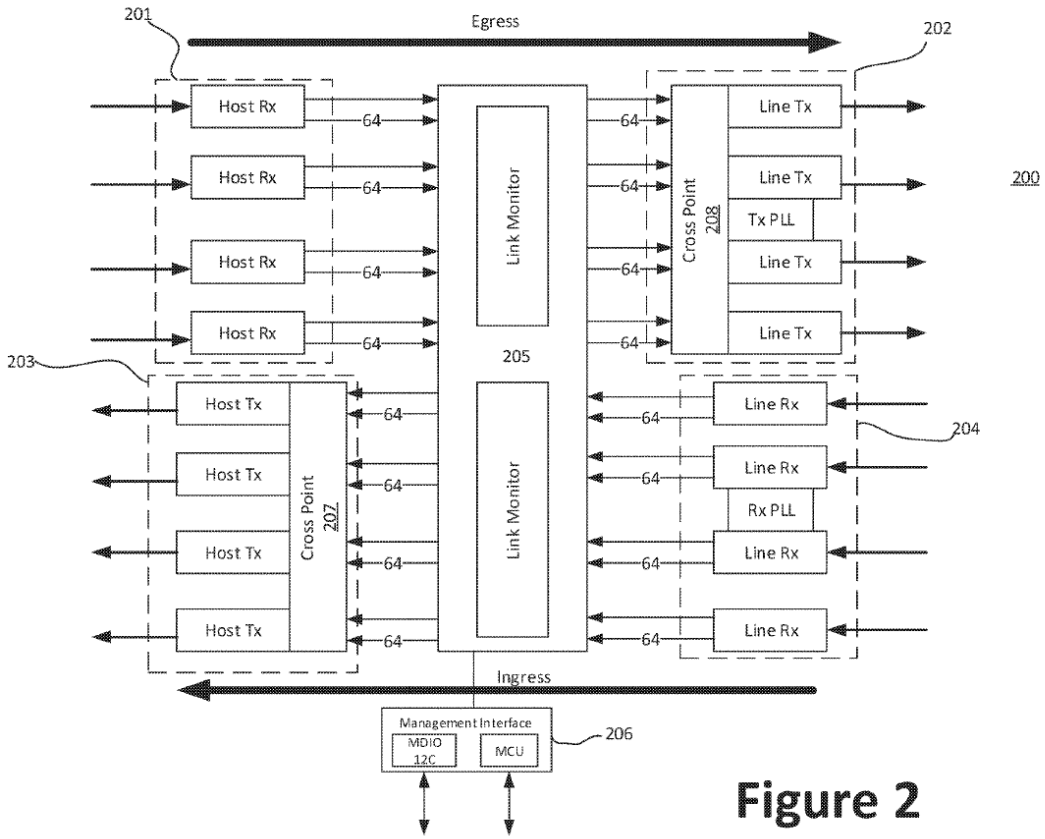
Claim 6	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

470

Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

Claim 6	Disclosure
	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

475

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

477

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

478

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

479

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

480

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

481

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

482

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

483

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

484

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

485

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

487

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

488

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

490

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

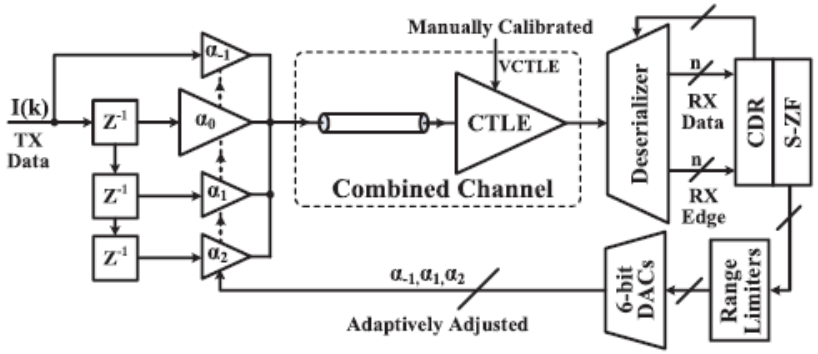


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

491

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

492

Claim 6	Disclosure

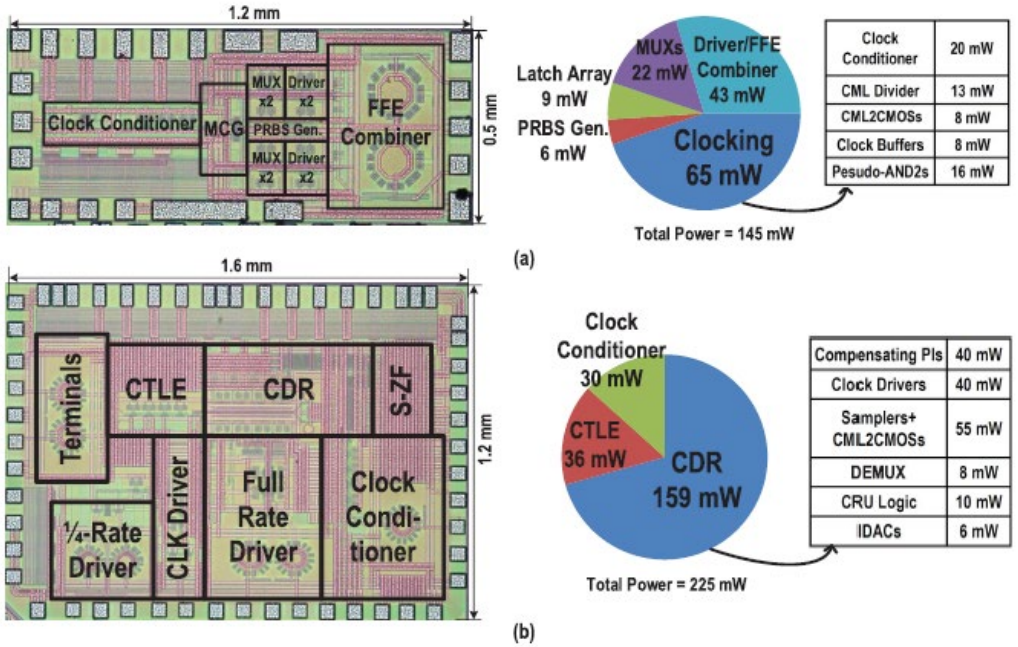


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

493

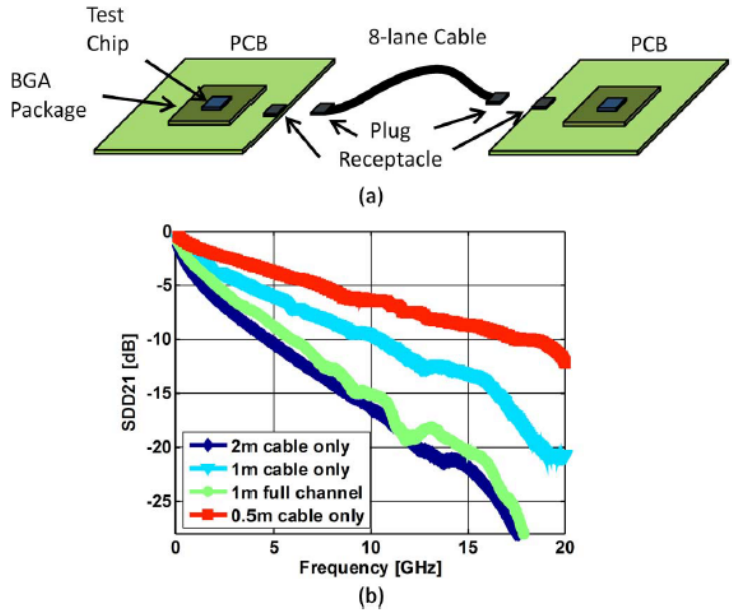
Claim 6	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "In the data transmission system 1 having the above configuration, the source 20 in the upstream device 2 generates a serial bit stream which is a video signal, for example. The source 20 transmits the generated serial bit stream from the downstream facing TX 20a to the retimer 21. The retimer 21 receives the serial bit stream transmitted from the downstream facing TX 20a in the upstream facing TX 21a. The retimer 21 performs a waveform shaping on the received serial bit stream and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 21b to the active cable 3." Kobayashi, 2:60-3:4.
6[b] performing clock and data recovery on the	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following

494

Claim 6	Disclosure
first electrical input signal with a first transceiver in the first connector to extract the first inbound data stream;	exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

495

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

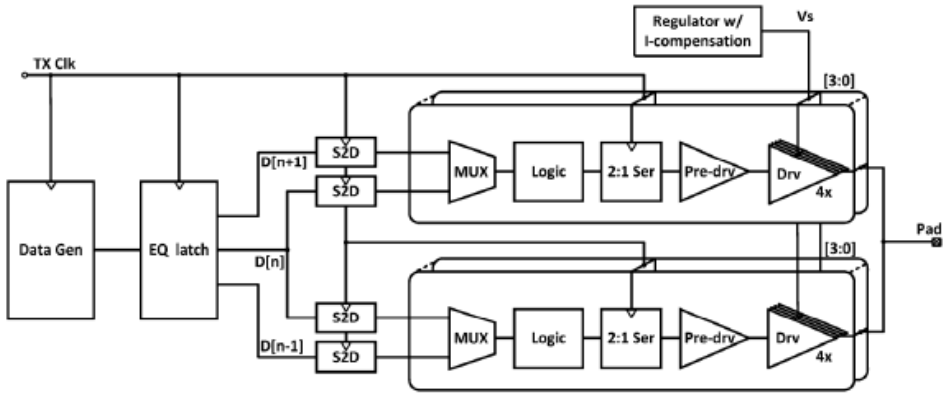
Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

497

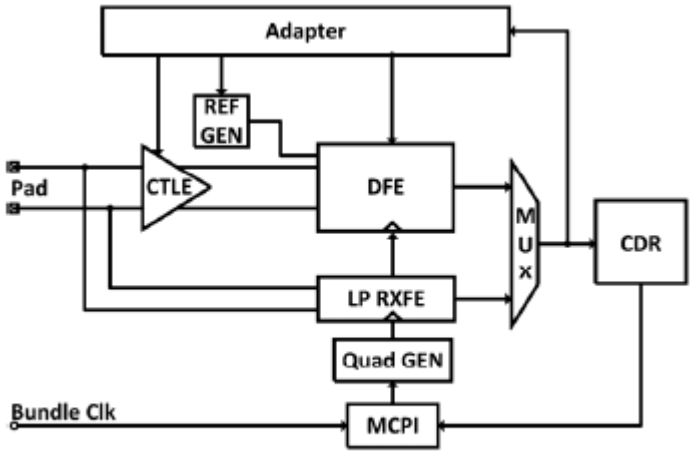
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

499

Claim 6	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

500

Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

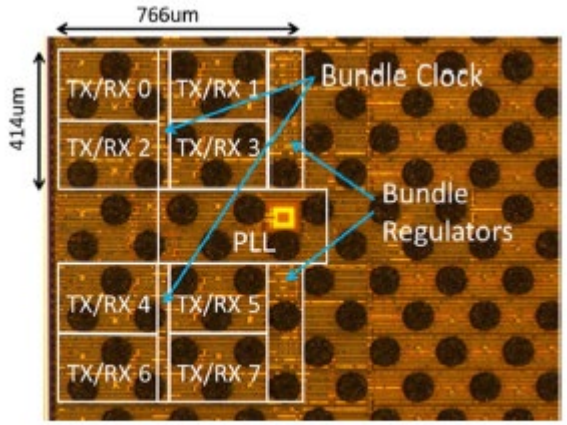
501

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

502

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

503

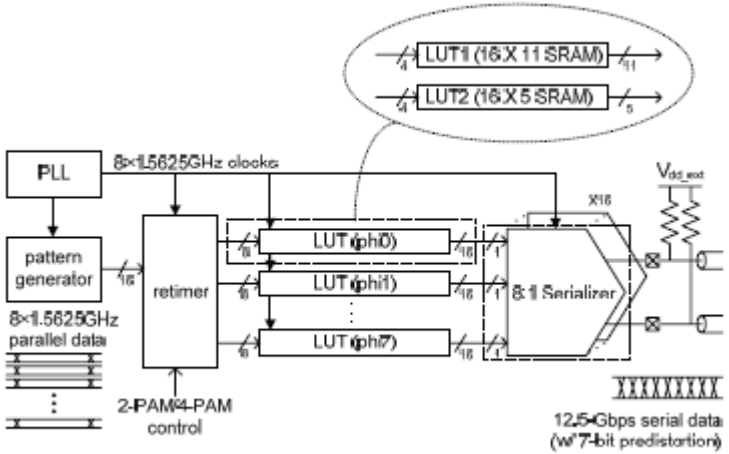
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

505

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

506

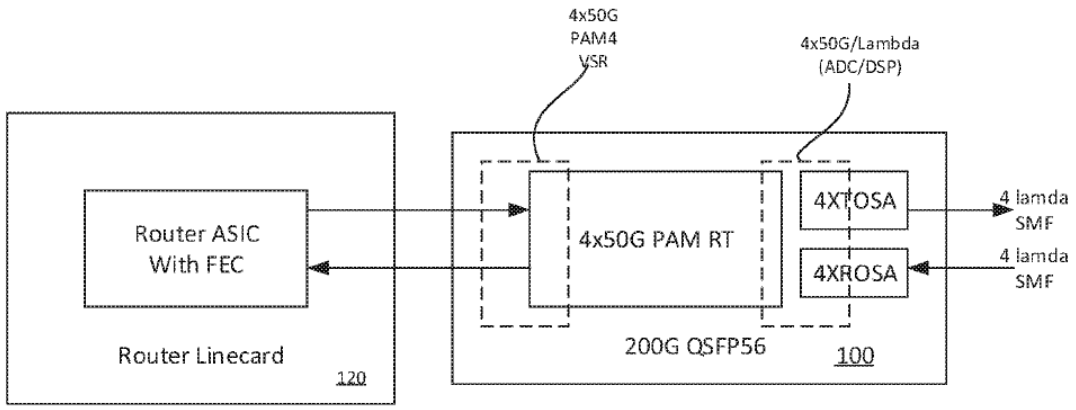
Claim 6	Disclosure
507	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

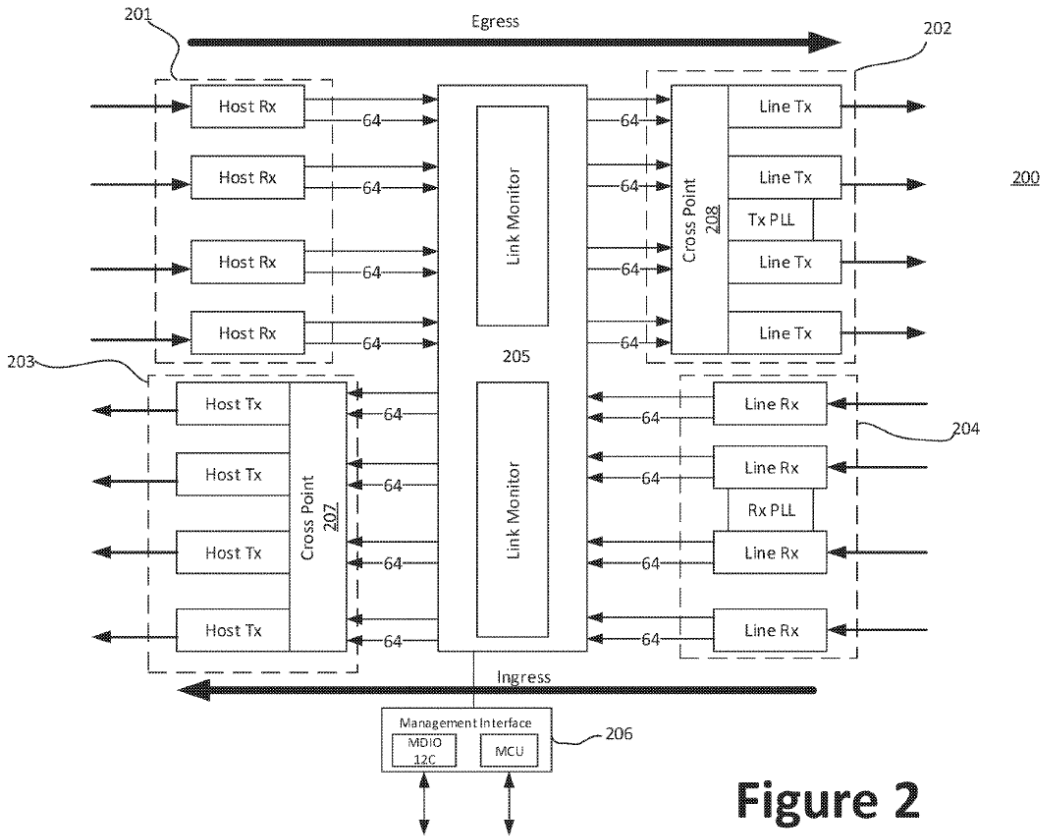
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

508

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

509

Claim 6	Disclosure
510	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

512

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

513

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

514

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

515

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

516

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

517

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

518

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

519

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

520

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

521

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

522

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

523

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

524

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

525

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

526

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

527

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

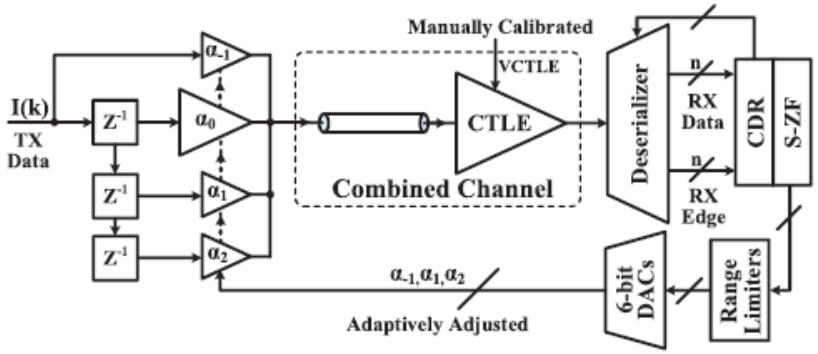


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

528

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

529

Claim 6	Disclosure

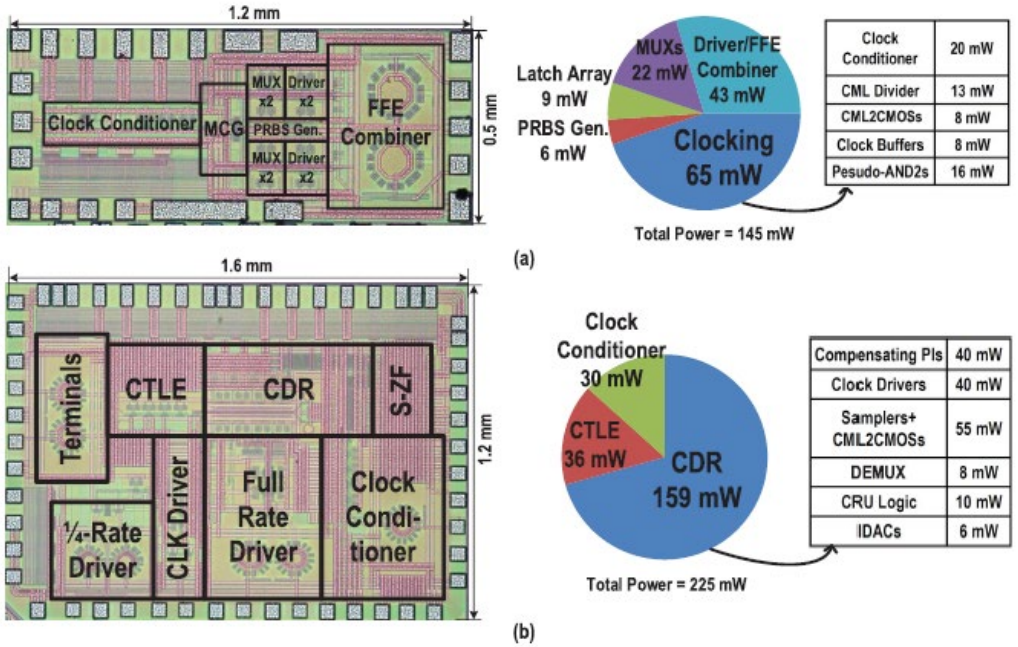


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

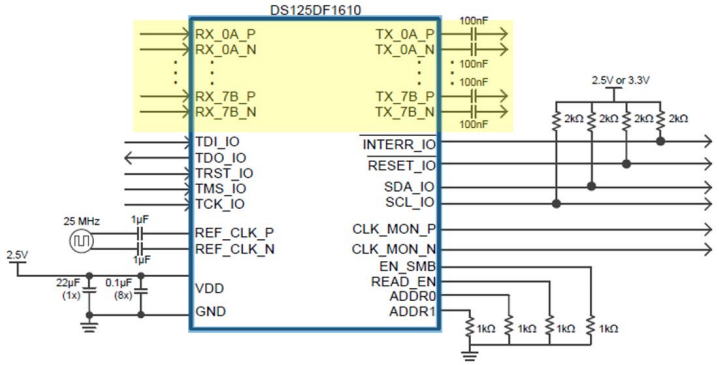
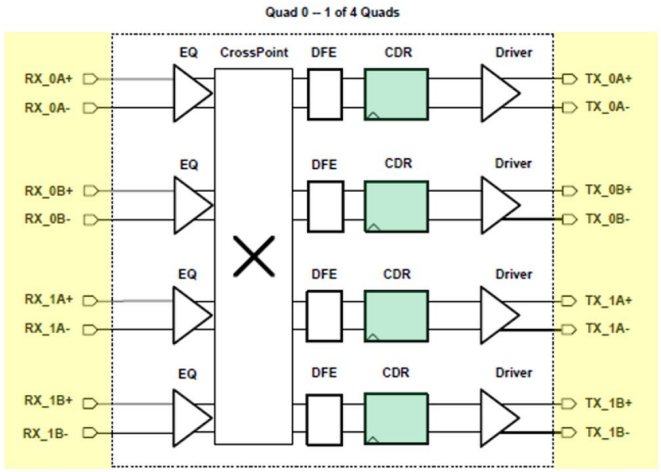
Zheng, 2973.

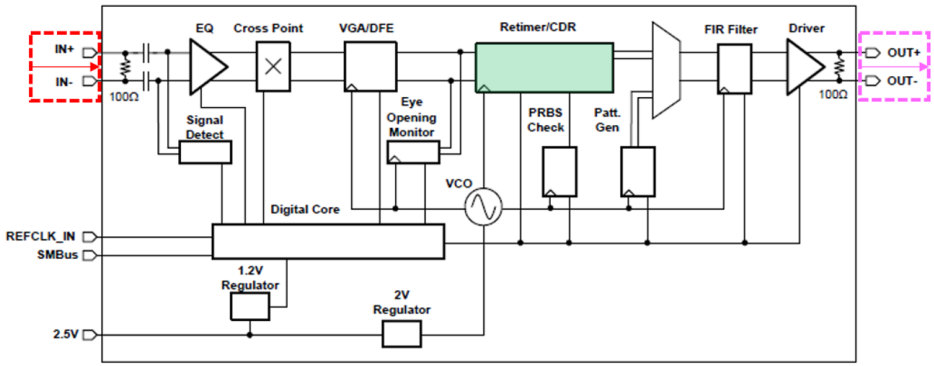
530

Claim 6	Disclosure
531	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local

Claim 6	Disclosure
	timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK" Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 circuit is illustrated in the Simplified Schematic and Figure 1 below. The highlighted pins illustrate sixteen channels of data, with each channel having two differential (Positive and Negative) input pins (RX_0A_P/N through RX_7B_P/N) to receive the channel of data and corresponding output pins (TX_0A_P/N through TX_7B_P/N) to transmit the channel of data after processing. *See, e.g.*, DS125DF1610 at §§ 3, 7.2. Thus, DS125DF1610's retimer is a transceiver that, when used in Tang's cable, receives one or more inbound data streams (on RX_0A_P/N through RX_7B_P/N), and transmits one or more transit data streams (via pins TX_0A_P/N through TX_7B_P/N).

532

Claim 6	Disclosure
533	**Simplified Schematic**   **Figure 1. DS125DF1610 Simplified Cross Point Diagram** Figure 1 above shows four channels and illustrates that DS125DF1610's retimer processes each data channel by, inter alia, performing "clock and data recovery (CDR)." *See id.* at §§ 3, 7.1. A channel's data path is shown in more detail in DS125DF1610 Figure 2 (below), which illustrates the input signal conveying an inbound data stream and the transit signal conveying a transit data stream in Tang's cable

Claim 6	Disclosure
	when DS125DF1610's retimer implements Tang's signal driver 44.  Figure 2. DS125DF1610 Simplified Data Path Diagram The Retimer/CDR circuit performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream. *See id*. at § 7.3.6. The '252 Patent specification makes clear that the input signal from the host can be pre-processed before CDR is performed on it. An ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data for output as the transit data stream because after extracting the input signal's timing and using that timing to sample and extract the input signal's data in an analog-to-digital CDR process, the retimer generates a new data stream for equalization by the FIR Filter (a digital-to-analog process) and to be driven for output by the differential driver (which produces an analog signal with configurable output voltage levels) to convey the output data stream (transit data stream when outputting over the cable) as an analog signal over the cable conductors, where the new signal is a "re-modulation" of the data. *See id*. at §§ 7.3.6-7.3.9.
6[c] re-modulating the first inbound data stream as a first transit data stream conveyed by a	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

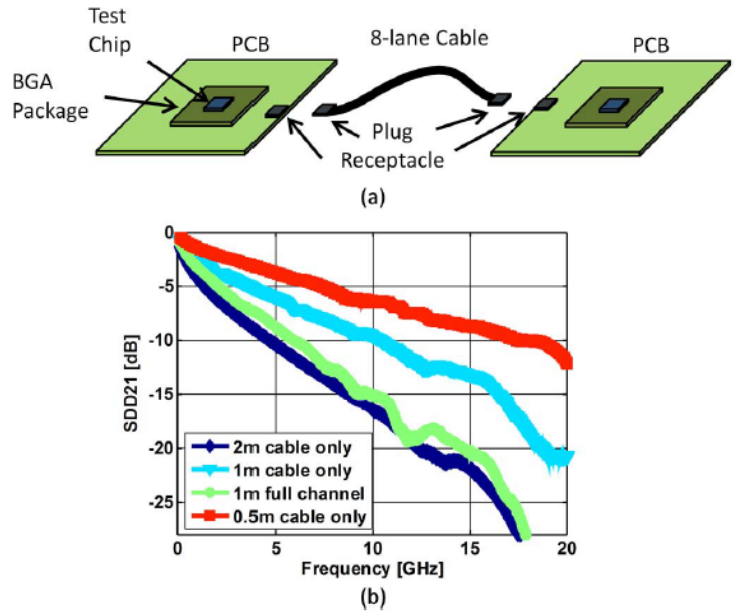
534

Claim 6	Disclosure
first electrical transit signal over a first of the conductor pairs;	*Abstract*—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

535

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

536

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

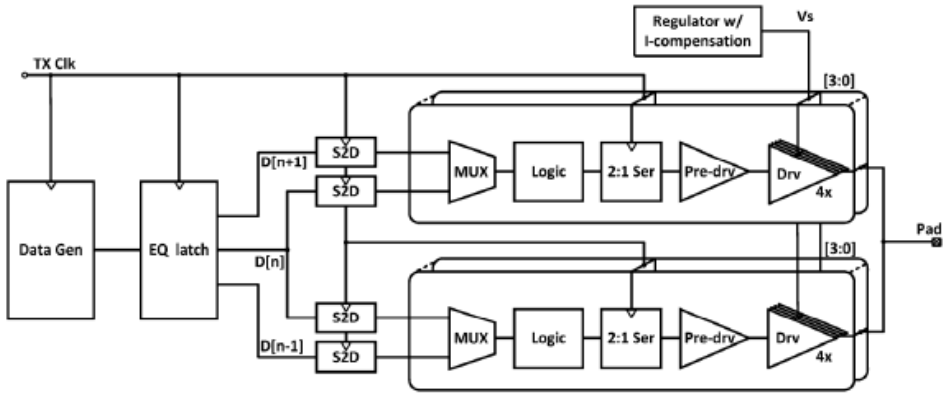
537

Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

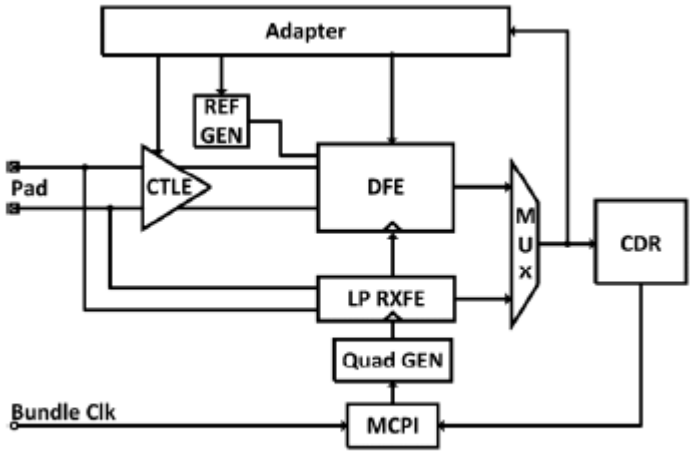
538

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

539

Claim 6	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

540

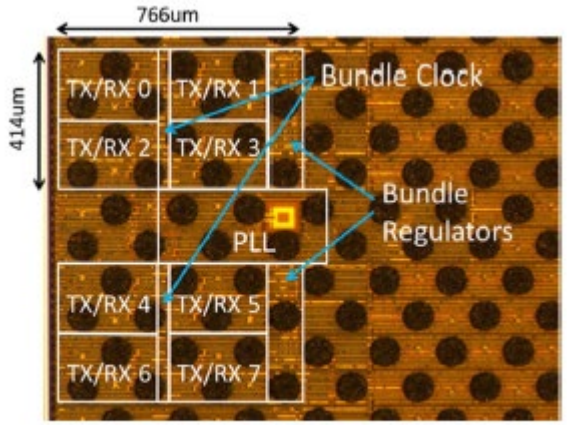
Claim 6	Disclosure
541	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

542

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

543

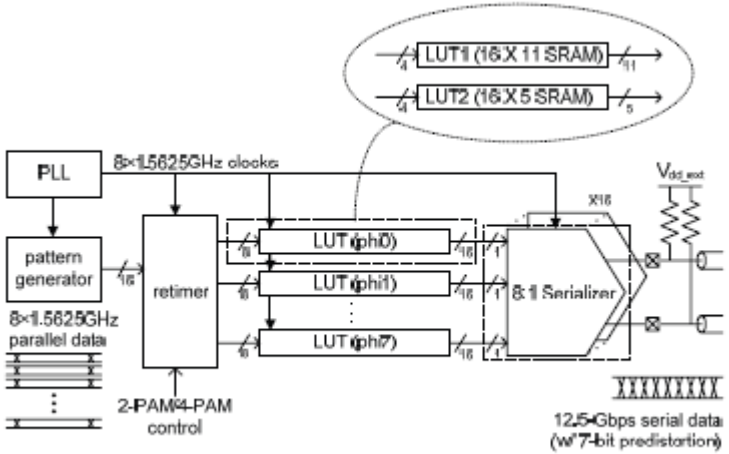
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

545

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

546

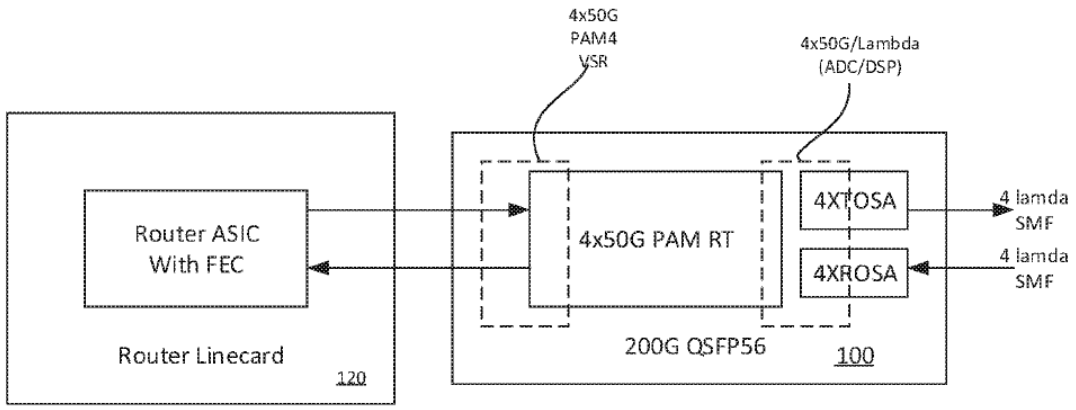
Claim 6	Disclosure
547	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

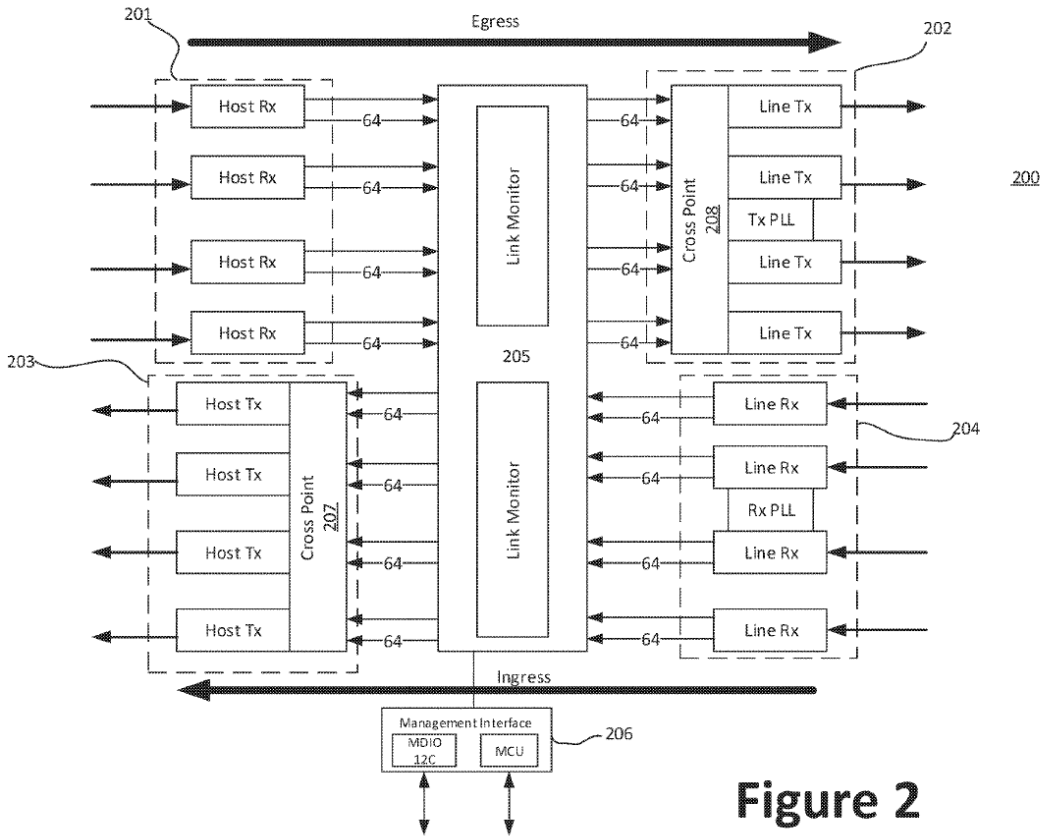
548

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

549

Claim 6	Disclosure
	 **Figure 1** Riani, Figure 1.

550

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

552

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

554

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

555

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

556

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

557

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

558

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

559

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

560

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

561

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

562

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

563

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

564

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

565

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

566

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

567

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

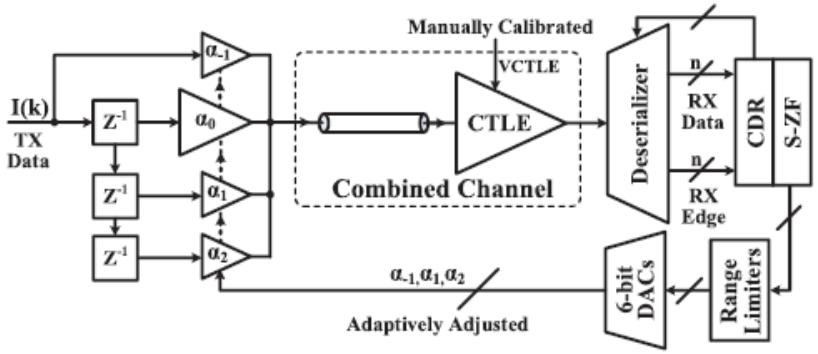


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

568

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

569

Claim 6	Disclosure

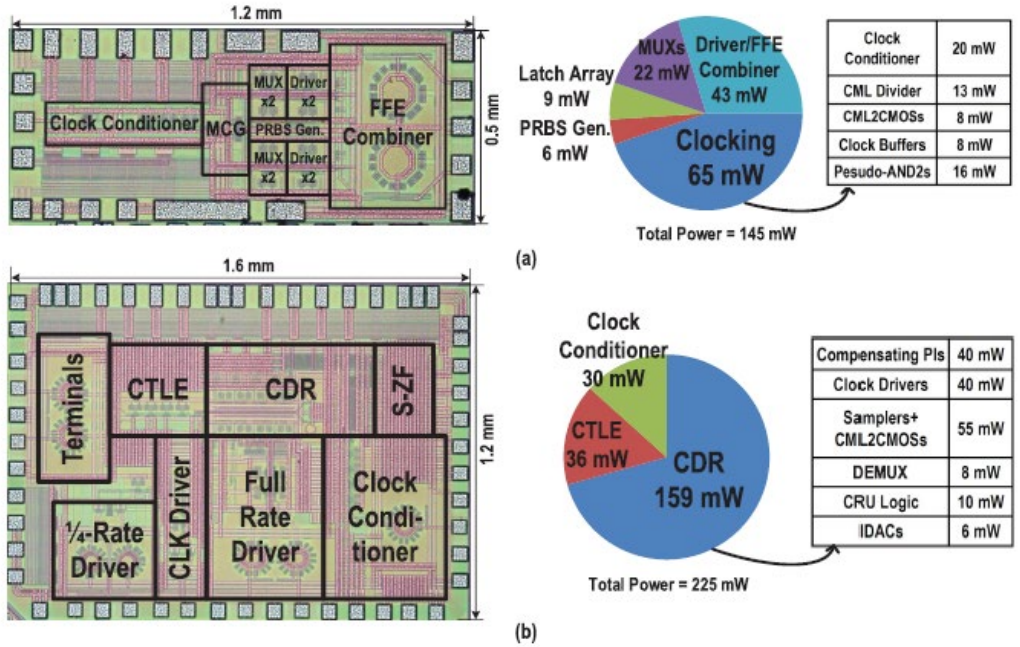


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

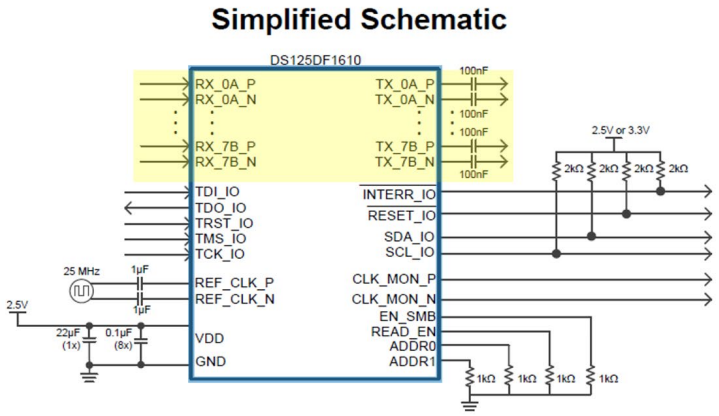
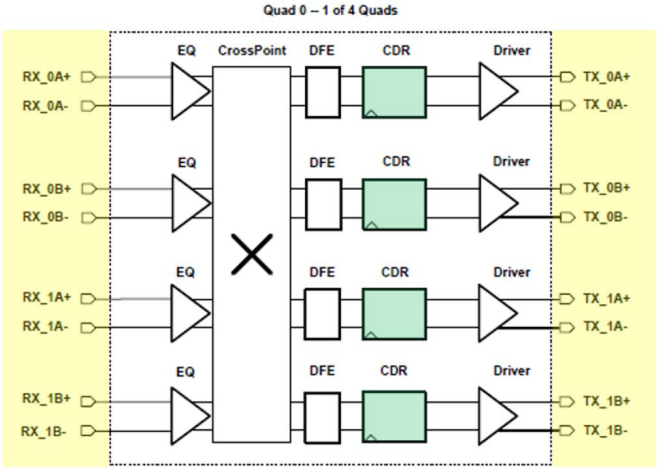
Zheng, 2973.

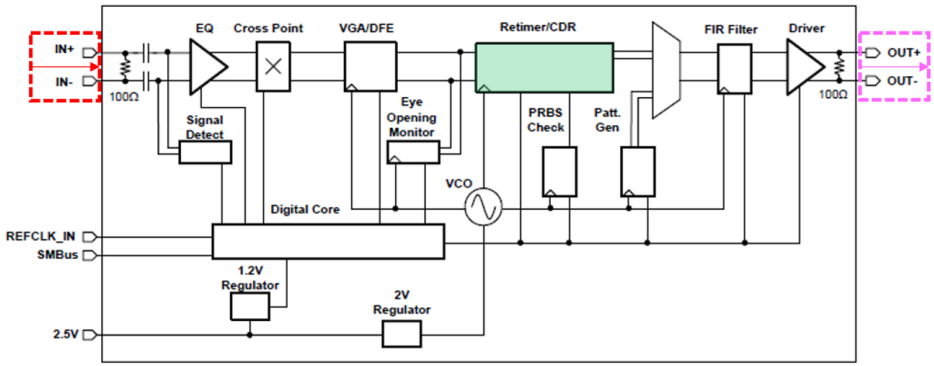
570

Claim 6	Disclosure
571	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local

Claim 6	Disclosure
	timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK" Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 circuit is illustrated in the Simplified Schematic and Figure 1 below. The highlighted pins illustrate sixteen channels of data, with each channel having two differential (Positive and Negative) input pins (RX_0A_P/N through RX_7B_P/N) to receive the channel of data and corresponding output pins (TX_0A_P/N through TX_7B_P/N) to transmit the channel of data after processing. *See, e.g.*, DS125DF1610 at §§ 3, 7.2. Thus, DS125DF1610's retimer is a transceiver that, when used in Tang's cable, receives one or more inbound data streams (on RX_0A_P/N through RX_7B_P/N), and transmits one or more transit data streams (via pins TX_0A_P/N through TX_7B_P/N).

572

Claim 6	Disclosure
573	  **Figure 1. DS125DF1610 Simplified Cross Point Diagram** Figure 1 above shows four channels and illustrates that DS125DF1610's retimer processes each data channel by, inter alia, performing "clock and data recovery (CDR)." *See id.* at §§ 3, 7.1. A channel's data path is shown in more detail in DS125DF1610 Figure 2 (below), which illustrates the input signal conveying an inbound data stream and the transit signal conveying a transit data stream in Tang's cable

Claim 6	Disclosure
	when DS125DF1610's retimer implements Tang's signal driver 44.  Figure 2. DS125DF1610 Simplified Data Path Diagram The Retimer/CDR circuit performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream. *See id*. at § 7.3.6. The '252 Patent specification makes clear that the input signal from the host can be pre-processed before CDR is performed on it. An ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data for output as the transit data stream because after extracting the input signal's timing and using that timing to sample and extract the input signal's data in an analog-to-digital CDR process, the retimer generates a new data stream for equalization by the FIR Filter (a digital-to-analog process) and to be driven for output by the differential driver (which produces an analog signal with configurable output voltage levels) to convey the output data stream (transit data stream when outputting over the cable) as an analog signal over the cable conductors, where the new signal is a "re-modulation" of the data. *See id*. at §§ 7.3.6-7.3.9s.
6[d] receiving with the second connector a second electrical input	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

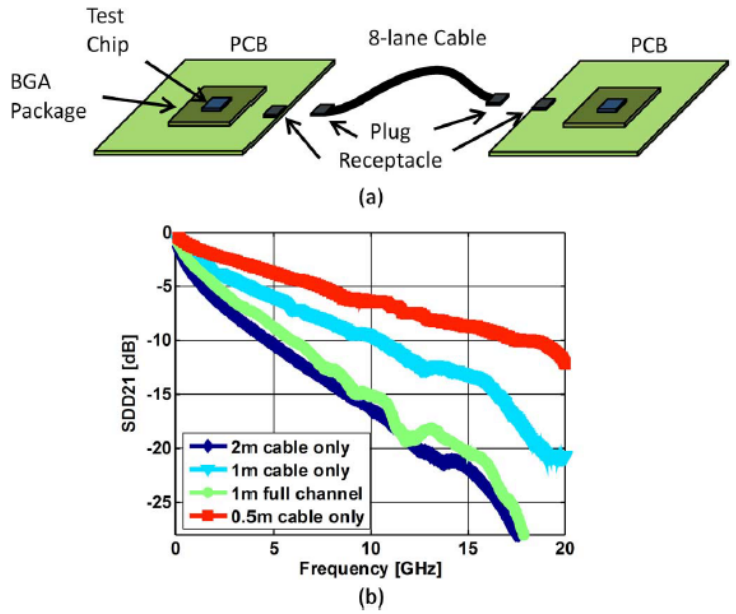
574

Claim 6	Disclosure
signal conveying a second inbound data stream from a second host device;	*Abstract*—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

575

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

576

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

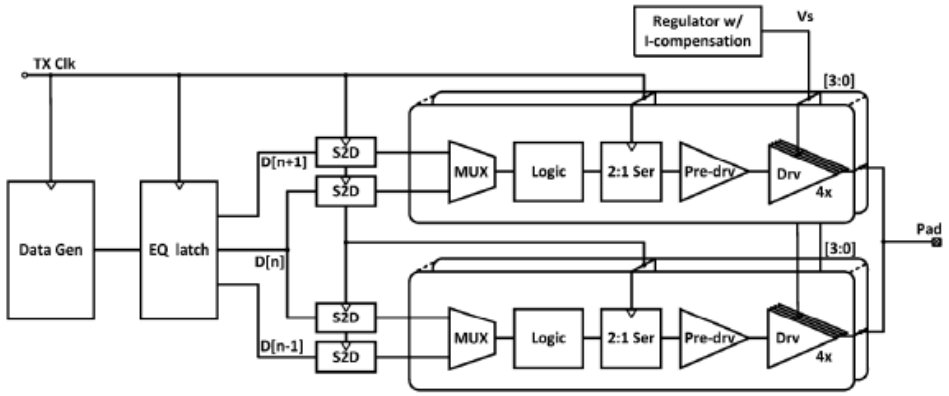
577

Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

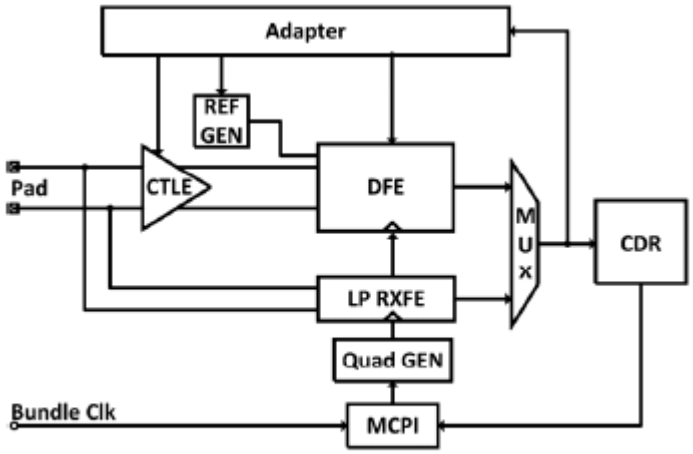
578

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

579

Claim 6	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

580

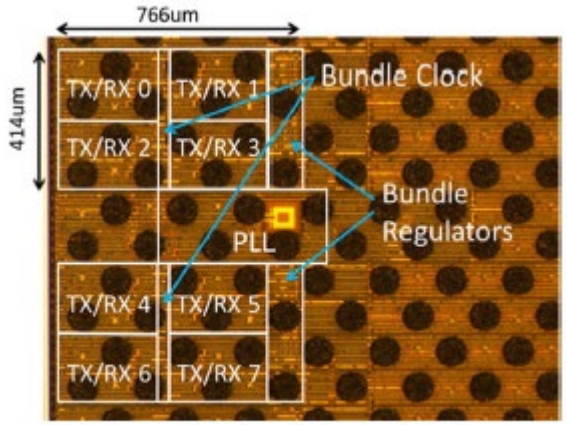
Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

582

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

583

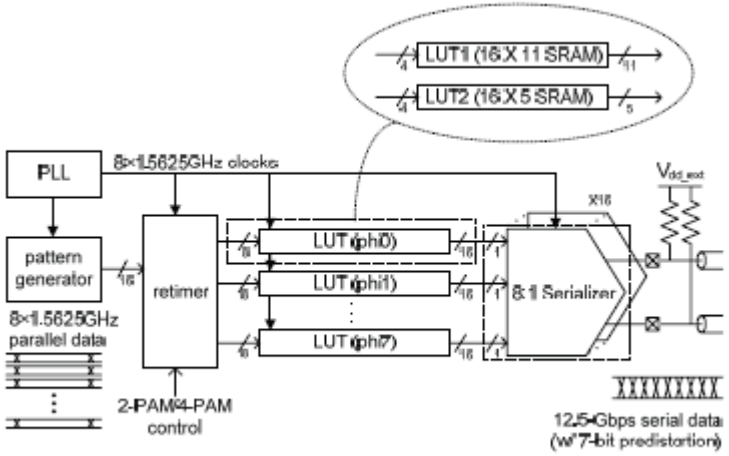
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

585

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

586

Claim 6	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

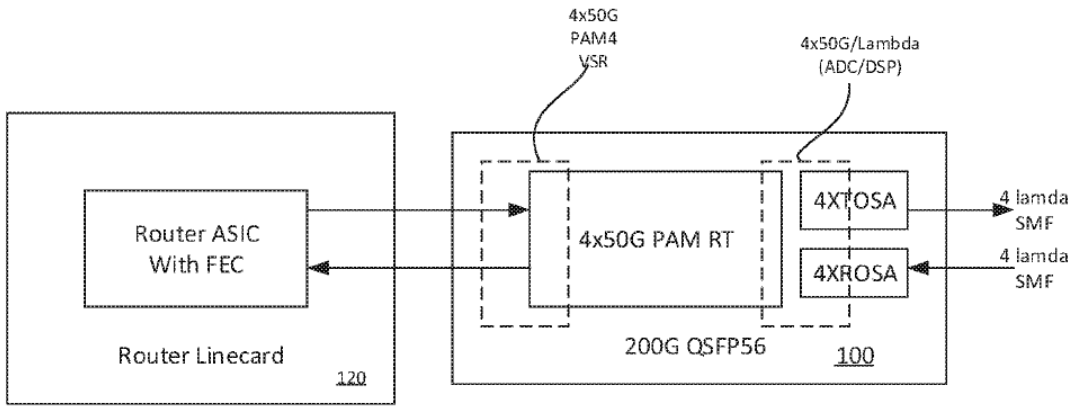
587

Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

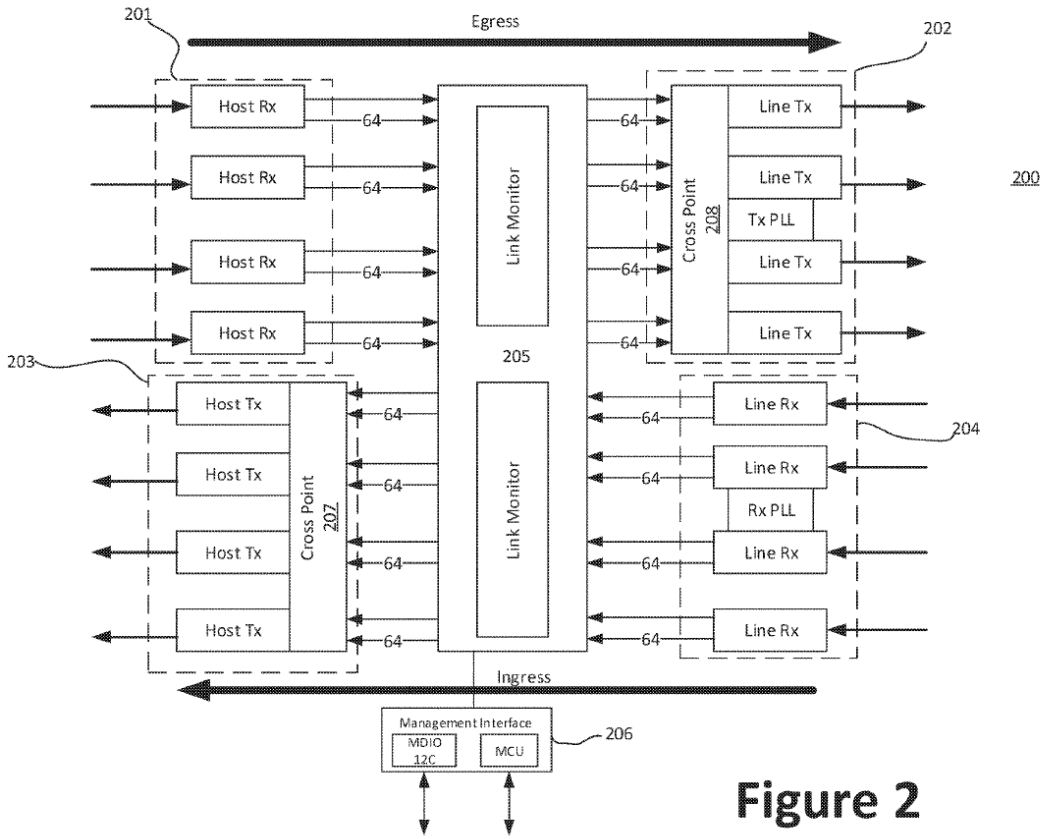
588

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

589

Claim 6	Disclosure
	 **Figure 1** Riani, Figure 1.

590

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

592

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

595

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

596

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

597

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

598

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

599

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

600

Claim 6	Disclosure
601	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

602

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

603

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

604

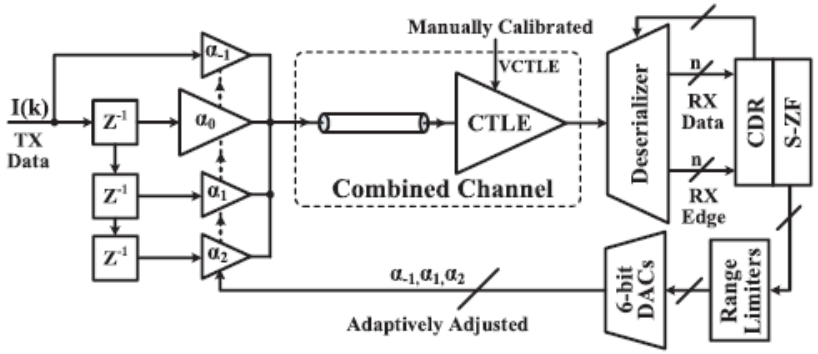
Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

605

Claim 6	Disclosure
609	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

607

Claim 6	Disclosure
608	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

609

Claim 6	Disclosure

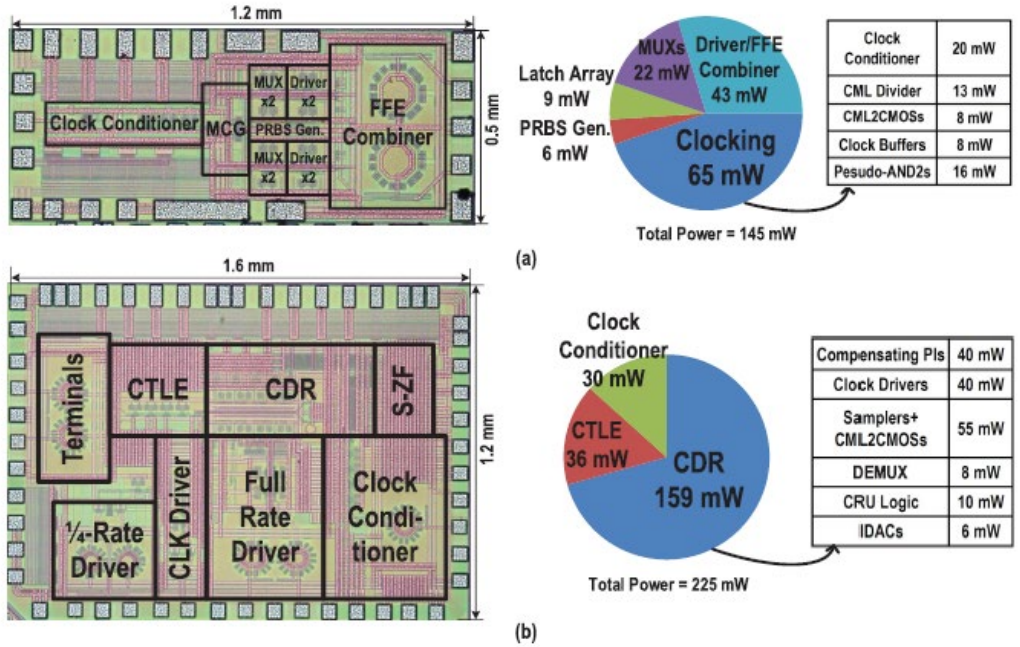


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

610

Claim 6	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The downstream device 4 receives the serial bit stream transmitted from the active cable 3 in the upstream facing RX 40a in the retimer 40. The retimer 40 performs a waveform shaping on the serial bit stream received by the upstream facing RX 40a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 40b to a sink 41. The sink 41 displays a video, for example, based on the received serial bit data." Kobayashi, 3:20-28.
6[e] performing clock and data recovery on the second electrical input signal with a second transceiver in the second	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

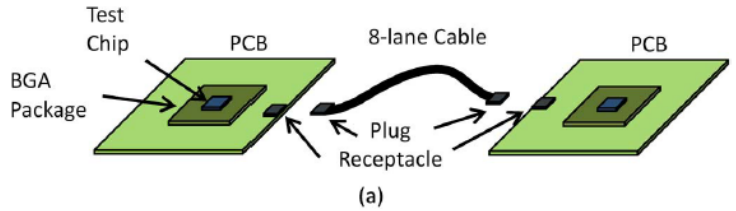
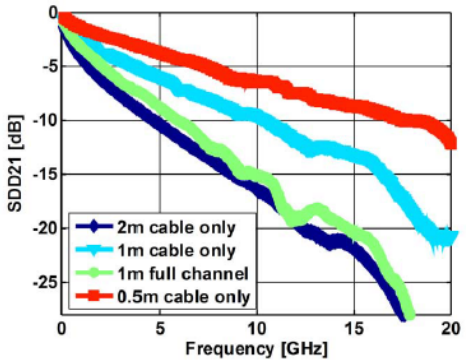
611

Claim 6	Disclosure
connector to extract the second inbound data stream;	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

612

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

613

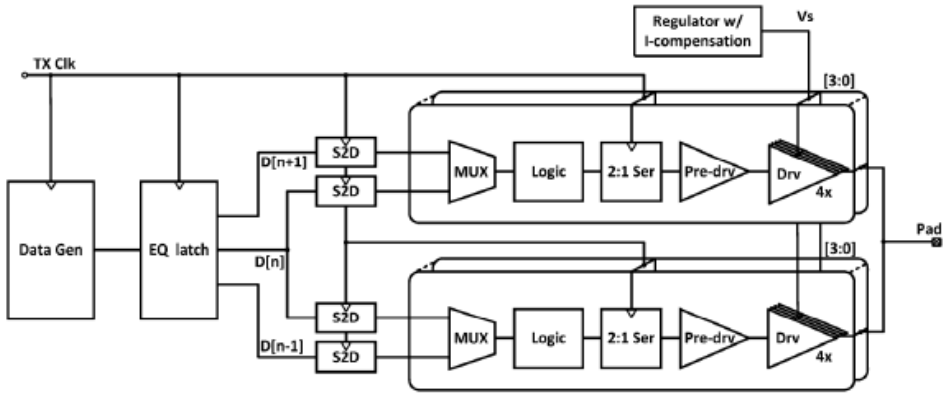
Claim 6	Disclosure
614	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

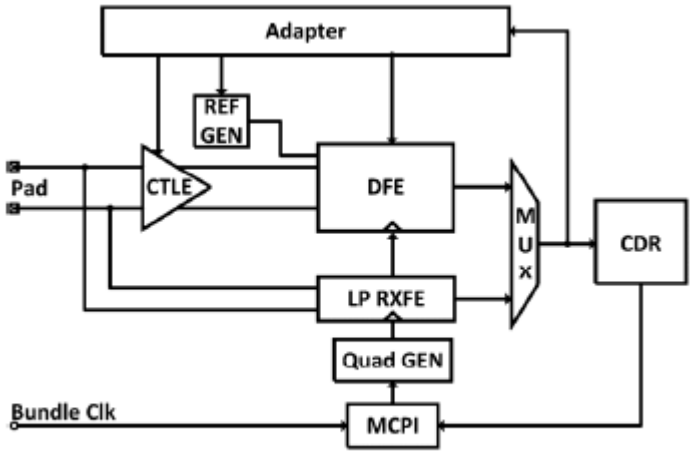
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

615

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

616

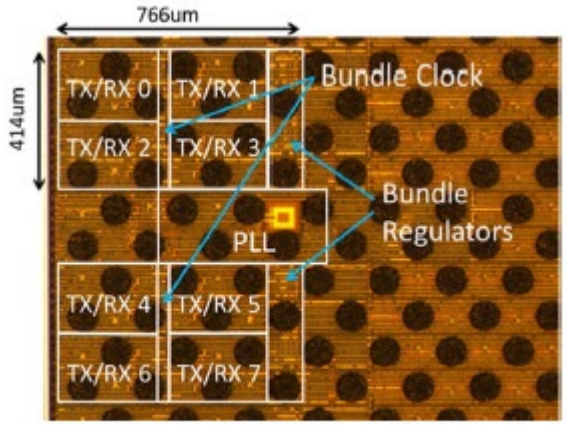
Claim 6	Disclosure
617	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 6	Disclosure
618	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
619	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

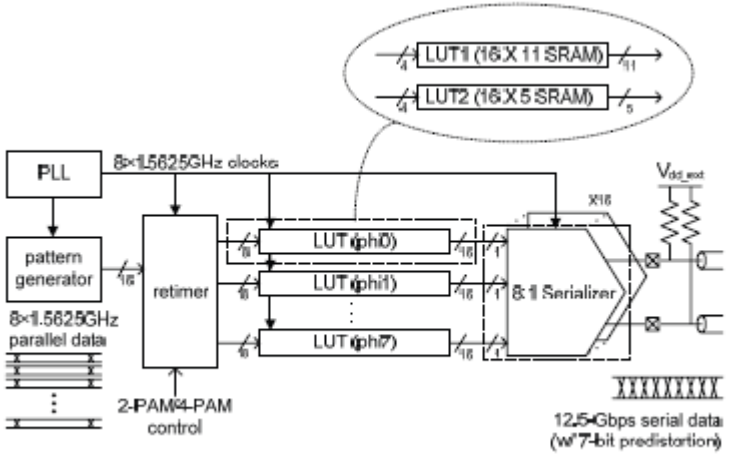
620

Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

621

Claim 6	Disclosure
622	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

Claim 6	Disclosure
623	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

Claim 6	Disclosure
624	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

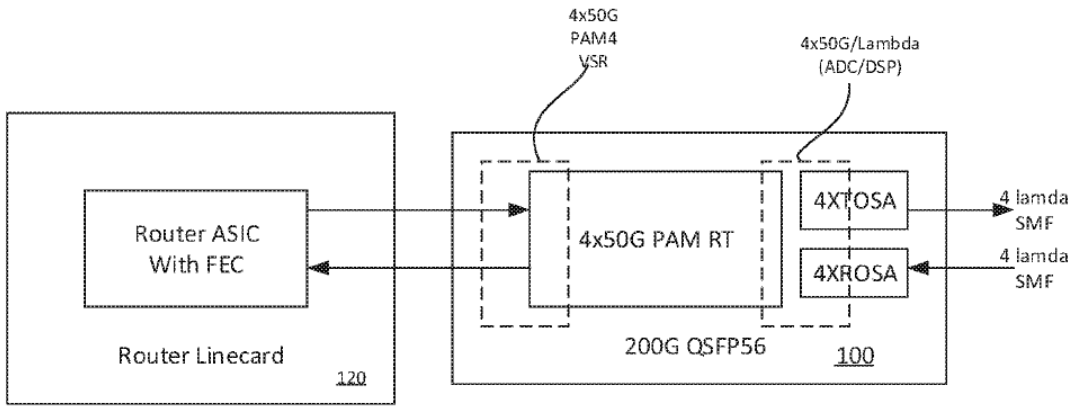
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

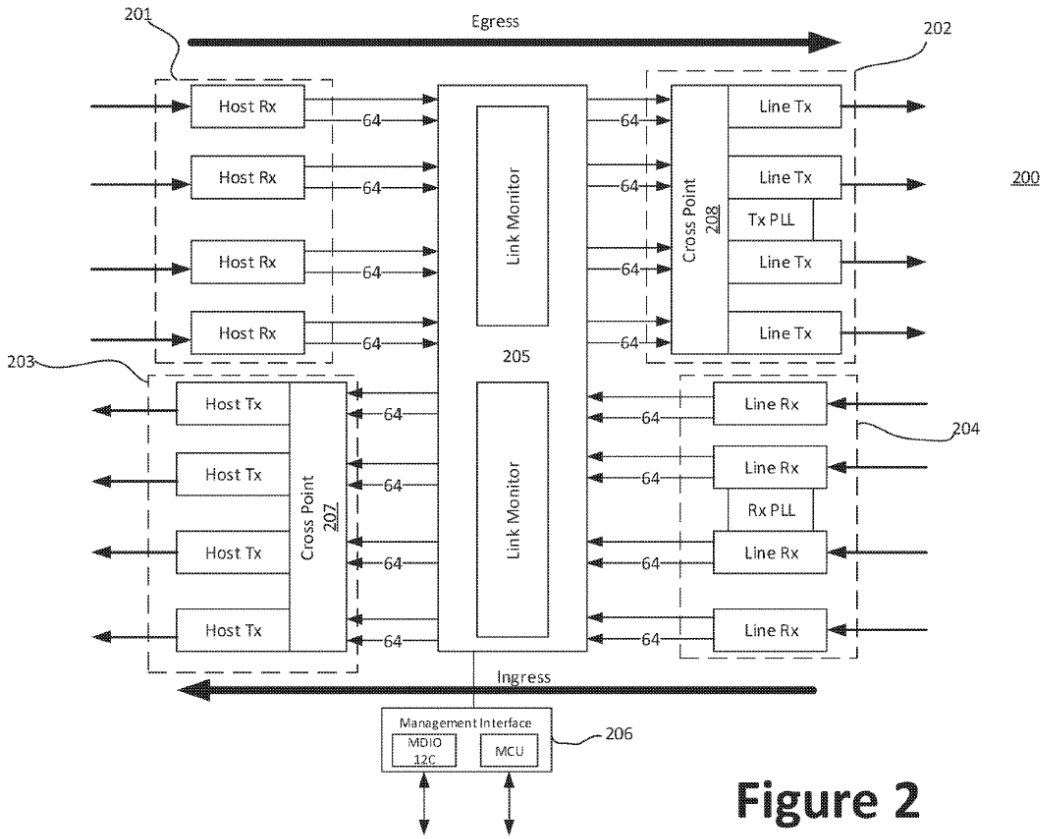
625

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

626

627

Claim 6	Disclosure
	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

629

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

630

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

631

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

632

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

633

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

634

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

635

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

636

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

637

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

638

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

639

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

640

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

641

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

642

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

643

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

644

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

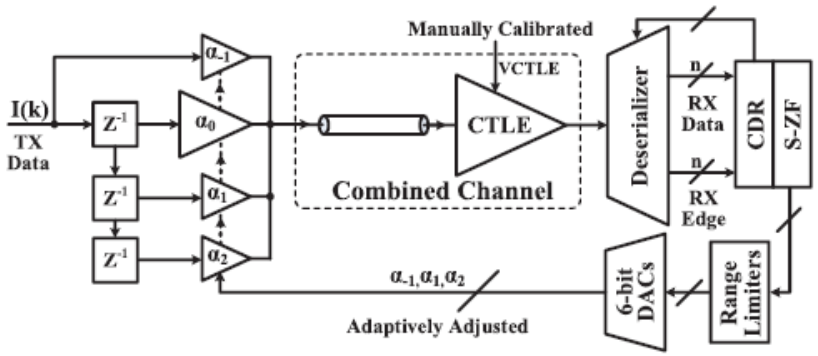


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

645

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

646

Claim 6	Disclosure

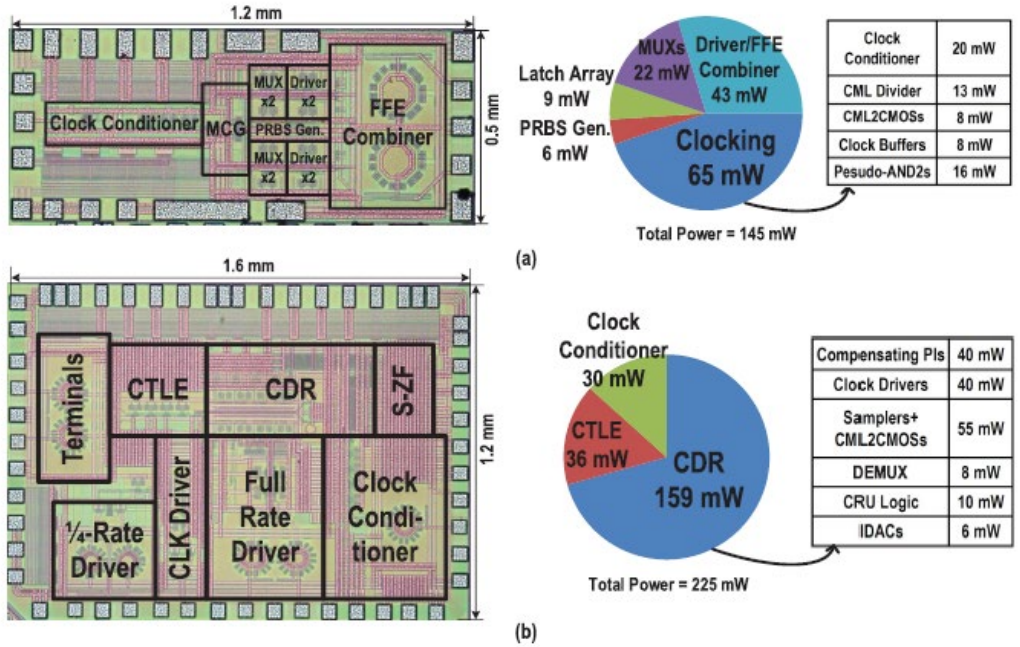


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

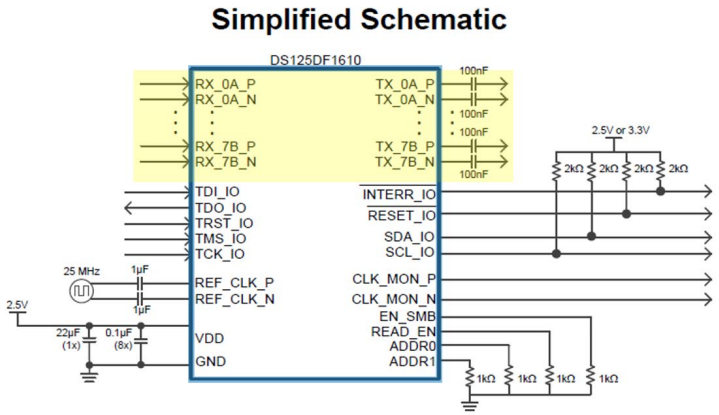
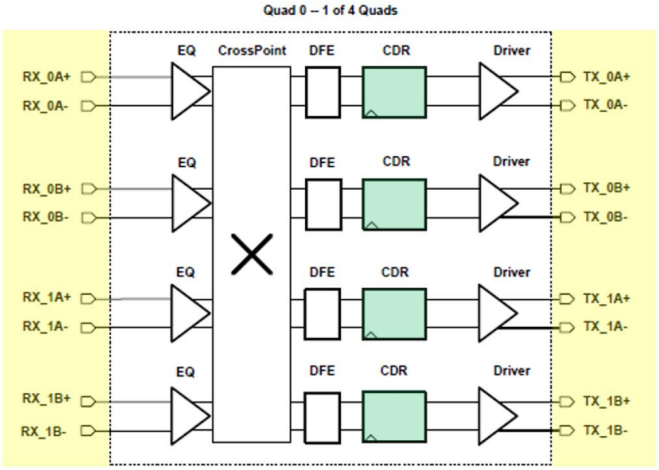
Zheng, 2973.

647

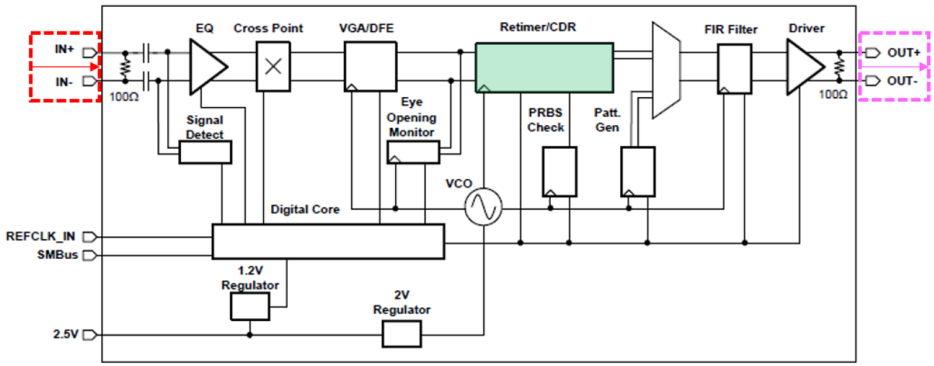
Claim 6	Disclosure
648	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

Claim 6	Disclosure
	Kobayashi, 4:4-22. <u>*DS125DF1610 Retimer Publication*</u> The DS125DF1610 circuit is illustrated in the Simplified Schematic and Figure 1 below. The highlighted pins illustrate sixteen channels of data, with each channel having two differential (Positive and Negative) input pins (RX_0A_P/N through RX_7B_P/N) to receive the channel of data and corresponding output pins (TX_0A_P/N through TX_7B_P/N) to transmit the channel of data after processing. *See, e.g.*, DS125DF1610 at §§ 3, 7.2. Thus, DS125DF1610's retimer is a transceiver that, when used in Tang's cable, receives one or more inbound data streams (on RX_0A_P/N through RX_7B_P/N), and transmits one or more transit data streams (via pins TX_0A_P/N through TX_7B_P/N).

649

Claim 6	Disclosure
	  **Figure 1. DS125DF1610 Simplified Cross Point Diagram** Figure 1 above shows four channels and illustrates that DS125DF1610's retimer processes each data channel by, inter alia, performing "clock and data recovery (CDR)." *See id.* at §§ 3, 7.1. A channel's data path is shown in more detail in DS125DF1610 Figure 2 (below), which illustrates the input signal conveying an inbound data stream and the transit signal conveying a transit data stream in Tang's cable

650

Claim 6	Disclosure
	when DS125DF1610's retimer implements Tang's signal driver 44.  Figure 2. DS125DF1610 Simplified Data Path Diagram The Retimer/CDR circuit performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream. *See id*. at § 7.3.6. The '252 Patent specification makes clear that the input signal from the host can be pre-processed before CDR is performed on it. An ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data for output as the transit data stream because after extracting the input signal's timing and using that timing to sample and extract the input signal's data in an analog-to-digital CDR process, the retimer generates a new data stream for equalization by the FIR Filter (a digital-to-analog process) and to be driven for output by the differential driver (which produces an analog signal with configurable output voltage levels) to convey the output data stream (transit data stream when outputting over the cable) as an analog signal over the cable conductors, where the new signal is a "re-modulation" of the data. *See id*. at §§ 7.3.6-7.3.9.
6[f] re-modulating the second inbound data stream as a second transit	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

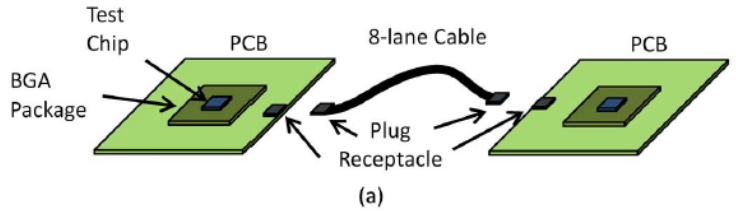
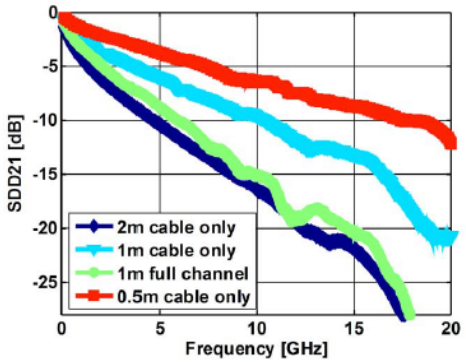
651

Claim 6	Disclosure
data stream conveyed by a second electrical transit signal over a second of the conductor pairs;	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

652

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

653

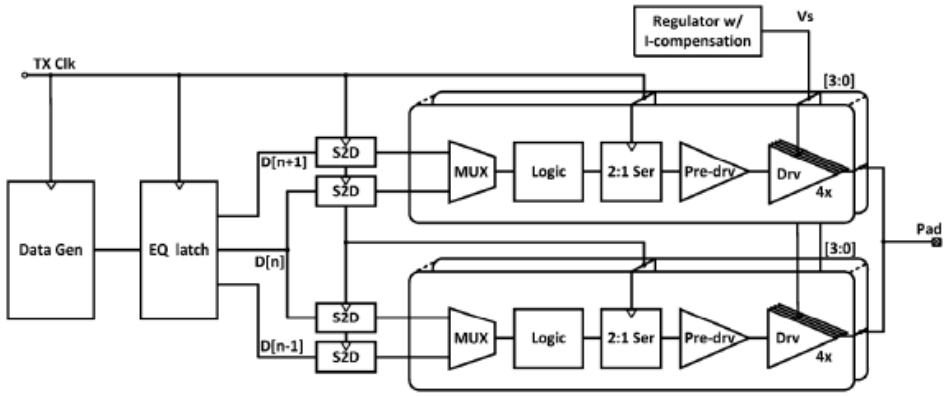
Claim 6	Disclosure
654	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

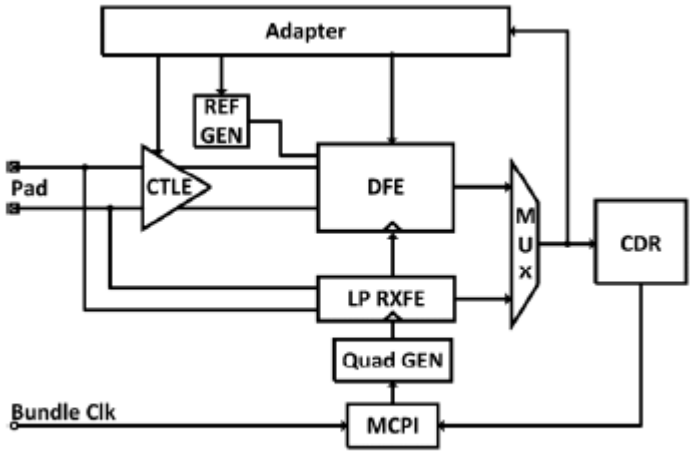
655

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

656

Claim 6	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

657

Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

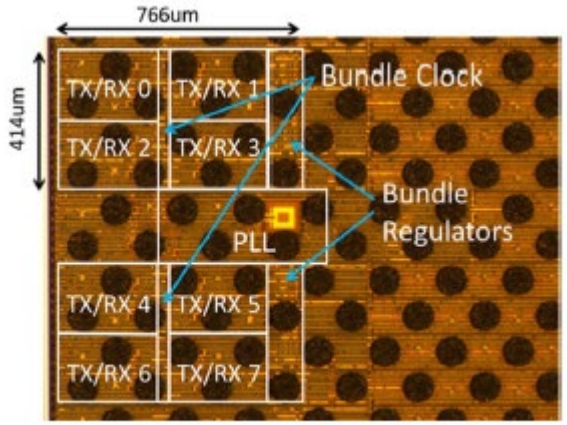
658

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

659

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

660

Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

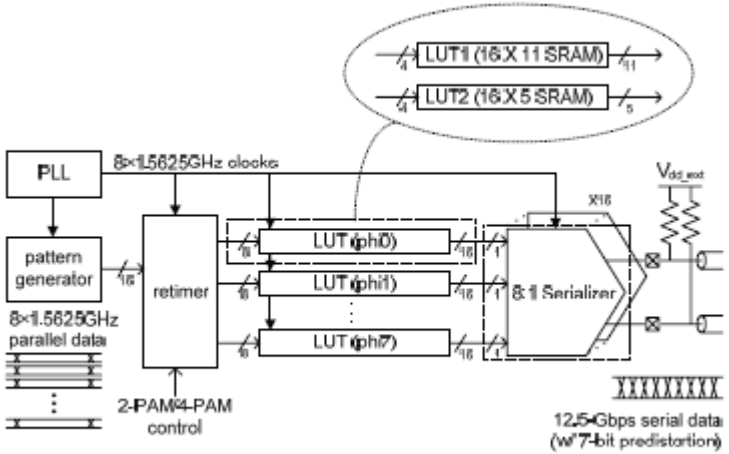
661

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

662

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

663

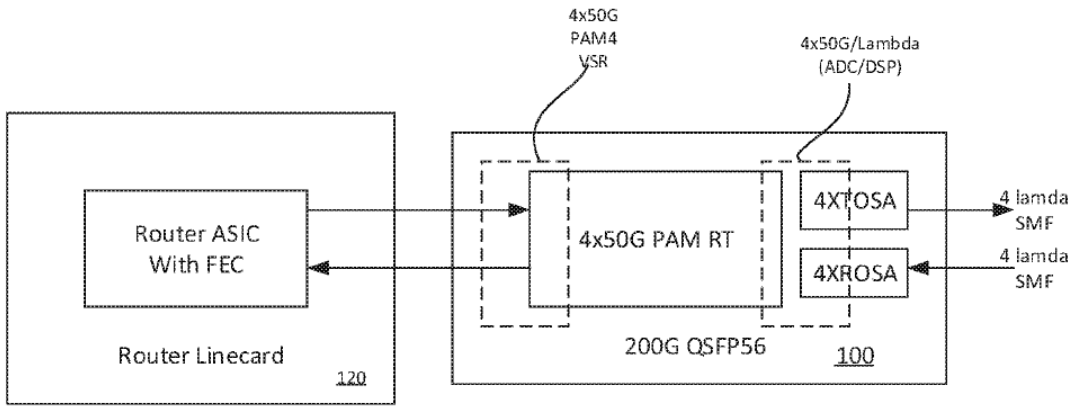
Claim 6	Disclosure
664	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

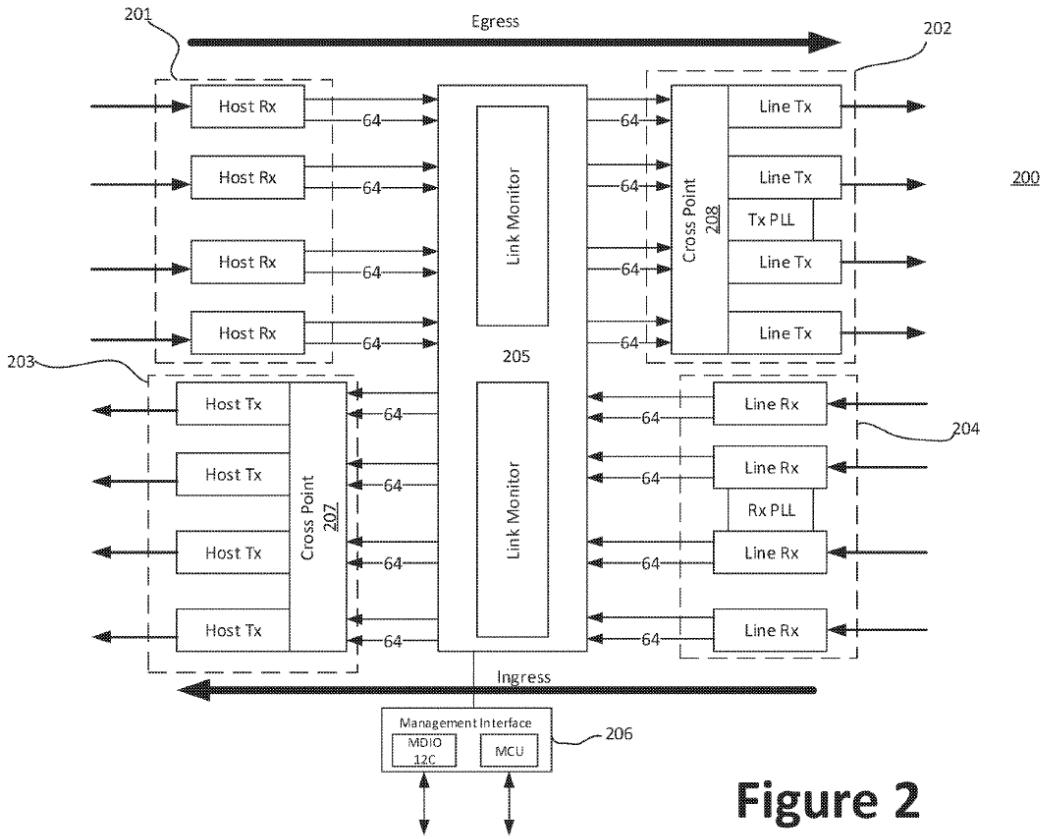
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

665

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

999

Claim 6	Disclosure
667	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
668	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

669

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

670

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

671

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

672

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

673

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

674

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

675

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

676

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

677

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

678

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

679

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

680

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

681

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

682

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

683

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

684

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

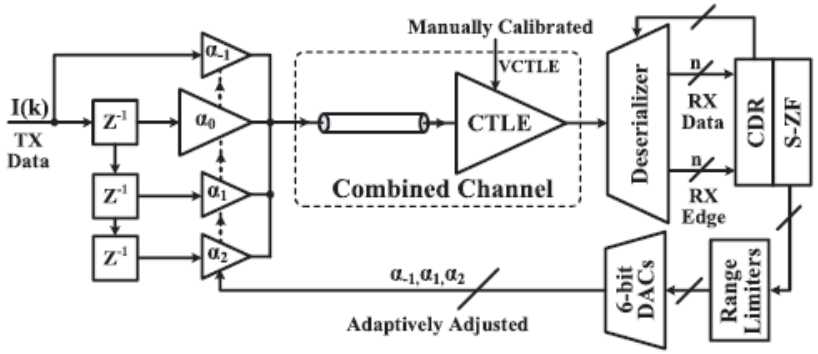


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

685

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

689

Claim 6	Disclosure

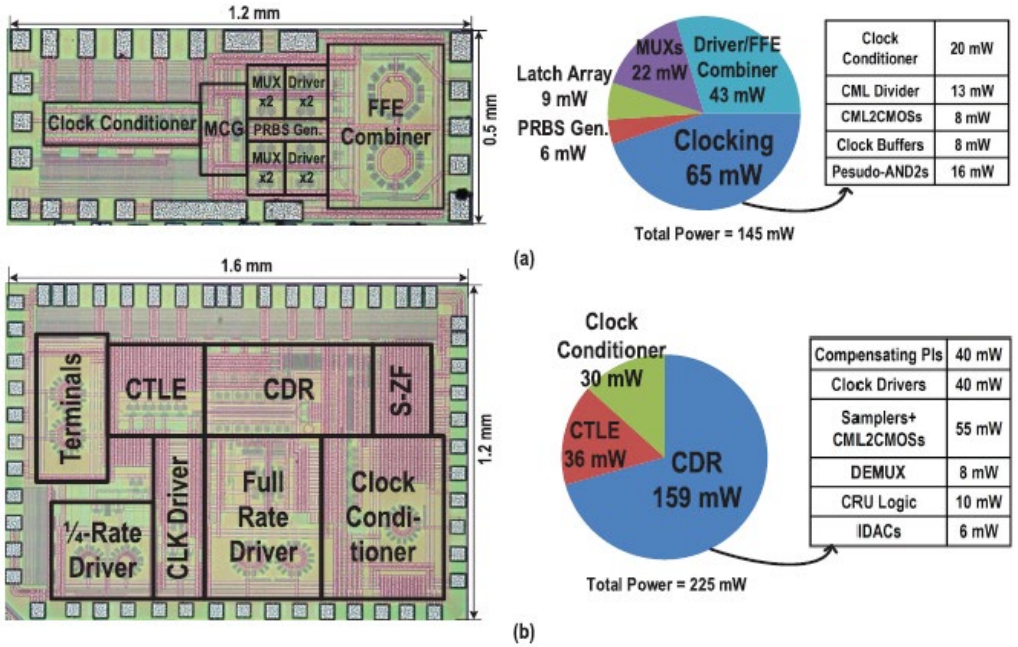


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

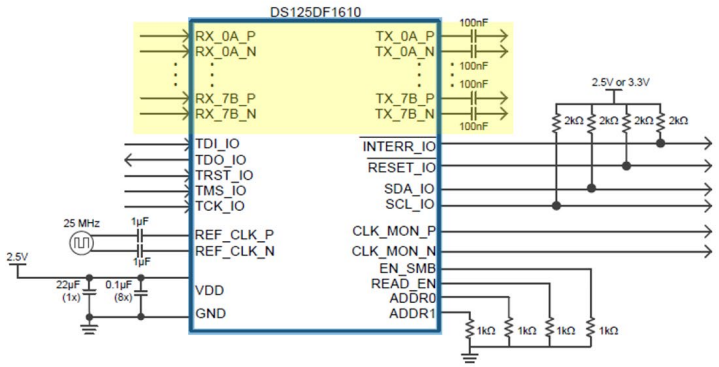
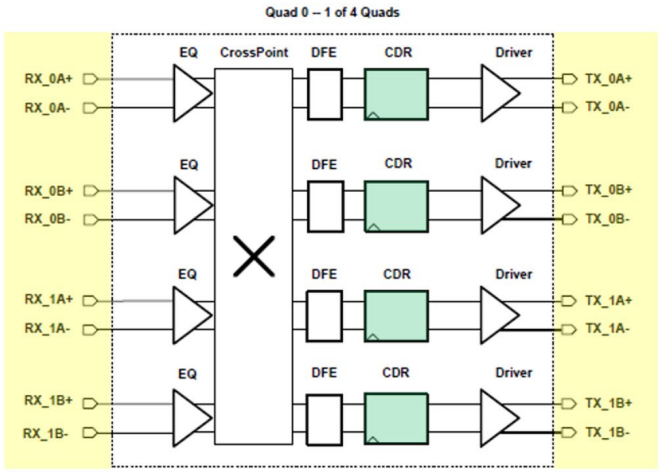
Zheng, 2973.

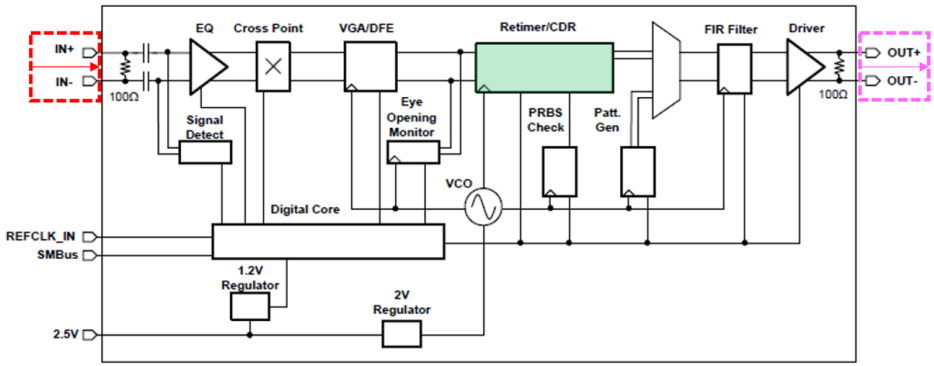
687

Claim 6	Disclosure
688	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

Claim 6	Disclosure
	Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 circuit is illustrated in the Simplified Schematic and Figure 1 below. The highlighted pins illustrate sixteen channels of data, with each channel having two differential (Positive and Negative) input pins (RX_0A_P/N through RX_7B_P/N) to receive the channel of data and corresponding output pins (TX_0A_P/N through TX_7B_P/N) to transmit the channel of data after processing. *See, e.g.*, DS125DF1610 at §§ 3, 7.2. Thus, DS125DF1610's retimer is a transceiver that, when used in Tang's cable, receives one or more inbound data streams (on RX_0A_P/N through RX_7B_P/N), and transmits one or more transit data streams (via pins TX_0A_P/N through TX_7B_P/N).

689

Claim 6	Disclosure
069	**Simplified Schematic**   **Figure 1. DS125DF1610 Simplified Cross Point Diagram** Figure 1 above shows four channels and illustrates that DS125DF1610's retimer processes each data channel by, inter alia, performing "clock and data recovery (CDR)." *See id.* at §§ 3, 7.1. A channel's data path is shown in more detail in DS125DF1610 Figure 2 (below), which illustrates the input signal conveying an inbound data stream and the transit signal conveying a transit data stream in Tang's cable

Claim 6	Disclosure
	when DS125DF1610's retimer implements Tang's signal driver 44.  Figure 2. DS125DF1610 Simplified Data Path Diagram The Retimer/CDR circuit performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream. *See id*. at § 7.3.6. The '252 Patent specification makes clear that the input signal from the host can be pre-processed before CDR is performed on it. An ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data for output as the transit data stream because after extracting the input signal's timing and using that timing to sample and extract the input signal's data in an analog-to-digital CDR process, the retimer generates a new data stream for equalization by the FIR Filter (a digital-to-analog process) and to be driven for output by the differential driver (which produces an analog signal with configurable output voltage levels) to convey the output data stream (transit data stream when outputting over the cable) as an analog signal over the cable conductors, where the new signal is a "re-modulation" of the data. *See id*. at §§ 7.3.6-7.3.9.
6[g] performing clock and data recovery on the first electrical transit	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

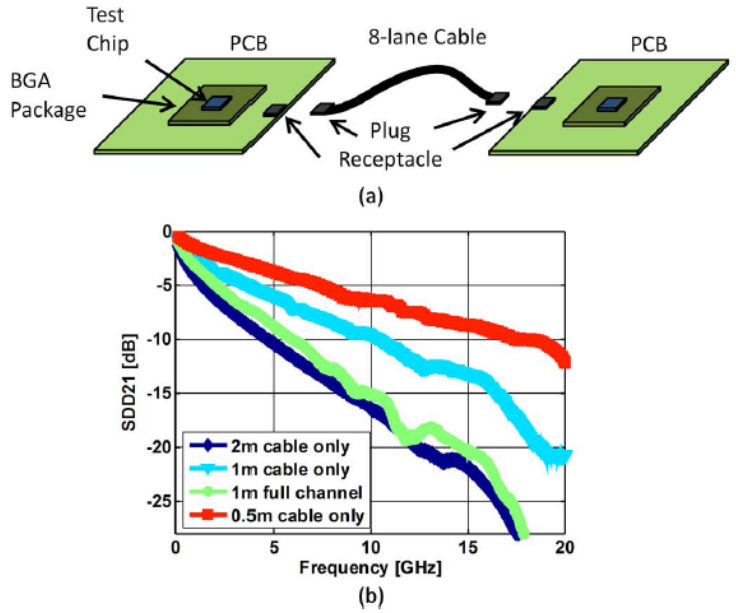
691

Claim 6	Disclosure
signal with the second transceiver to extract the first transit data stream;	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

692

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

693

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

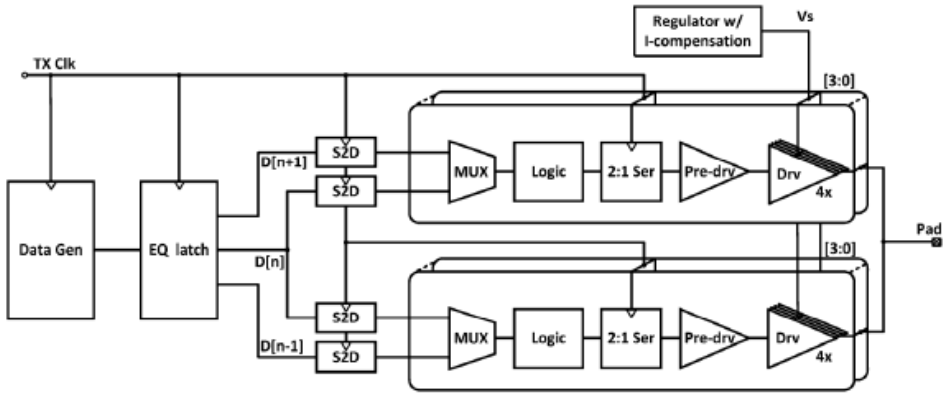
694

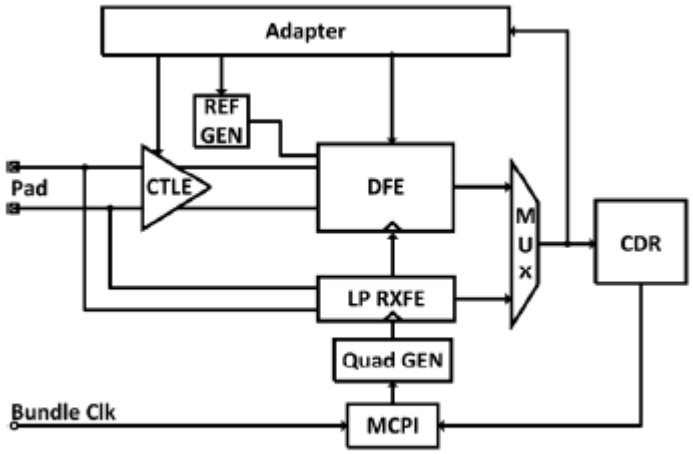
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

695

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

969

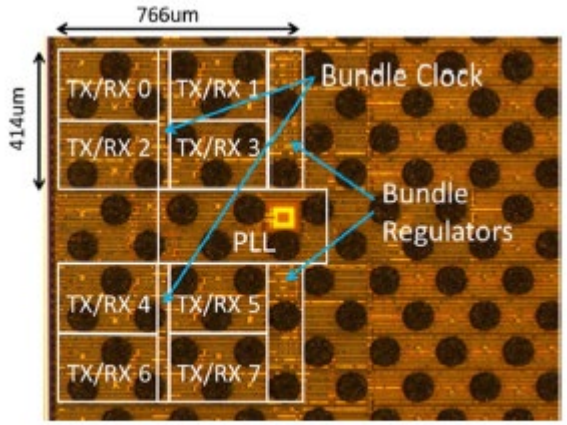
Claim 6	Disclosure
697	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 6	Disclosure
869	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
669	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

700

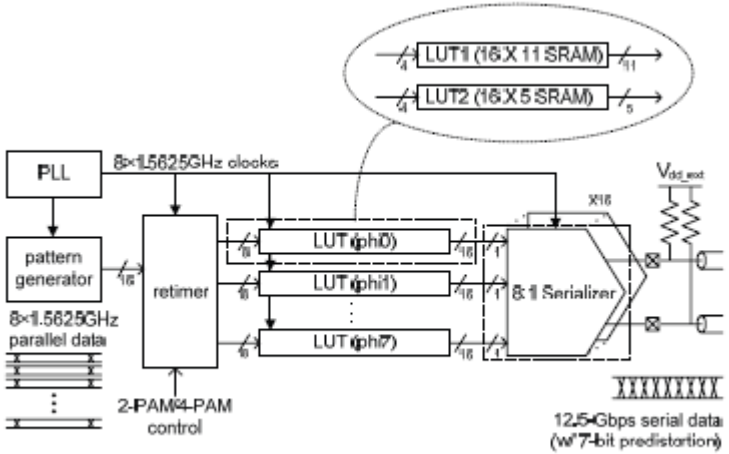
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

701

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

703

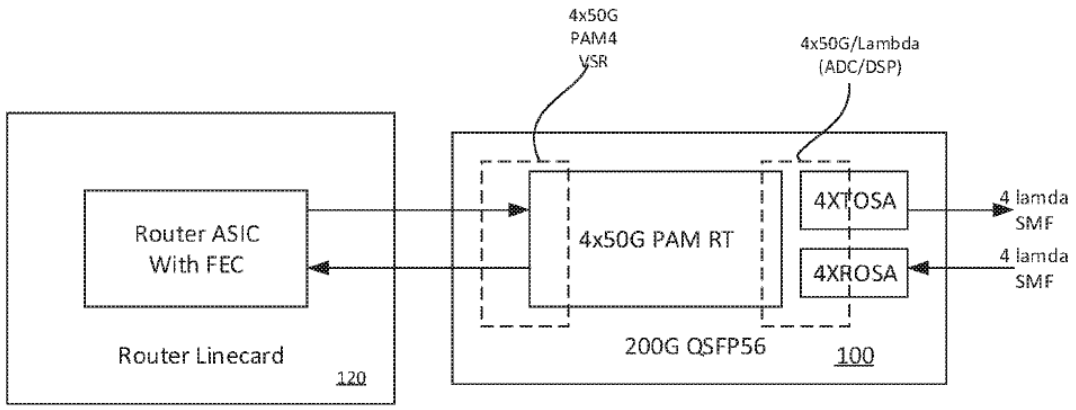
Claim 6	Disclosure
704	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

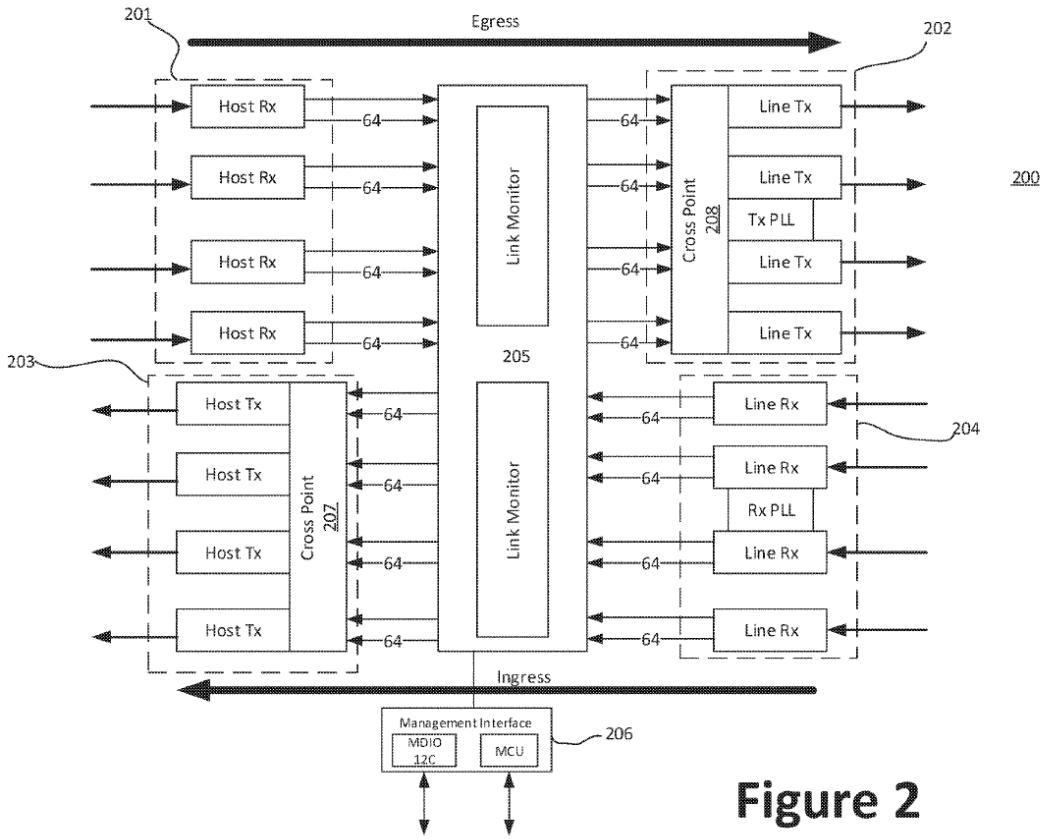
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

705

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

706

Claim 6	Disclosure
707	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
708	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

709

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

711

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

712

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

713

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

714

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

715

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

716

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

717

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

718

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

719

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

720

Claim 6	Disclosure
721	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

722

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

723

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

724

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

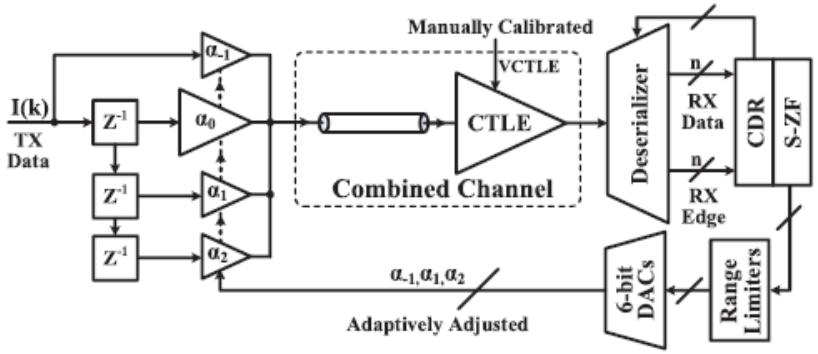


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

725

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

726

Claim 6	Disclosure

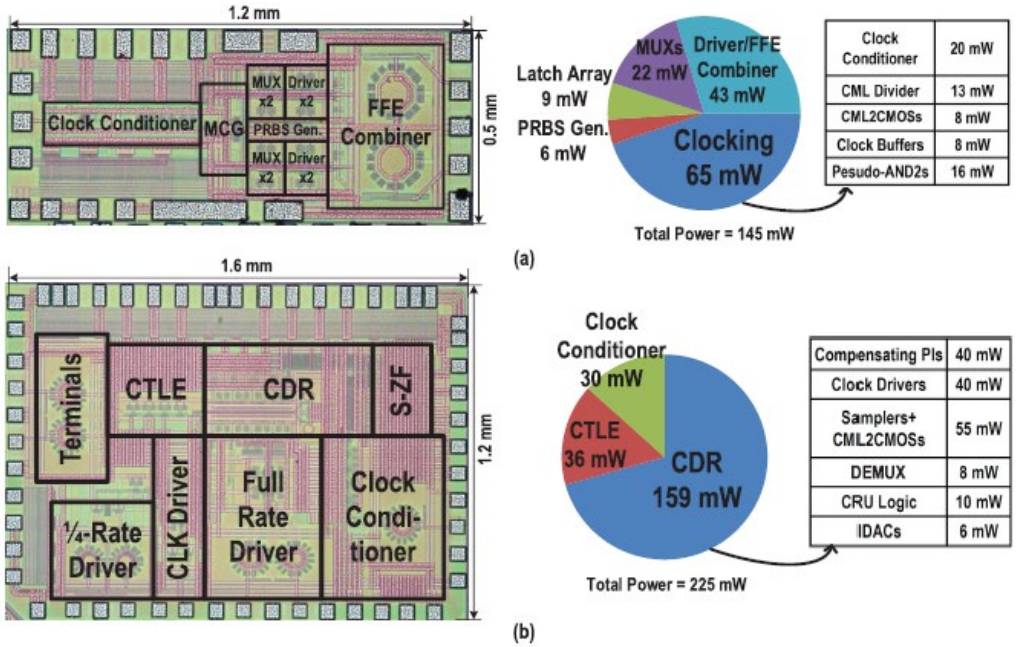
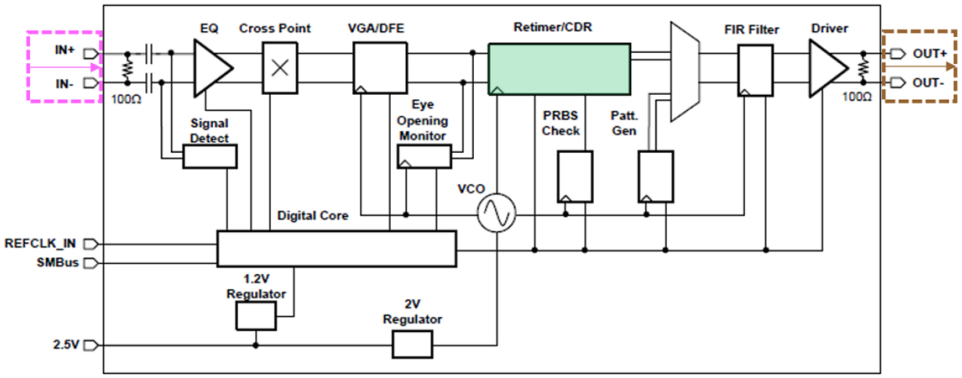


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

727

Claim 6	Disclosure
728	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

Claim 6	Disclosure
	Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 retimer performs CDR for respective transit signals conveying transit data streams received from electrical conductors to generate output signals conveying outbound data streams that exit cable 10 to respective hosts, as shown in DS125DF1610 Fig. 2 below. For similar reasons to those discussed in claim limitation 1[c], an ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data as the outbound data after extracting the input signal's data and generating a newly-timed signal that is a "re-modulation" of the data.  **Figure 2. DS125DF1610 Simplified Data Path Diagram**
6[h] re-modulating the first transit data stream as a second outbound data stream conveyed by a second electrical output signal to the second host	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

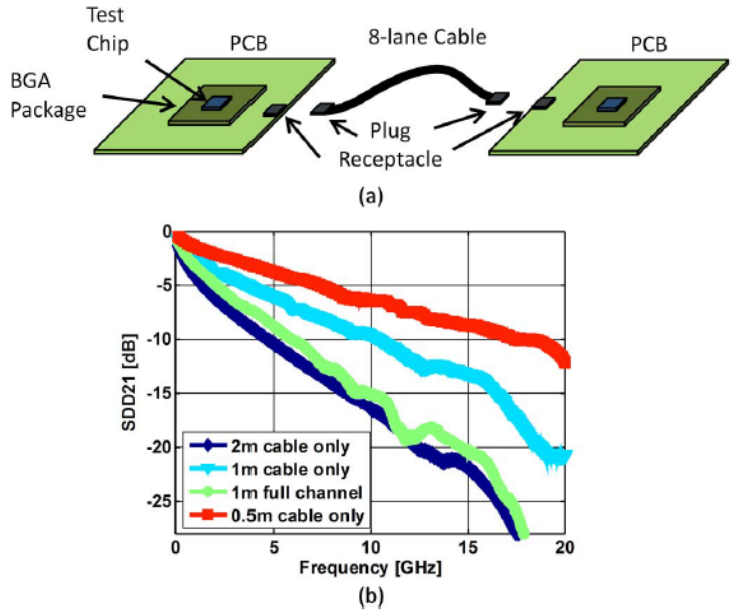
729

Claim 6	Disclosure
device;	*Abstract*—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

730

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

731

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

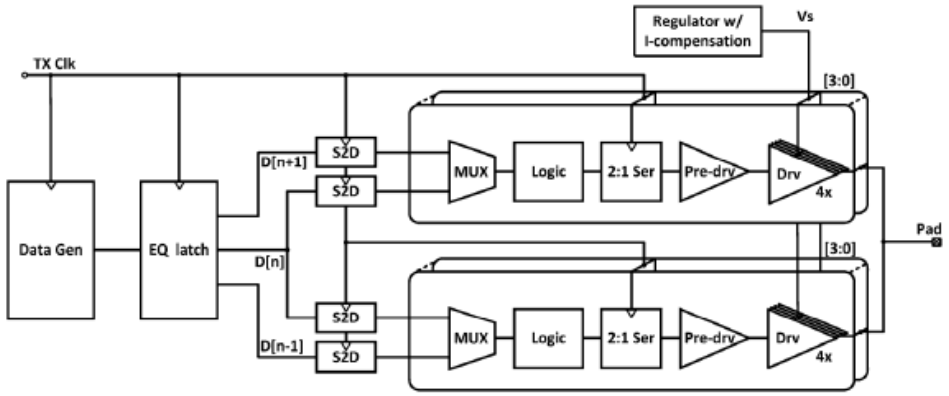
732

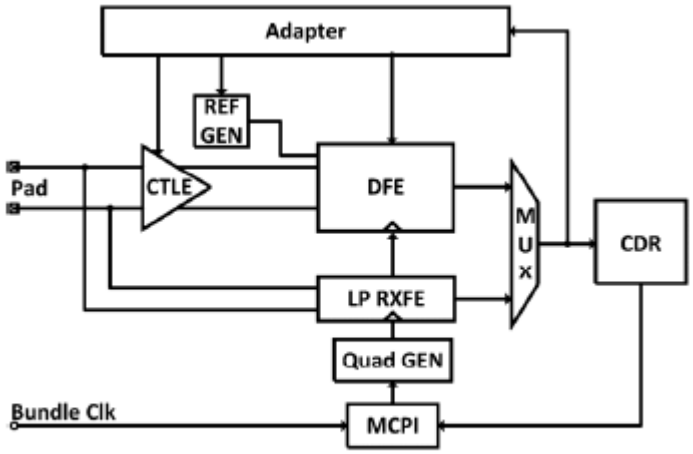
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

733

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

734

Claim 6	Disclosure
735	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

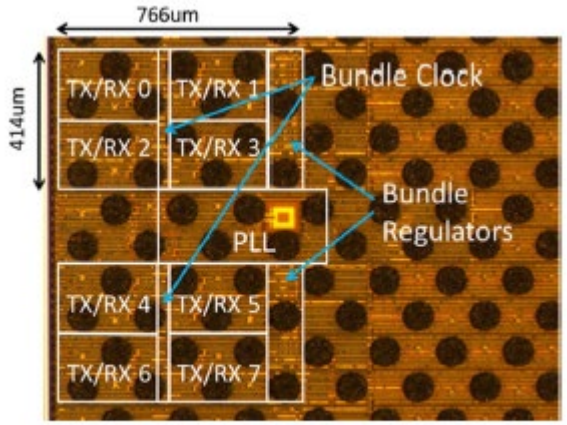
Claim 6	Disclosure
736	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

737

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

738

Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

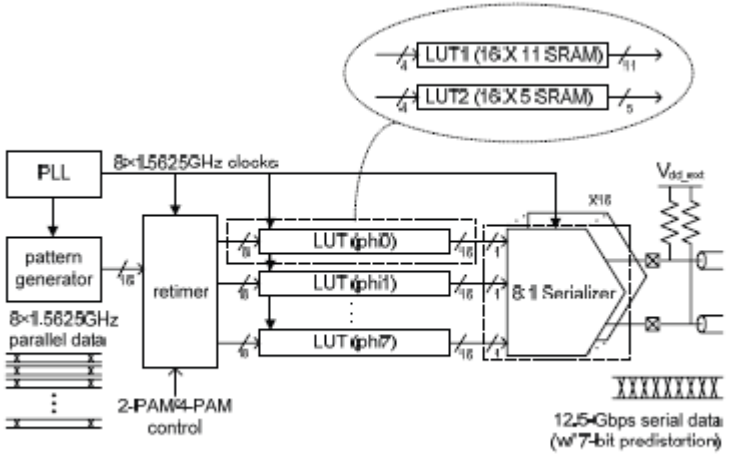
739

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

740

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

741

Claim 6	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

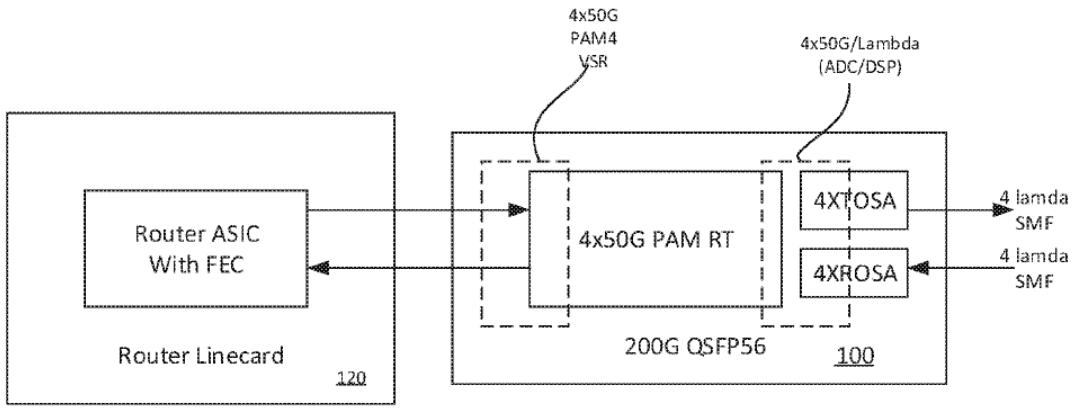
742

Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

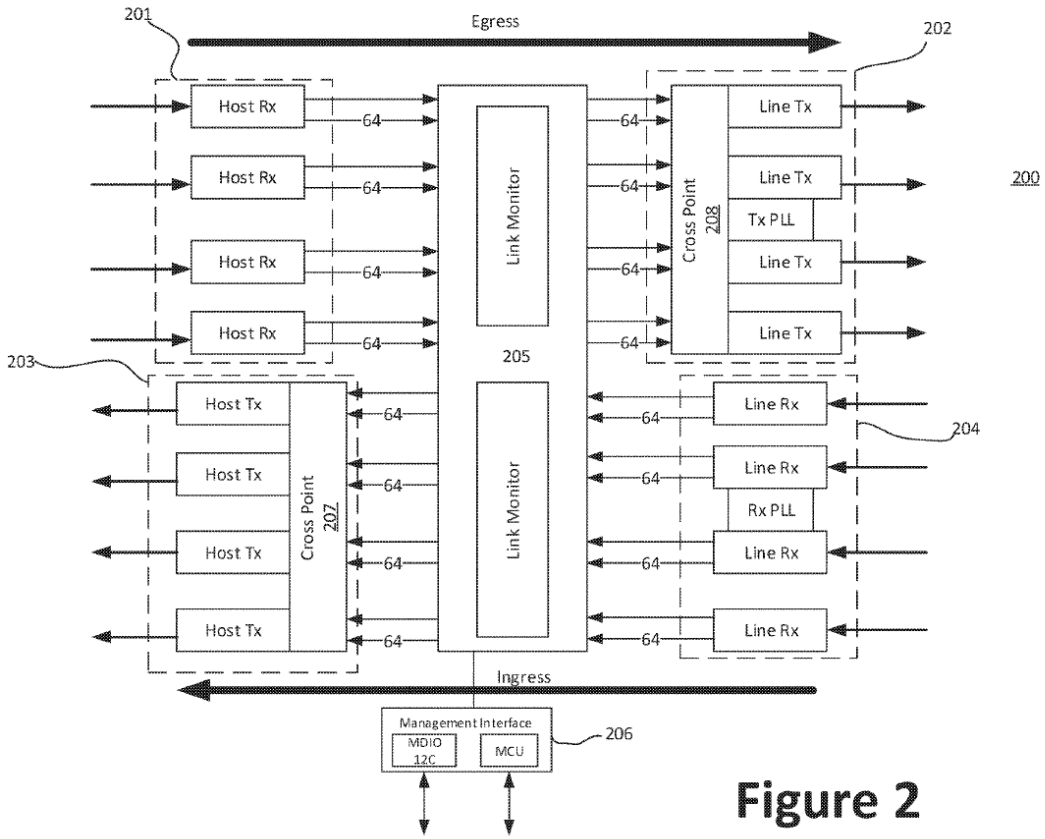
743

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

744

Claim 6	Disclosure
	 **Figure 1** Riani, Figure 1.

745

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

747

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

748

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

749

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

750

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

751

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

752

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

753

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

754

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

757

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

758

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

759

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

760

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

761

Claim 6	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

762

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

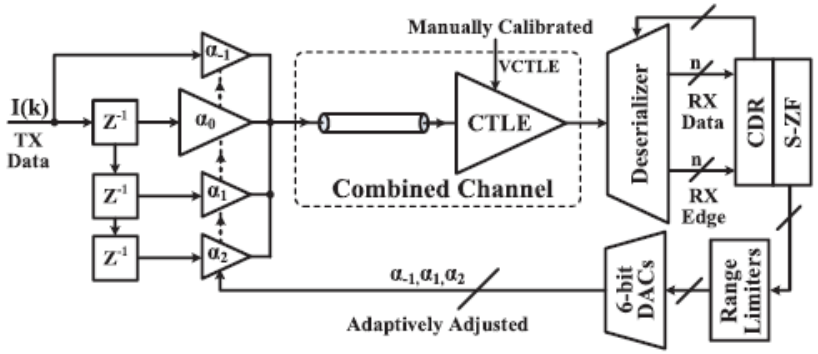


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

763

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

764

Claim 6	Disclosure

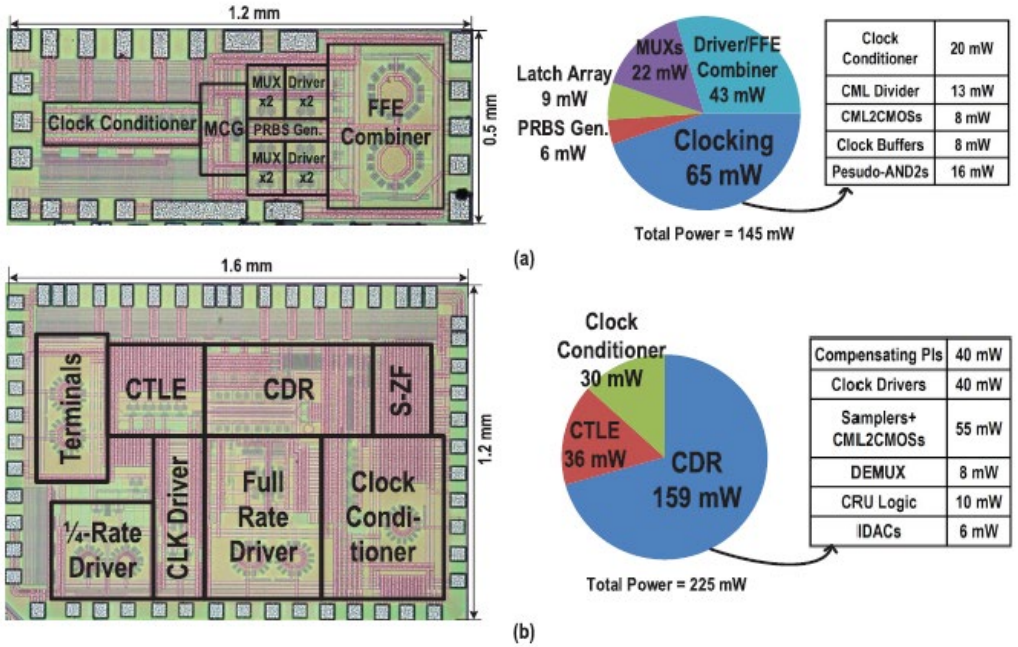
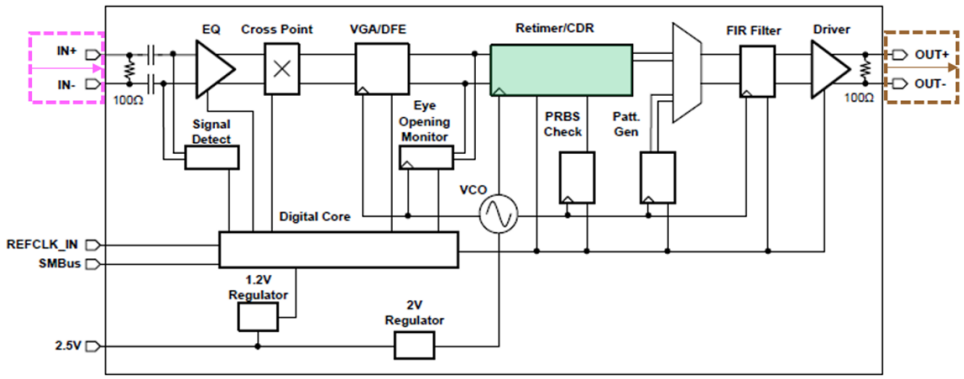


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

765

Claim 6	Disclosure
766	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

Claim 6	Disclosure
	Kobayashi, 4:4-22. <u>*DS125DF1610 Retimer Publication*</u> The DS125DF1610 retimer performs CDR for respective transit signals conveying transit data streams received from electrical conductors to generate output signals conveying outbound data streams that exit cable 10 to respective hosts, as shown in DS125DF1610 Fig. 2 below. For similar reasons to those discussed in claim limitation 1[c], an ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data as the outbound data after extracting the input signal's data and generating a newly-timed signal that is a "re-modulation" of the data.  **Figure 2. DS125DF1610 Simplified Data Path Diagram**
6[i] performing clock and data recovery on the second electrical transit signal with the first transceiver to extract the	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

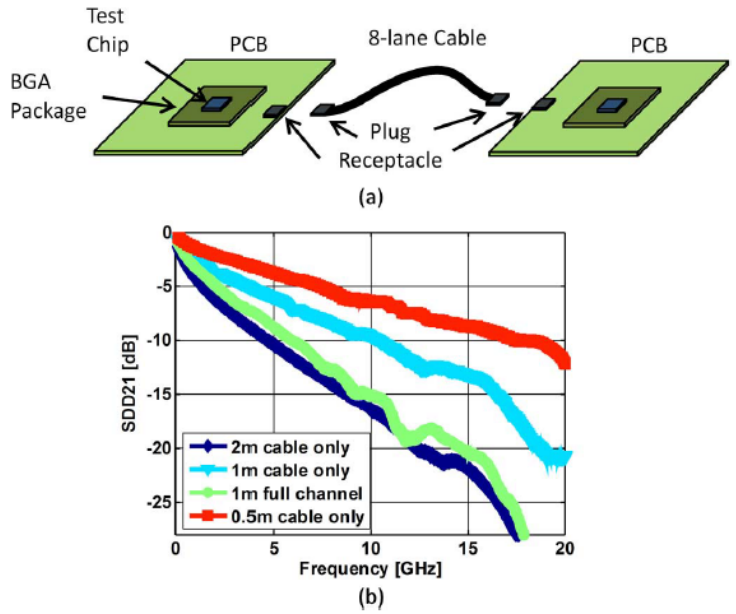
767

Claim 6	Disclosure
second transit data stream; and	*Abstract*—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

768

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

769

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

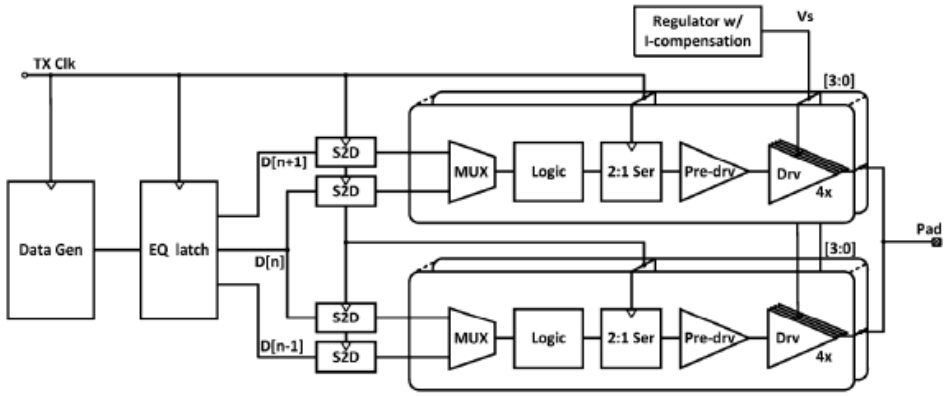
770

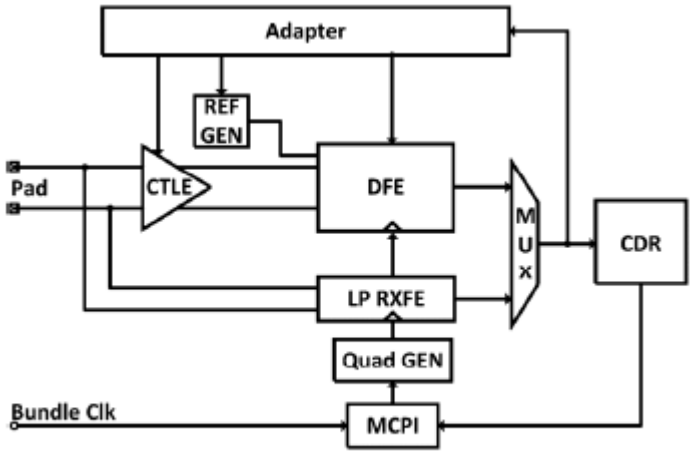
Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

771

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

772

Claim 6	Disclosure
773	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

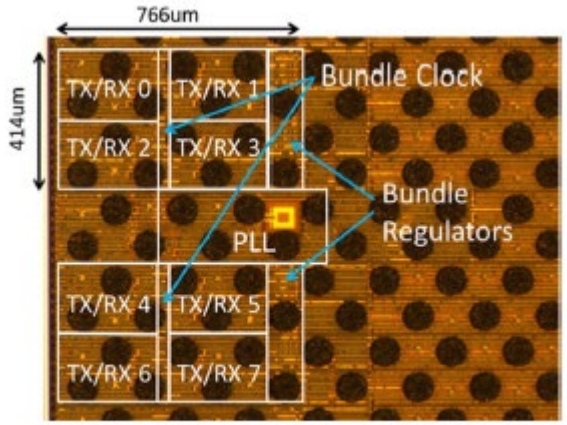
774

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

775

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

776

Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

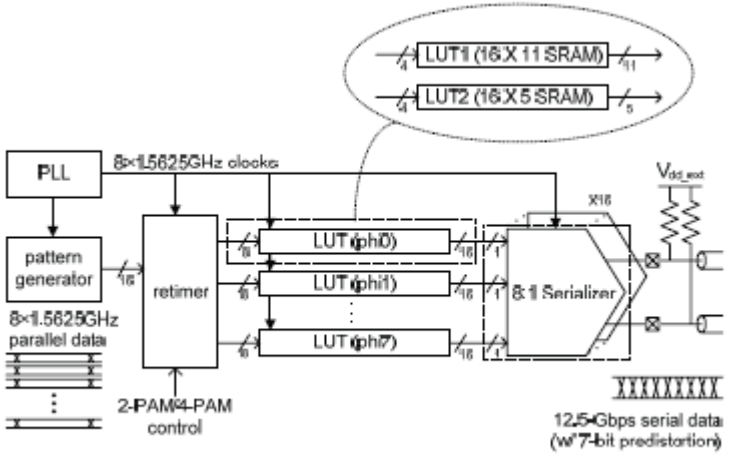
777

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

778

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

779

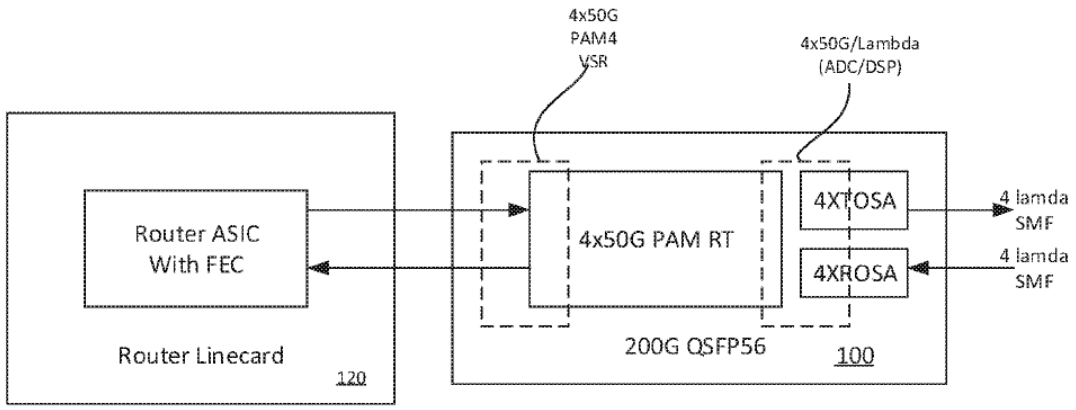
Claim 6	Disclosure
780	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

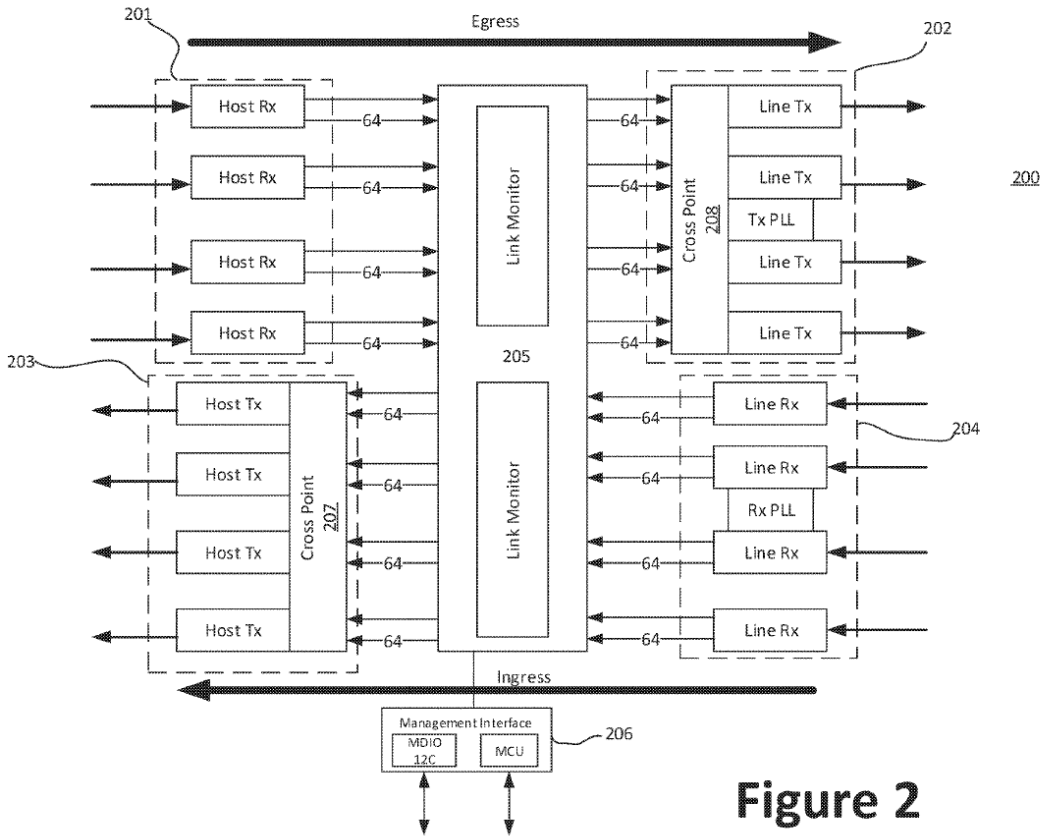
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

781

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

782

Claim 6	Disclosure
783	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

785

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

786

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

787

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

788

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

789

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

790

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

791

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

792

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

794

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

795

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

796

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

797

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

798

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

799

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

800

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

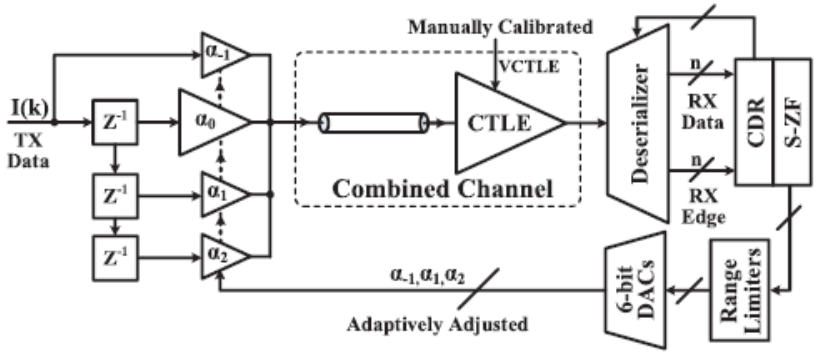


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

801

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

802

Claim 6	Disclosure

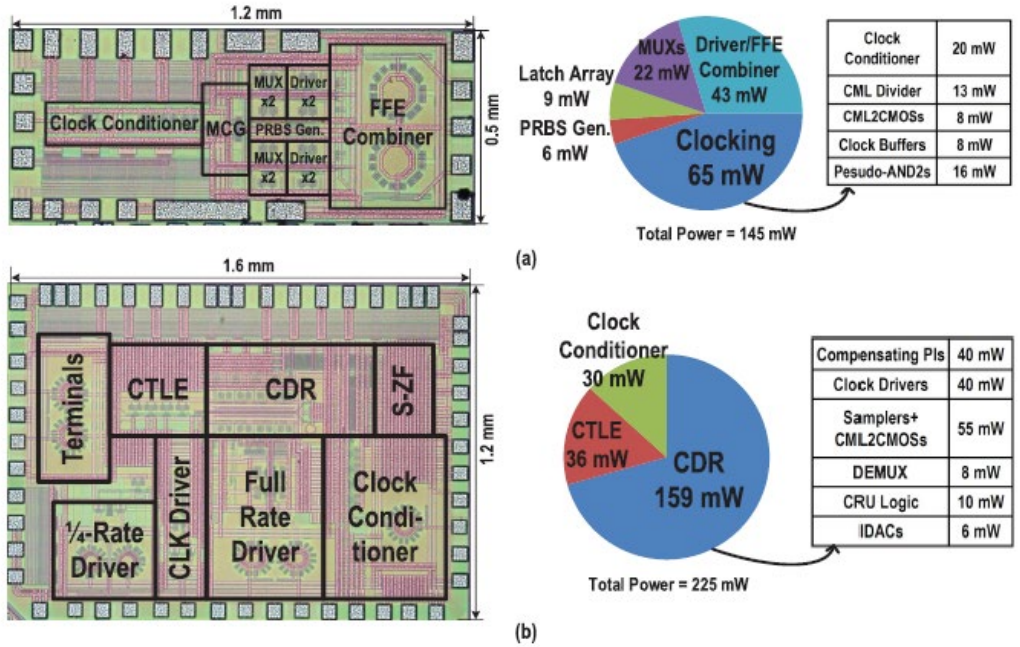


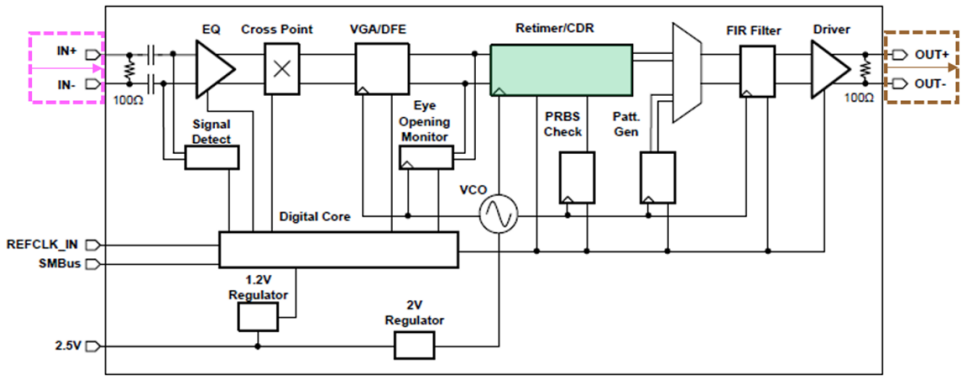
Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

803

Claim 6	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

804

Claim 6	Disclosure
	Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 retimer performs CDR for respective transit signals conveying transit data streams received from electrical conductors to generate output signals conveying outbound data streams that exit cable 10 to respective hosts, as shown in DS125DF1610 Fig. 2 below. For similar reasons to those discussed in claim limitation 1[c], an ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data as the outbound data after extracting the input signal's data and generating a newly-timed signal that is a "re-modulation" of the data.  **Figure 2. DS125DF1610 Simplified Data Path Diagram**
6[j] re-modulating the second transit data stream as a first outbound data stream conveyed by a first	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

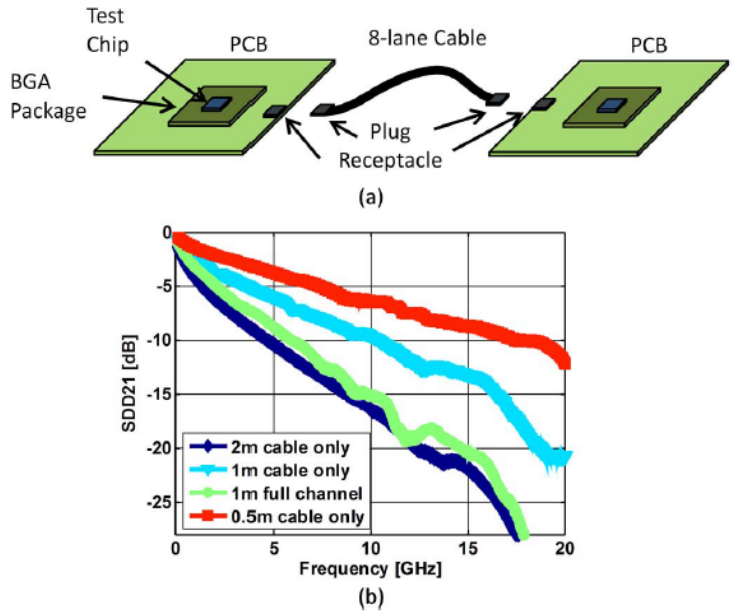
805

Claim 6	Disclosure
electrical output signal to the first host device,	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

808

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

807

Claim 6	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

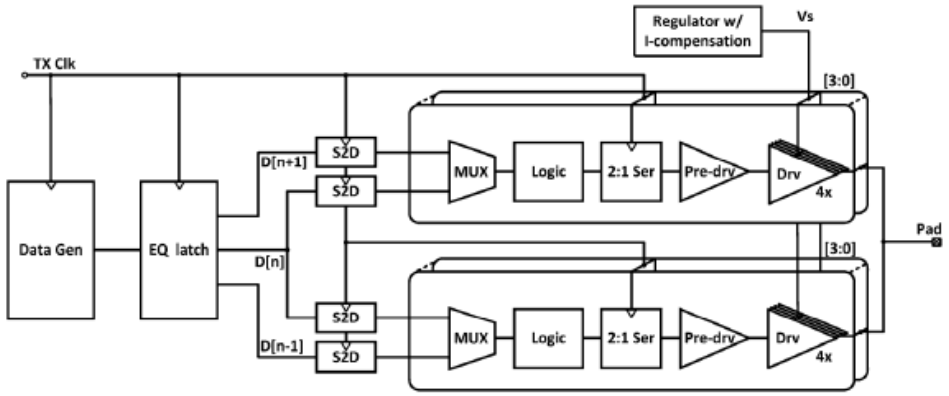
808

Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

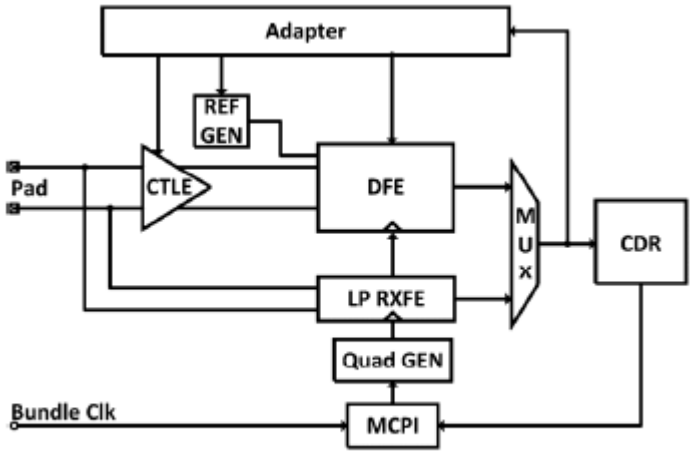
808

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

810

Claim 6	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

811

Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

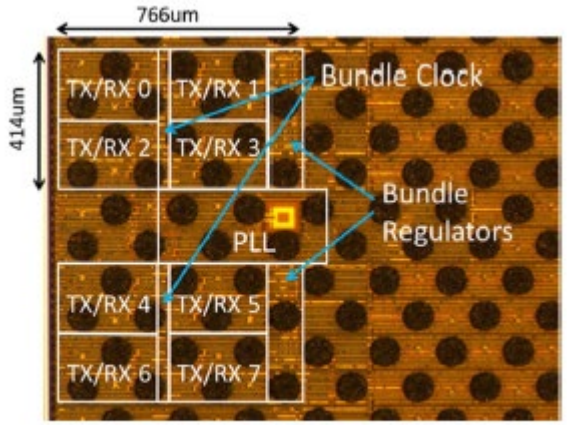
812

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

813

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

814

Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

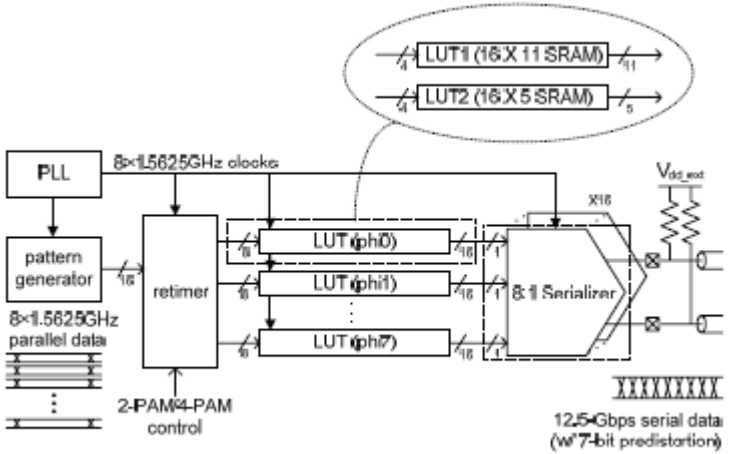
815

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

816

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

817

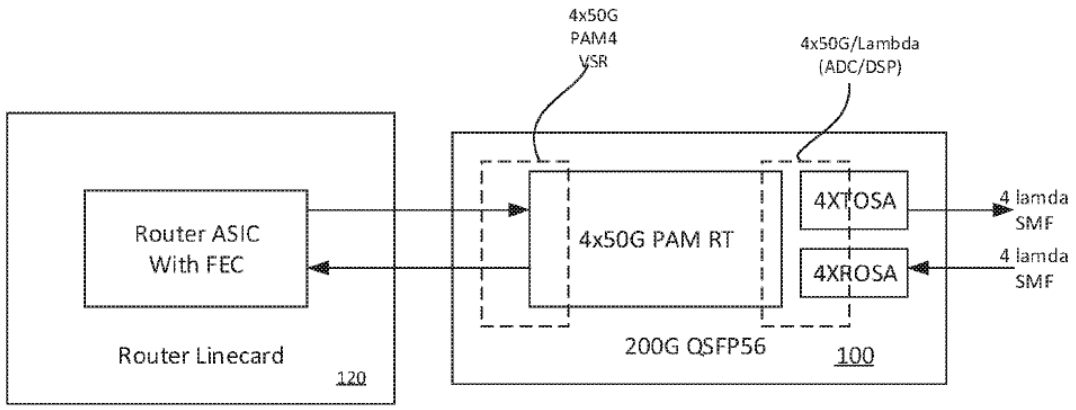
Claim 6	Disclosure
818	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

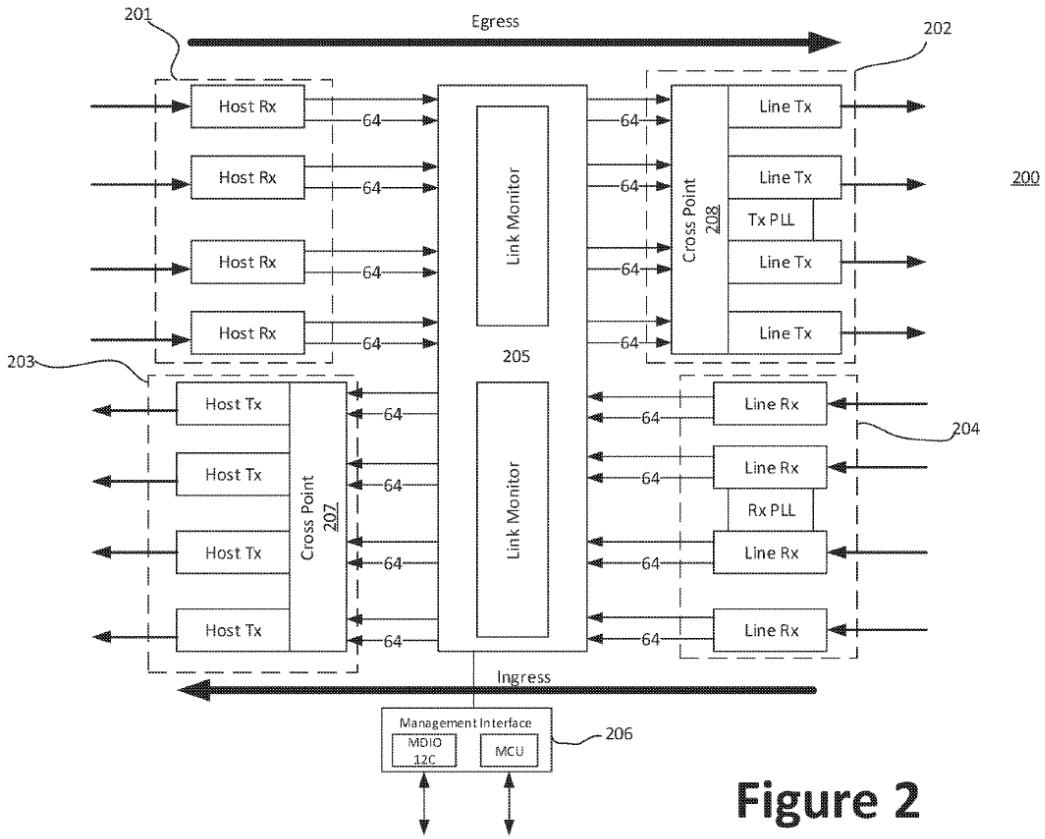
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

819

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

820

Claim 6	Disclosure
821	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

822

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

823

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

824

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

825

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

826

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

827

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

828

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

829

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

830

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

831

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

832

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

833

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

834

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

835

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

836

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

837

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

838

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

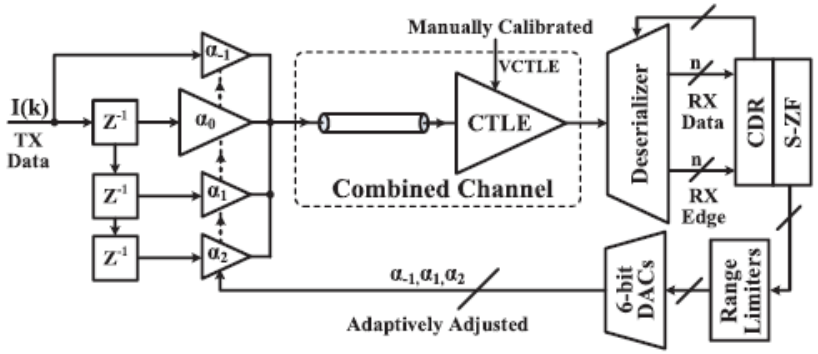


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

839

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

840

Claim 6	Disclosure

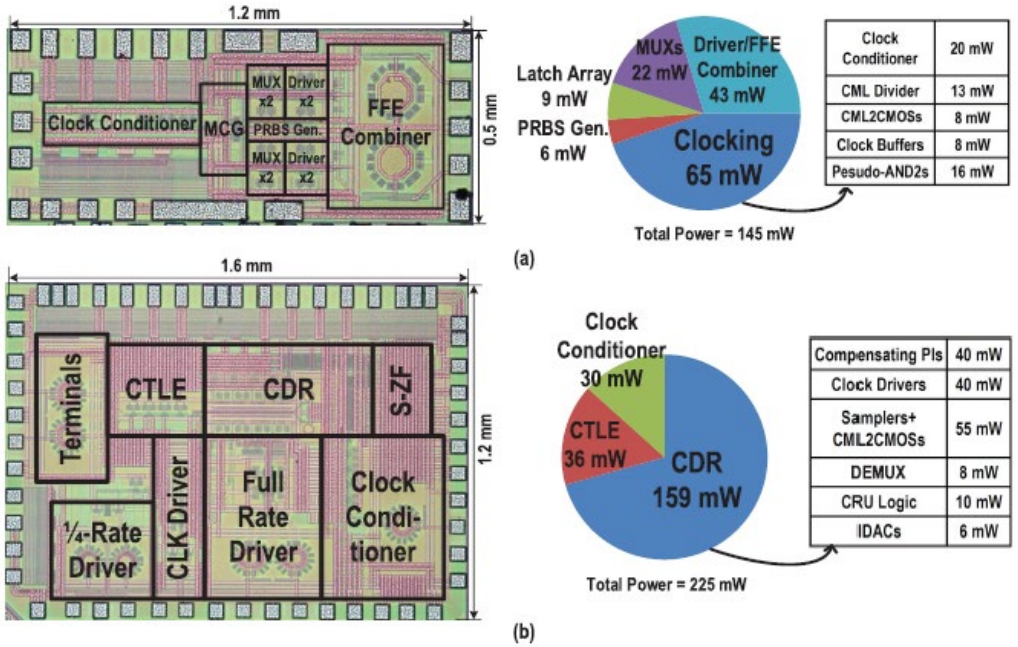


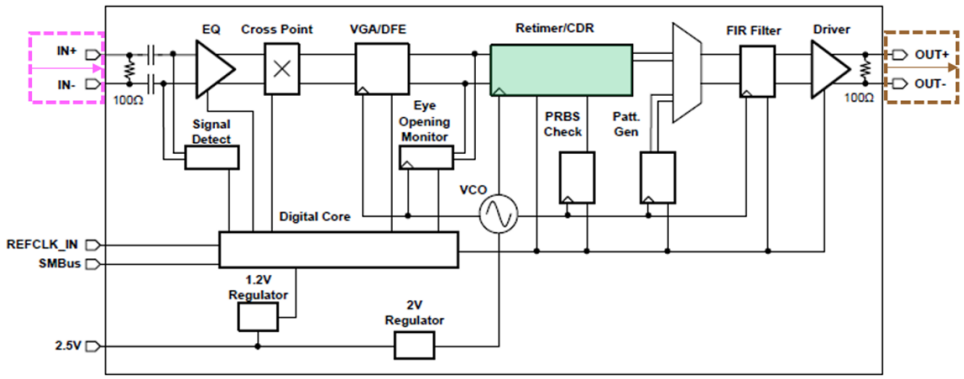
Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

841

Claim 6	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

842

Claim 6	Disclosure
	Kobayashi, 4:4-22. *DS125DF1610 Retimer Publication* The DS125DF1610 retimer performs CDR for respective transit signals conveying transit data streams received from electrical conductors to generate output signals conveying outbound data streams that exit cable 10 to respective hosts, as shown in DS125DF1610 Fig. 2 below. For similar reasons to those discussed in claim limitation 1[c], an ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data as the outbound data after extracting the input signal's data and generating a newly-timed signal that is a "re-modulation" of the data.  **Figure 2. DS125DF1610 Simplified Data Path Diagram**
6[k] wherein said re-modulating the first transit data stream, said remodulating the second transit data stream, said	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

Claim 6	Disclosure
performing clock and data recovery on the first electrical input signal, and said performing clock and data recovery on the second electrical input signal, each employ fixed, cable-independent, equalization parameters.	Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

844

Claim 6	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

845

Claim 6	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

846

Claim 6	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

847

Claim 6	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

848

Claim 6	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

849

Claim 6	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value | Bit pattern | Output Amplitude 0 | 0000 | 100-400 mV (P-P) 1 | 0001 | 300-600 mV (P-P) 2 | 0010 | 400-800 mV (P-P) 3 | 0011 | 600-1200 mV (P-P) 4-14 | 0100-1110 | Reserved 15 | 1111 | Custom CMIS, 41.

850

Claim 6	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

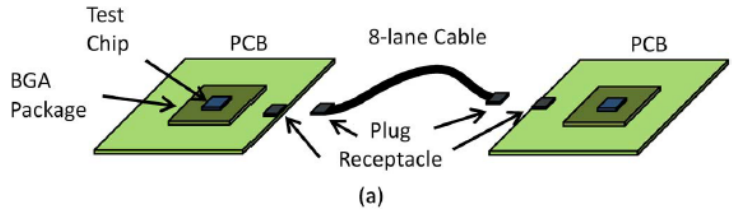
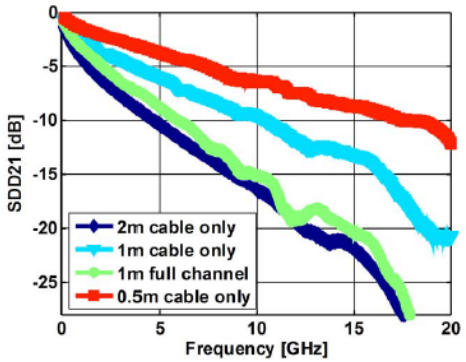
851

Claim 6	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

852

Claim 6	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

853

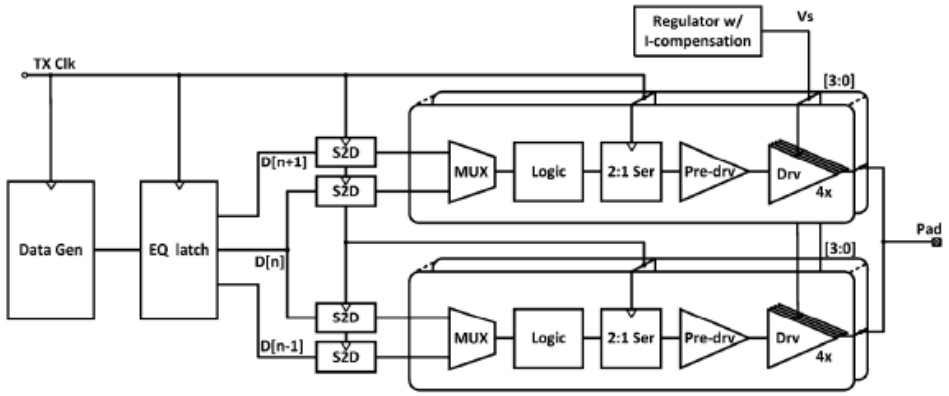
Claim 6	Disclosure
854	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 6	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

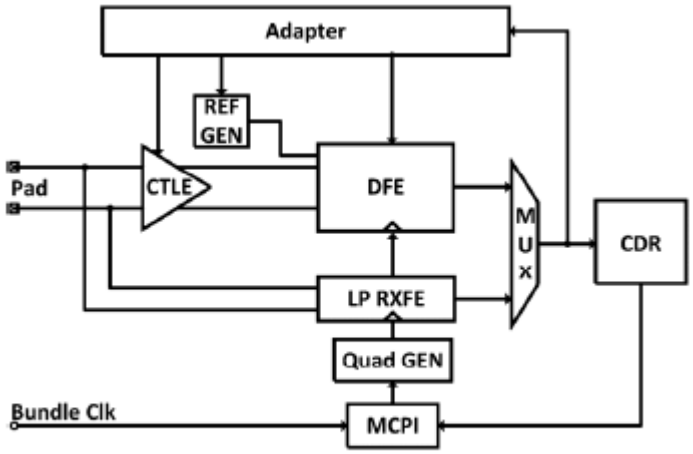
855

Claim 6	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

856

Claim 6	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

857

Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

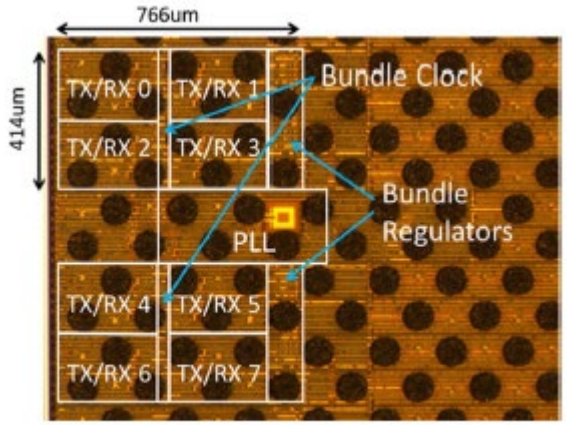
858

Claim 6	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

859

Claim 6	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

860

Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

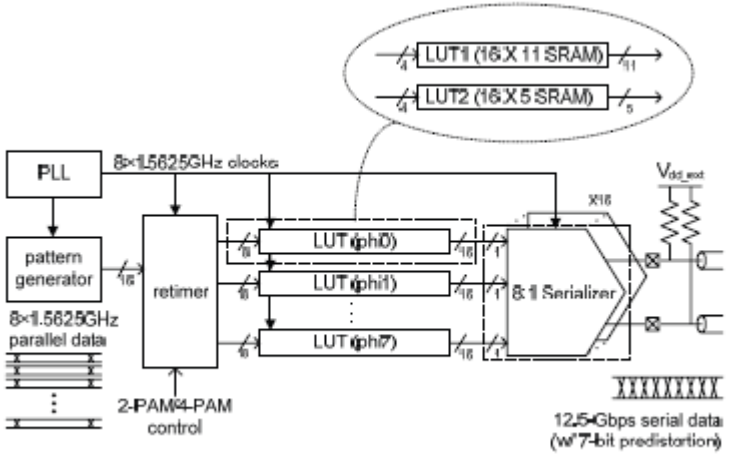
861

Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

862

Claim 6	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

863

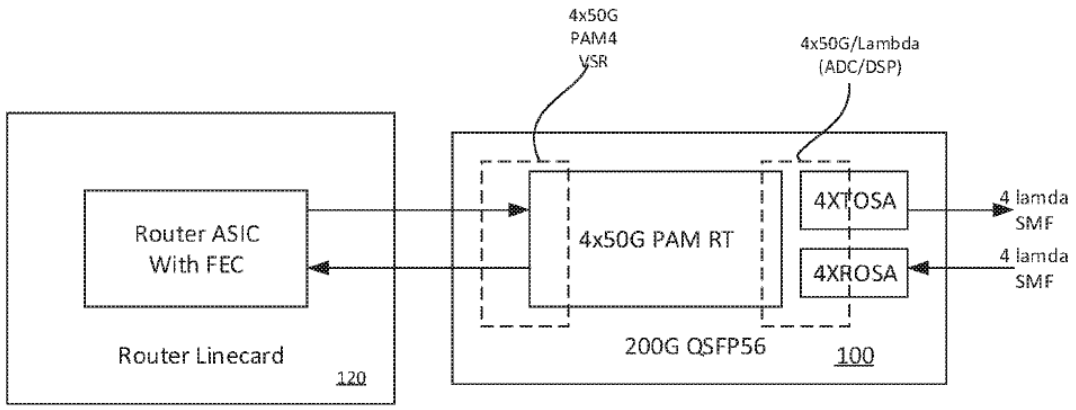
Claim 6	Disclosure
864	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

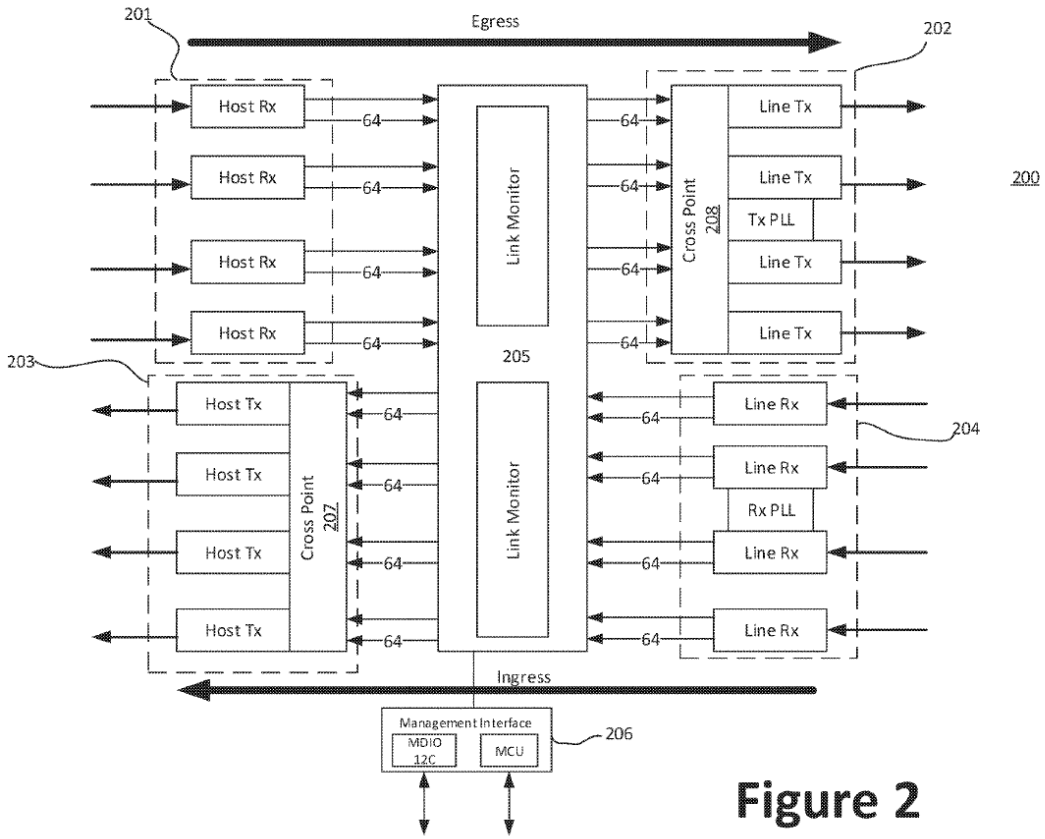
Claim 6	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

865

Claim 6	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

866

Claim 6	Disclosure
867	 **Figure 1** Riani, Figure 1.

Claim 6	Disclosure
	 Figure 2 Riani, Figure 2.

868

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

869

Claim 6	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

870

Claim 6	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

871

Claim 6	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

872

Claim 6	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

873

Claim 6	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

874

Claim 6	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

875

Claim 6	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

876

Claim 6	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

877

Claim 6	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

878

Claim 6	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

879

Claim 6	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

880

Claim 6	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

881

Claim 6	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

882

Claim 6	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

883

Claim 6	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

884

Claim 6	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

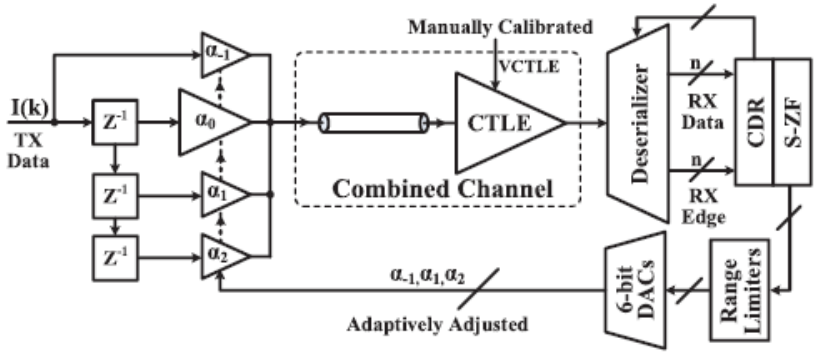


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

885

Claim 6	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

886

Claim 6	Disclosure

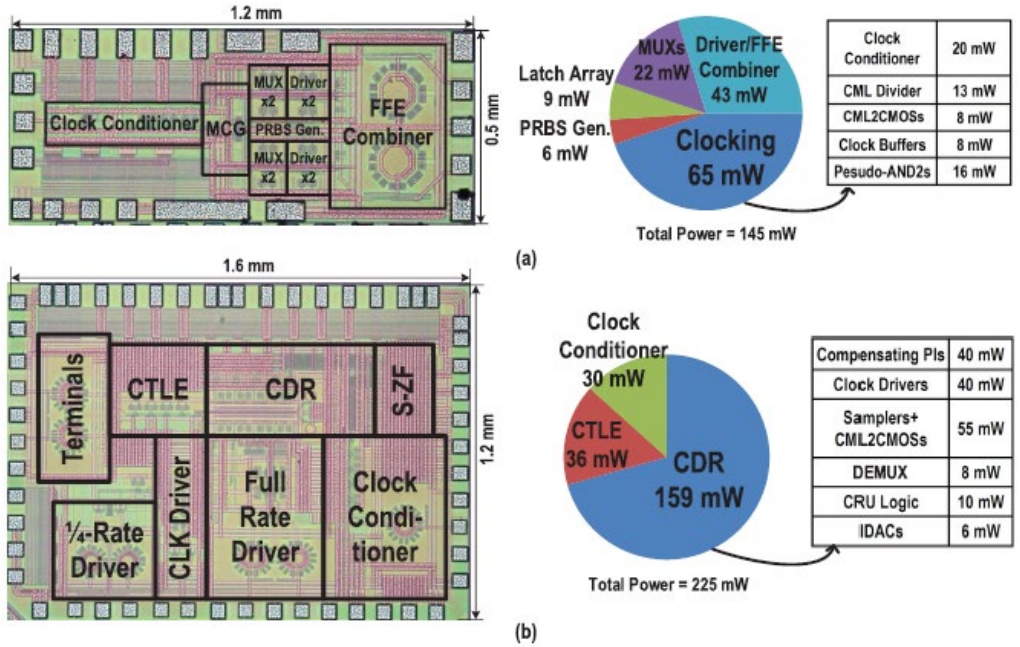


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

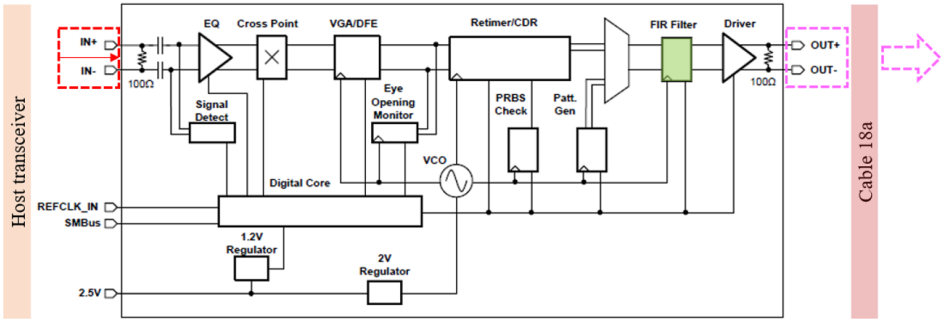
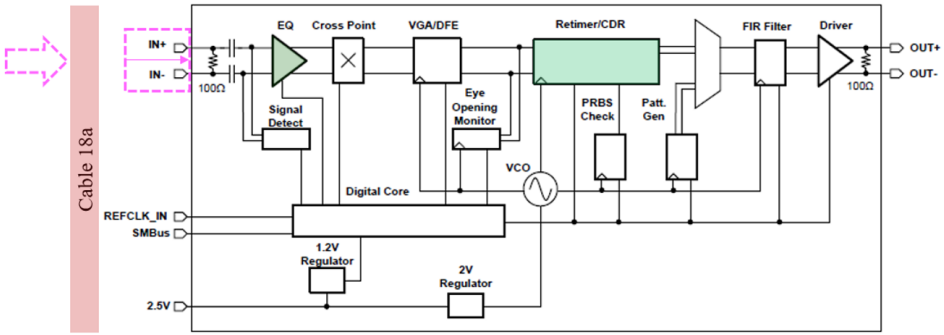
Zheng, 2973.

887

Claim 6	Disclosure
888	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The active cable 3 receives the serial bit stream transmitted from the upstream device 2 in the upstream facing RX 30a in one retimer 30. The retimer 30 performs a waveform shaping on the serial bit stream received by the upstream facing RX 30a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 30b to the other retime 31. The retimer 31 receives the serial bit stream transmitted from the retimer 30 in the upstream facing RX 31a. The retimer 31 performs a waveform shaping on the serial bit stream received by the upstream facing RX 31a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 31b to the downstream device 4." See column 3, lines 5-19. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing

Claim 6	Disclosure
	TX 52 includes a serializer and a driver." Kobayashi, 4:4-9. "The exchange module 202 is configured to communicate capabilities information with a remote link partner including host loss information, RS-FEC ability information, and RS-FEC request information. The exchange module 202 may transmit local host loss information to the link partner and receive remote host loss information from the link partner. In some embodiments, the exchange module 202 may be included in or otherwise established by the autonegotiation module 104 and thus the capabilities information may be exchanged during execution of an autonegotiation protocol. In some embodiments, the exchange module 202 may be included in or otherwise established by the PMD/PMA module 106 and thus the capabilities information may be exchanged during a startup protocol performed by the PMD/PMA module 106." Lusted, paragraph 23. "The autonegotiation module 104 is configured to auto-negotiate line transmission speed, mode of operation, and other communication parameters with a link partner when the communication lane 102 is brought up. Additionally, in some embodiments the autonegotiation module 104 may be configured to exchange capabilities information with the link partner including host insertion loss information, RS-FEC ability information, and RS-FEC request information. The autonegotiation module 104 may encode the capabilities information, for example, in an autonegotiation base page or in a next page. Although illustrated in FIG. 1 as coupled between the PMD/PMA 106 and the communication lanes 102, in other embodiments the autonegotiation module 104 may be included at different positions in the Ethernet port logic 100. For example, in some embodiments the autonegotiation module 104 may be coupled between the PMD/PMA 106 and the FEC 108." Lusted, paragraph 17. "In view of the above transmission characteristics of twisted-pair, current Ethernet chips are usually equipped with preemphasis technique of a transmitting end and equalization technique of a receiving end to compensate the difference in transmission attenuation of physical layer signals on different frequencies on the same medium during Ethernet baseband transmission. Main functions of the preemphasis technique of the transmitting end and the equalization technique of the receiving end include: amplifying signals on different frequencies differently at the transmitting end to compensate different transmission attenuation in

Claim 6	Disclosure
	the twisted pair; applying the equalization technique to different frequency components in signals received at the receiving end to compensate the attenuation." Yu, paragraph 34. "An embodiment of the present invention provides an equalizer setting apparatus, which uses a filter circuit 101, an analog-to-digital converter 102, a digital memory 103, a digital comparator 104 and a first digital accumulator 105, the filter circuit 101, the analog-to-digital converter 102, the digital memory 103, the digital comparator 104 and the first digital accumulator 105 are connected in sequence, the output end of the equalizer 201 is connected with the input end of the filter circuit 101, the output end of the first digital accumulator 105 is connected with the set value input end of the equalizer 201, the output end of the analog-to-digital converter 102 is connected with the input end of the digital comparator 104, so that the problems that the setting value of the equalizer is accurately set in the prior art, which is time-consuming and labor-consuming are solved, the method has the advantages of achieving accurate setting of the set value of the equalizer in a single time, fast, simple, convenient, accurate and low-cost manner and saving labor force to a certain extent." Yin, paragraph 26 (English translation). *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit

Claim 6	Disclosure
	signal."  Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram *U.S. Patent No. 7,762,727 to Aronson* Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-

891

Claim 6	Disclosure
	pair or coaxial cables like Lugthart '414. *See, e.g., id*. at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id*. at 14:16-15:24.

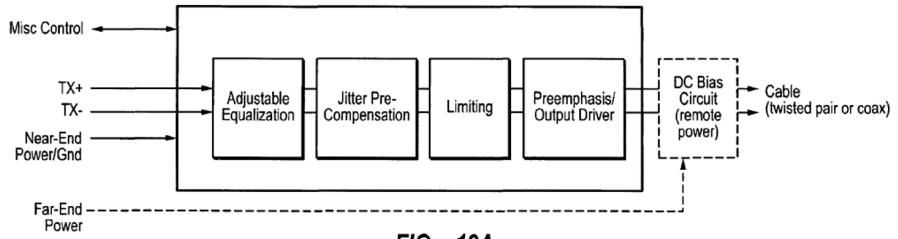


FIG. 12A

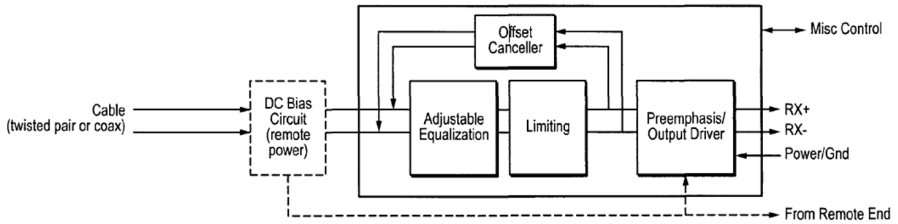


FIG. 12B

Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23.

Aronson discloses the host-side equalization and pre-equalization use fixed cable-independent parameters to address host-side losses. *See, e.g.*, Aronson at 12:10-17, 14:16-23, 15:17-24. "Fixed" parameters reduced complexity and manufacturing costs. See, e.g., U.S. Patent No. 7,239,665 at 1:7-18 (discussing slow

892

Claim 6	Disclosure
	convergence with adaptive equalization); U.S. Patent App. Pub. No. 2011/0228821 at [0009]-[0011] (discussing tradeoffs between fixed and adaptive equalization); U.S. Patent No. 6,621,862 at 1:52-55 (large gate count for adaptive equalization can cause "high chip costs"); U.S. Patent App. Pub. No. 2006/0045176 at [0034] ("The adaptive equalizer can be trapped in a failure state (i.e., incorrect state) and not able to recover from the failure state."); U.S. Patent App. Pub. No. 2017/0187463 at [0058] ("[E]qualization techniques… come with different engineering tradeoffs, including trade-offs in power consumption, performance, and/or cost."). **5. The PCIe 3.0 Equalization Adaptation Protocol** The PCIe 3.0 Base Specification [10] requires that a compliant Tx be able to produce any of the output equalization (TxEQ) levels shown in Figure 21, whenever requested to do so by the receiving port, during link equalization training. The same specification describes a protocol of TxEQ adaptation, where the host (upstream agent) or the end device (downstream agent) can each request that any of the TxEQ values in Figure 21 be presented to them by the other agent on the link. The intent of adaptation is to allow both agents to adjust the link partner's TxEQ to an optimal value for their receiver, and for the specific channel and operating conditions at hand. Samaan, p. 24.

893

G. DEPENDENT CLAIM 7

Claim 7	Disclosure
7[a] The communication method of claim 6, wherein cable-dependent equalization parameters are employed for at least one of: re-modulating the first inbound data stream,	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

894

Claim 7	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

895

Claim 7	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

968

Claim 7	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

897

Claim 7	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

898

Claim 7	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

668

Claim 7	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

900

Claim 7	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	RQD
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

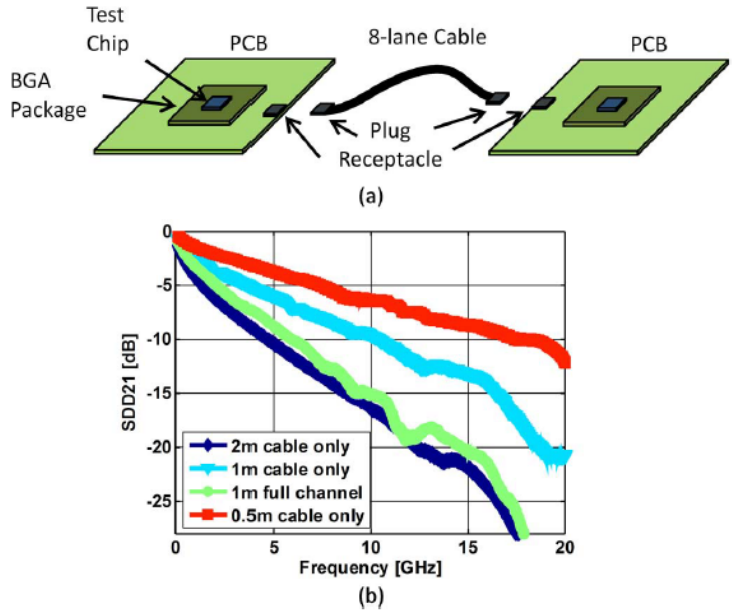
901

Claim 7	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

902

Claim 7	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

903

Claim 7	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

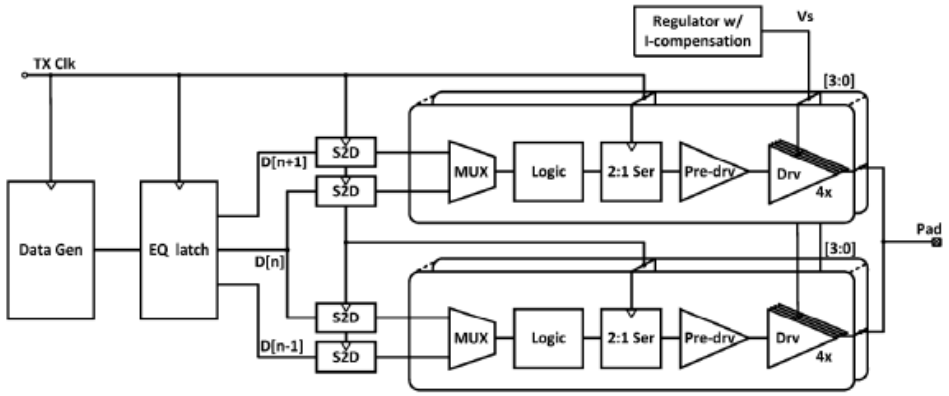
904

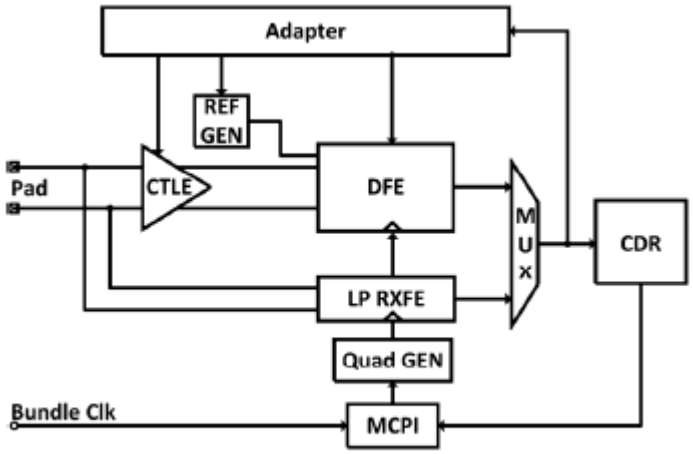
Claim 7	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

905

Claim 7	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

906

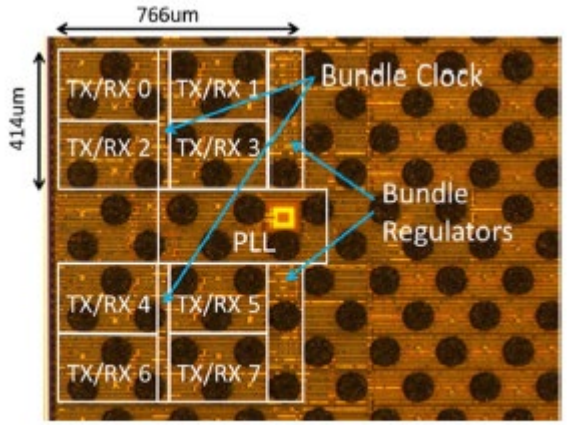
Claim 7	Disclosure
907	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 7	Disclosure
908	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 7	Disclosure
606	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 7	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

910

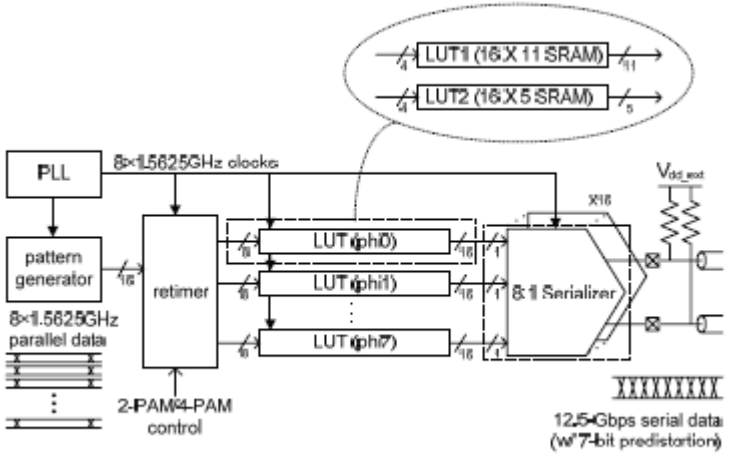
Claim 7	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 7	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

912

Claim 7	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

913

Claim 7	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

914

Claim 7	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

915

Claim 7	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

916

Claim 7	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

917

Claim 7	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

Claim 7	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 7	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

920

Claim 7	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

921

Claim 7	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

922

Claim 7	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

923

Claim 7	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

924

Claim 7	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

925

Claim 7	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

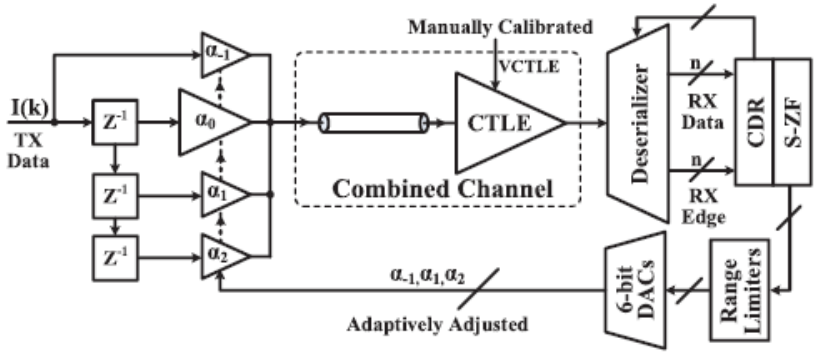


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

926

Claim 7	Disclosure
927	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

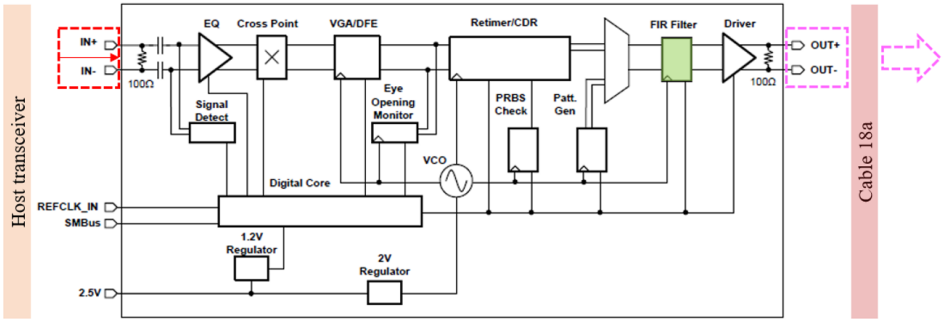
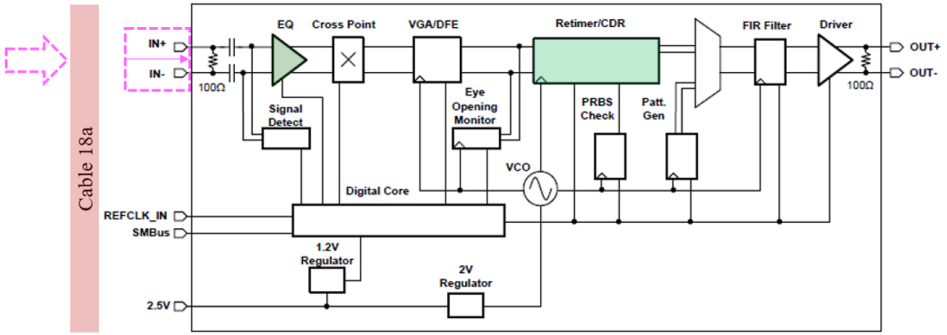
Claim 7	Disclosure
928	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The link activation module 204 is configured to determine a total insertion loss for the connection with the remote host based on the local host loss, the remote host loss, and cable loss. The link activation module 204 compares the total insertion loss to a FEC limit and to a specification limit. The link activation module 204 may bring up the communication link with the link partner without FEC enabled if the total insertion loss is below the specification limit and below the FEC limit. Additionally, the link activation module 204 may bring up the communication link with the link partner with FEC enabled if the total insertion loss is below the specification limit but not below the FEC limit. Further, the link activation module 204 may not bring up the communication link with the link partner if the total insertion loss is not below the specification limit (e.g., the link activation module 204 may keep the communication link down or bring the communication link down)." Lusted, paragraph 24.

Claim 7	Disclosure
	"The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. "FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to

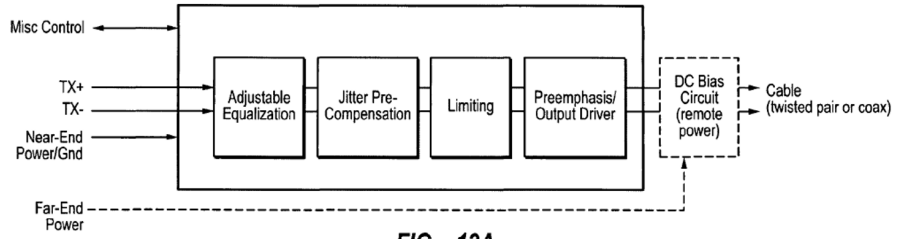
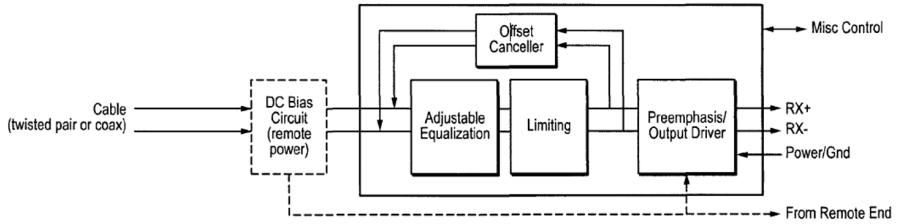
929

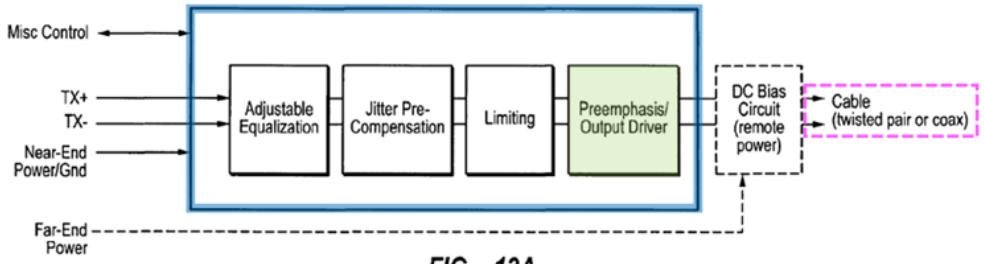
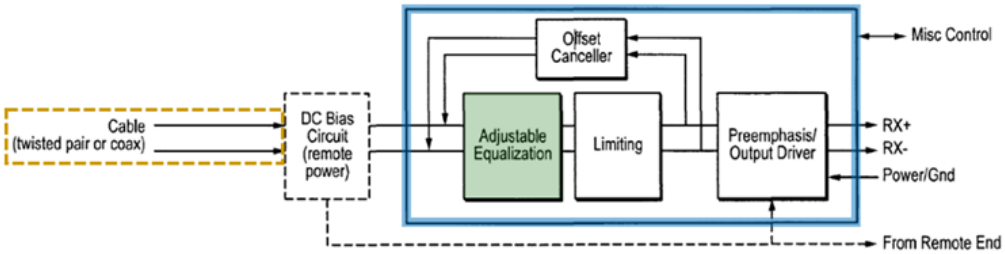
Claim 7	Disclosure
	the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."

930

Claim 7	Disclosure
	 Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram *U.S. Patent No. 7,762,727 to Aronson* Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id.* at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id.* at 14:16-15:24.

931

Claim 7	Disclosure
	 FIG. 12A  FIG. 12B Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23. Using "cable dependent" parameters for equalization would have been the conventional and obvious way to address losses over the cable. *See, e.g.*, Aronson at 14:33-43 (pre-emphasis/output driver performs pre-equalization to overcome "loss on the cable" using cable-dependent parameters that "match the particular length [and/or] the characteristics of the copper cabling"), Fig. 12A (below).

Claim 7	Disclosure
	 FIG. 12A Using "cable dependent" parameters for equalization would have been the conventional and obvious way to address losses over the cable. Aronson at 15:1-8 (equalization "compensate[s] for the cable" losses using cable dependent parameters based on "cable length and characteristics"), Fig. 12B (below).  FIG. 12B
7[b] re-modulating the second inbound data stream,	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

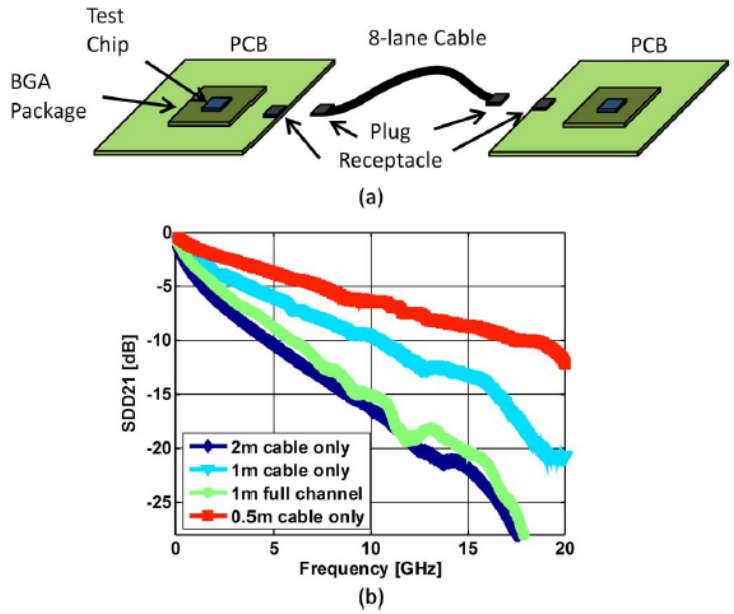
933

Claim 7	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

934

Claim 7	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

935

Claim 7	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

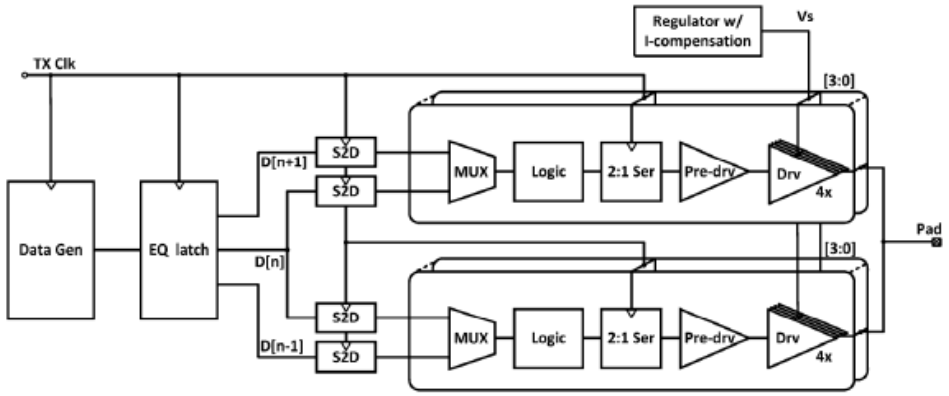
936

Claim 7	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

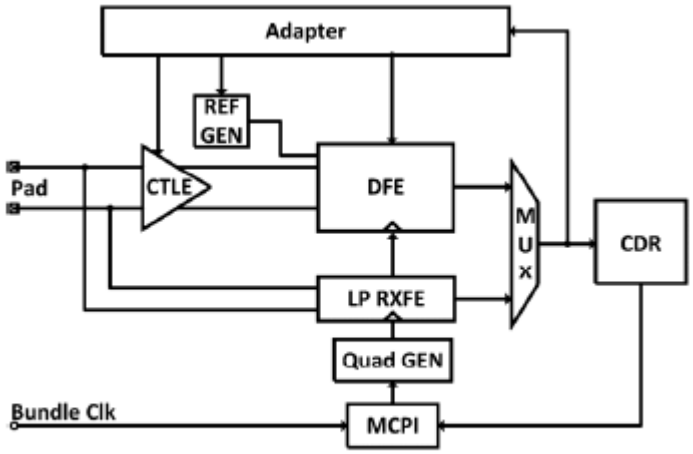
937

Claim 7	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

938

Claim 7	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

939

Claim 7	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

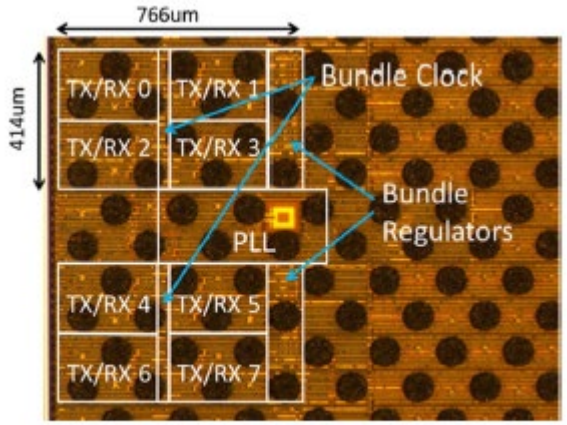
940

Claim 7	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

941

Claim 7	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

942

Claim 7	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

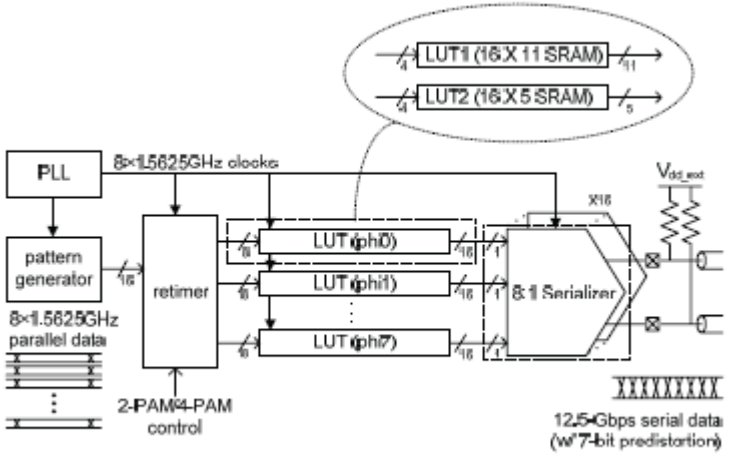
943

Claim 7	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

944

Claim 7	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

945

Claim 7	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

946

Claim 7	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

947

Claim 7	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

948

Claim 7	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

Claim 7	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

950

Claim 7	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

951

Claim 7	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

952

Claim 7	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

953

Claim 7	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/ LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

954

Claim 7	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

955

Claim 7	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

956

Claim 7	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

957

Claim 7	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

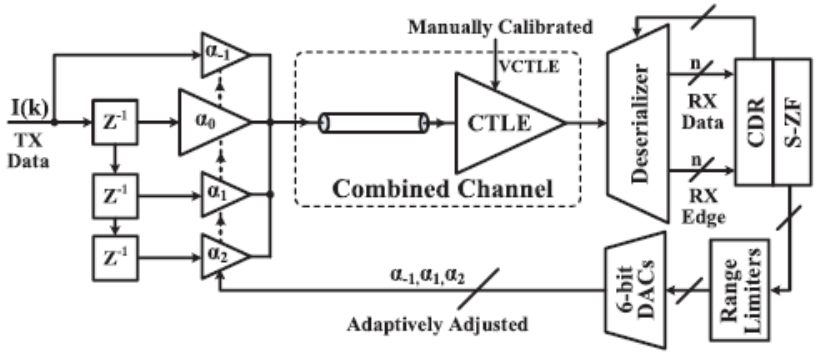


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

958

Claim 7	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

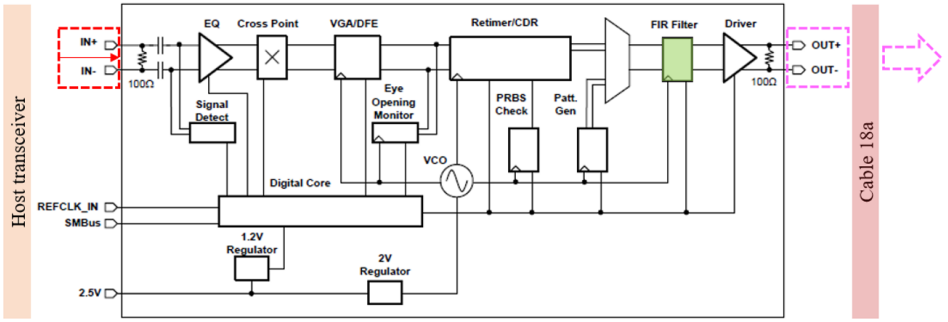
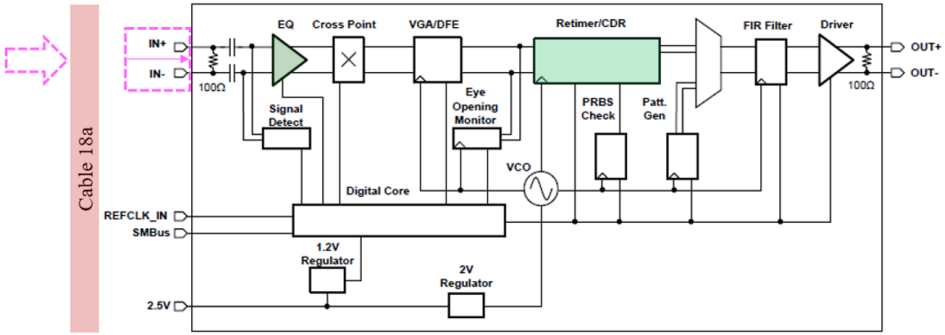
959

Claim 7	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20.

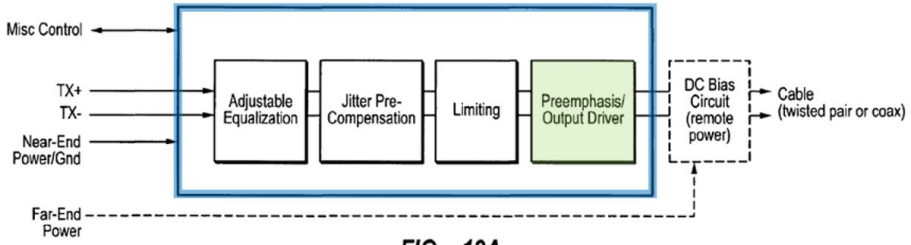
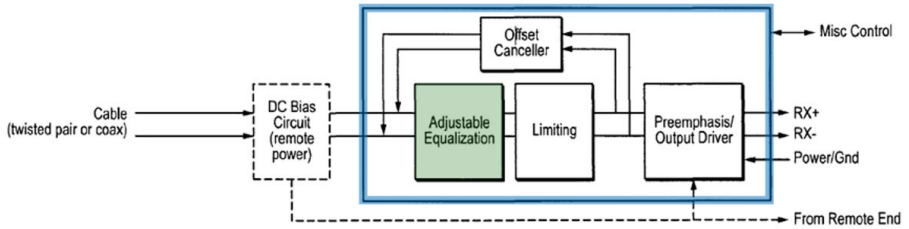
960

Claim 7	Disclosure
	"FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."

961

Claim 7	Disclosure
	 Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram *U.S. Patent No. 7,762,727 to Aronson* Aronson discloses cable-side equalization and pre-equalization that use cable-dependent parameters to address cable-side losses. *See, e.g.*, Aronson, 14:33-43, 15:1-8. Aronson's transmitter IC's line-side pre-emphasis/output driver overcomes "loss on the cable" using cable-dependent parameters that "match the particular length [and/or] the characteristics of the copper cabling." *See id*. at 14:33-43, Fig. 12A (below). Similarly, Aronson's receiver IC equalizes received line-side signals to "compensat[e] for the cable" losses using cable-dependent parameters based on "cable length and characteristics." *See id*. at 15:1-8, Fig. 12B

962

Claim 7	Disclosure
	(below).  FIG. 12A  FIG. 12B
7[c] said performing clock and data recovery on the first electrical transit signal, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

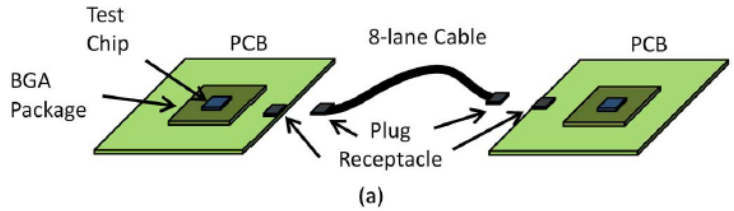
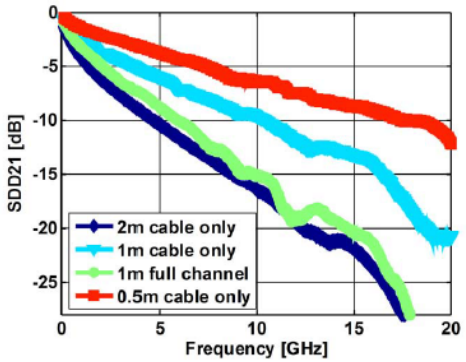
963

Claim 7	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

964

Claim 7	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

965

Claim 7	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

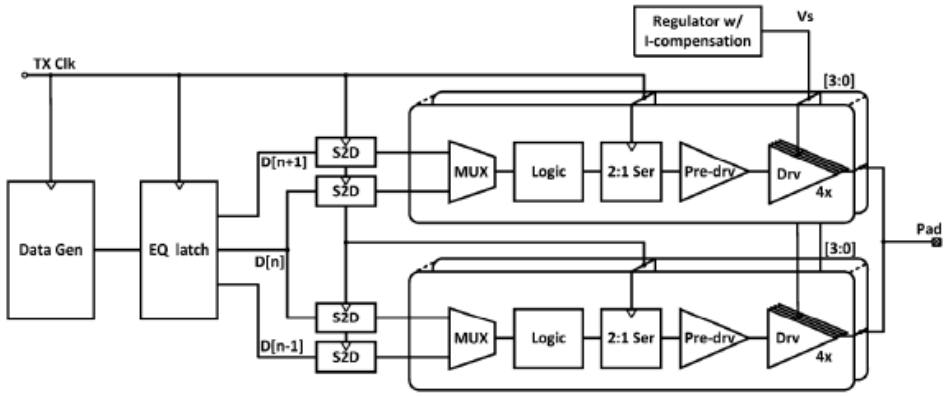
996

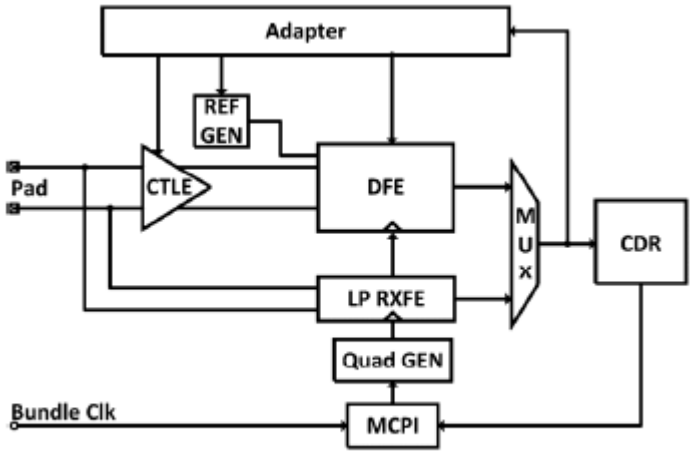
Claim 7	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

967

Claim 7	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

968

Claim 7	Disclosure
696	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

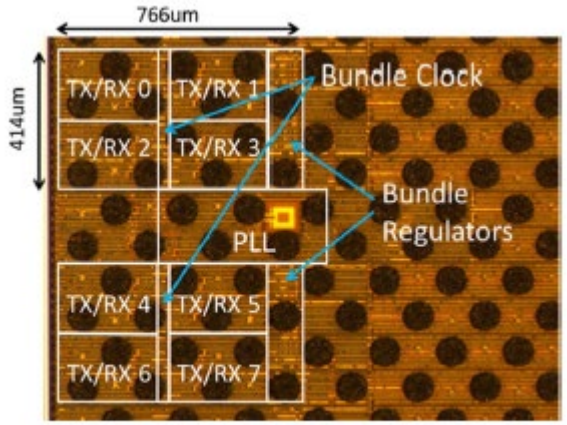
Claim 7	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

970

Claim 7	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

971

Claim 7	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

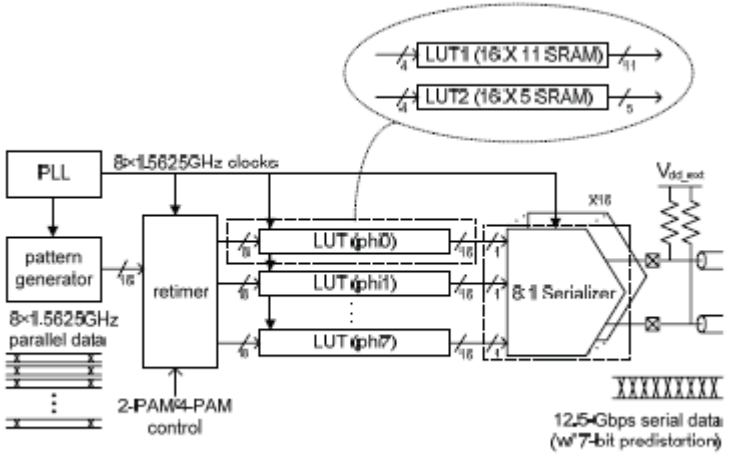
Claim 7	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 7	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

974

Claim 7	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

975

Claim 7	Disclosure
976	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 7	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

977

Claim 7	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

978

Claim 7	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

979

Claim 7	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

980

Claim 7	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 7	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

982

Claim 7	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

983

Claim 7	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

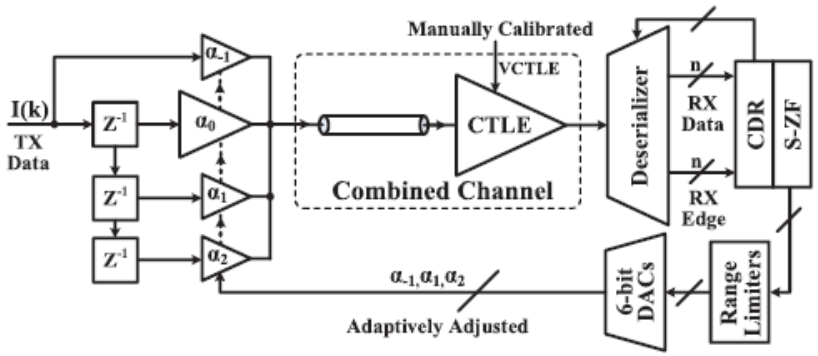
Claim 7	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

985

Claim 7	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

986

Claim 7	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

Claim 7	Disclosure
988	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 7	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

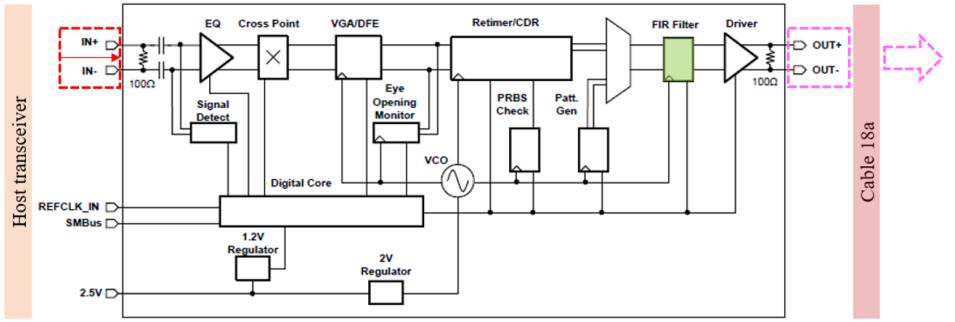
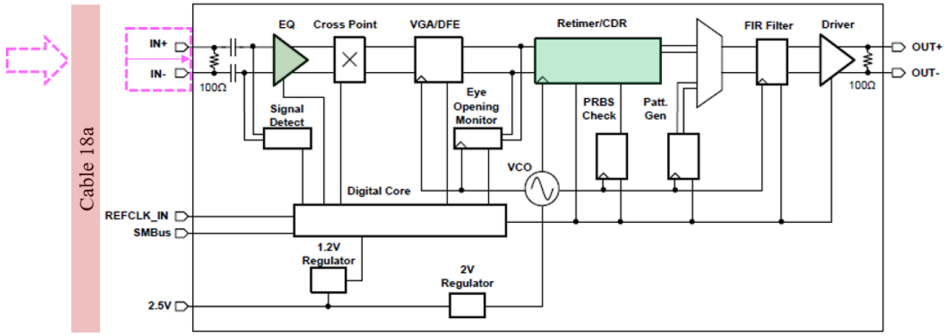
686

Claim 7	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20.

990

Claim 7	Disclosure
	"FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."

991

Claim 7	Disclosure
	 Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram
7[d] said performing clock and data recovery on the second electrical input signal,	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

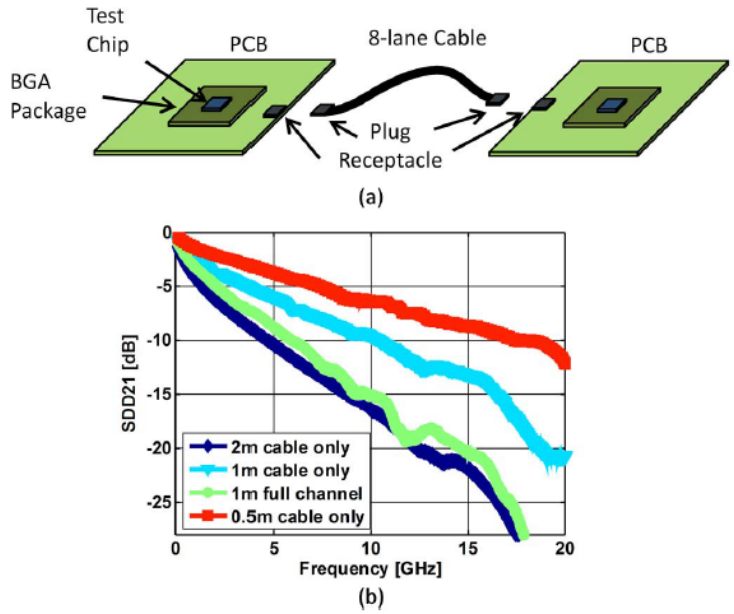
992

Claim 7	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

993

Claim 7	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

994

Claim 7	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

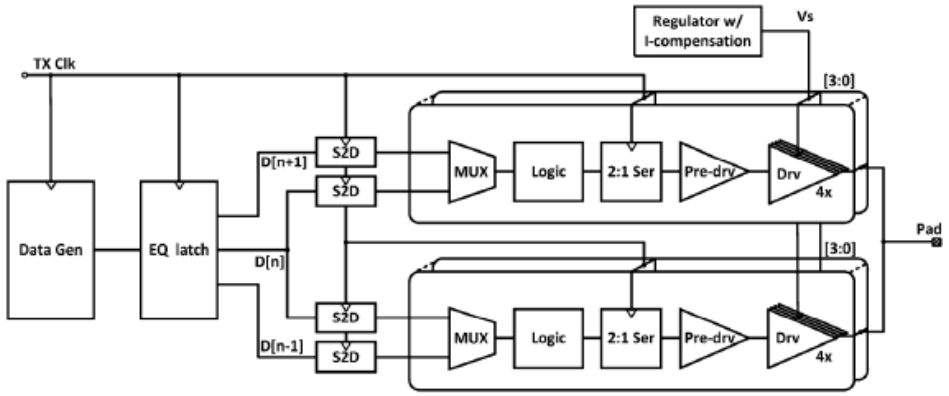
995

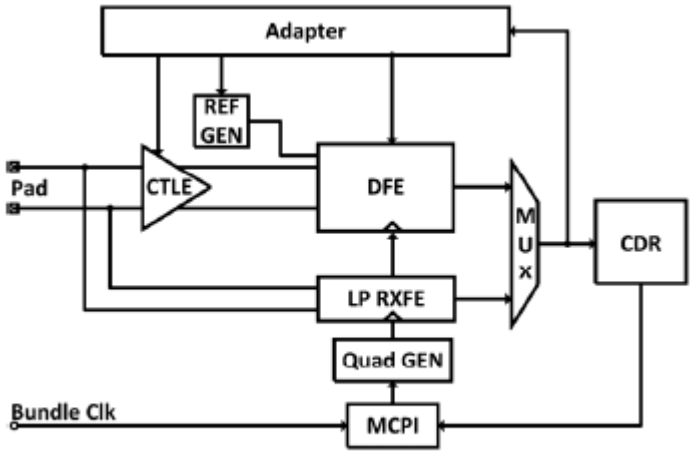
Claim 7	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

966

Claim 7	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

997

Claim 7	Disclosure
866	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

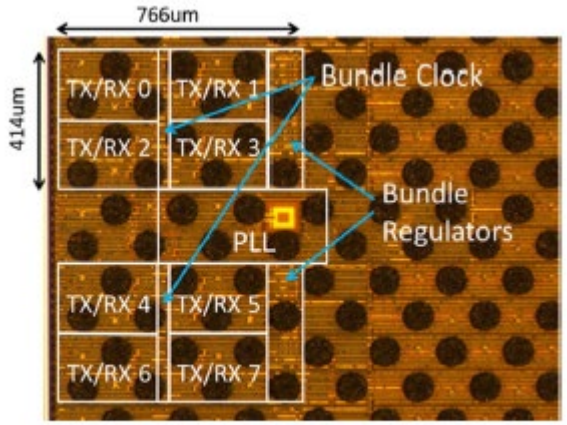
Claim 7	Disclosure
666	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 7	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1000

Claim 7	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1001

Claim 7	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

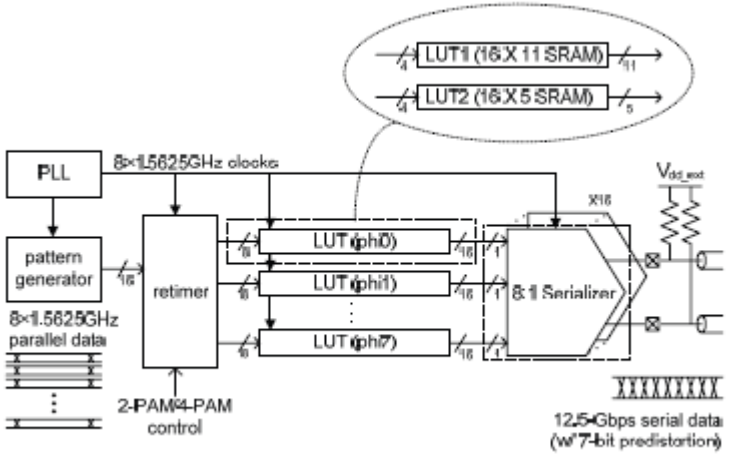
1002

Claim 7	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1003

Claim 7	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1004

Claim 7	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1005

Claim 7	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1006

Claim 7	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1007

Claim 7	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1008

Claim 7	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1009

Claim 7	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

1010

Claim 7	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1011

Claim 7	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1012

Claim 7	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1013

Claim 7	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1014

Claim 7	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1015

Claim 7	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1016

Claim 7	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

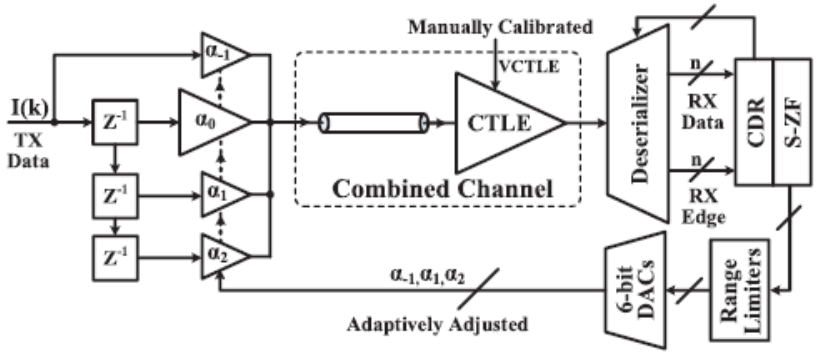


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1017

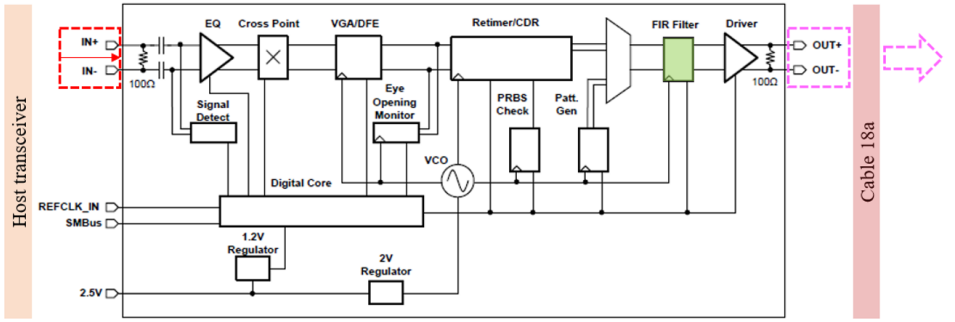
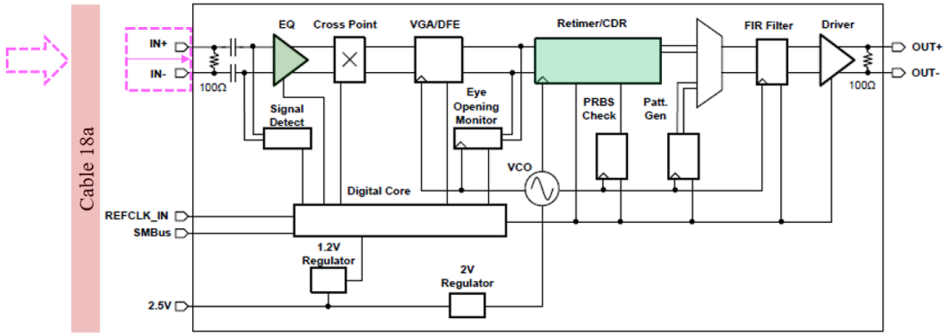
Claim 7	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1018

Claim 7	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20.

1019

Claim 7	Disclosure
1020	"FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."

Claim 7	Disclosure
	 Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram
7[e] each employ fixed, cable-independent, equalization parameters.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

1021

Claim 7	Disclosure
	Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

1022

Claim 7	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1023

Claim 7	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

1024

Claim 7	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1025

Claim 7	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1026

Claim 7	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1027

Claim 7	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1028

Claim 7	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

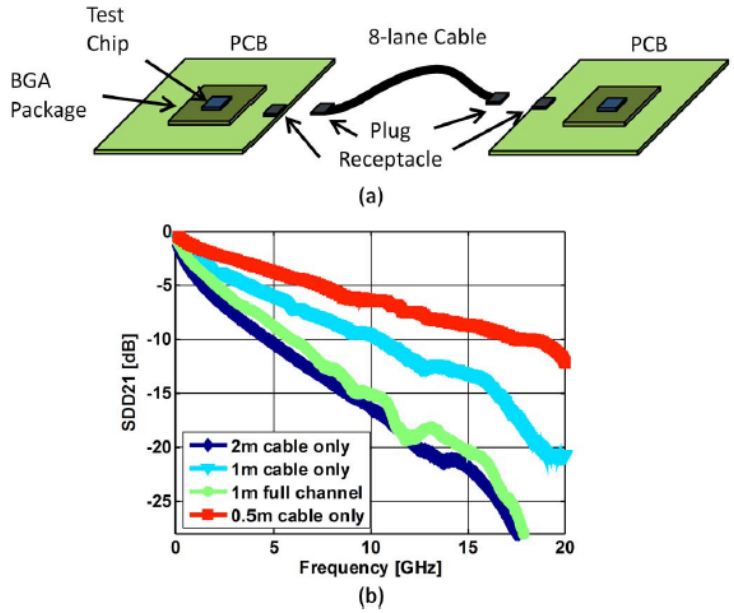
1029

Claim 7	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1030

Claim 7	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1031

Claim 7	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

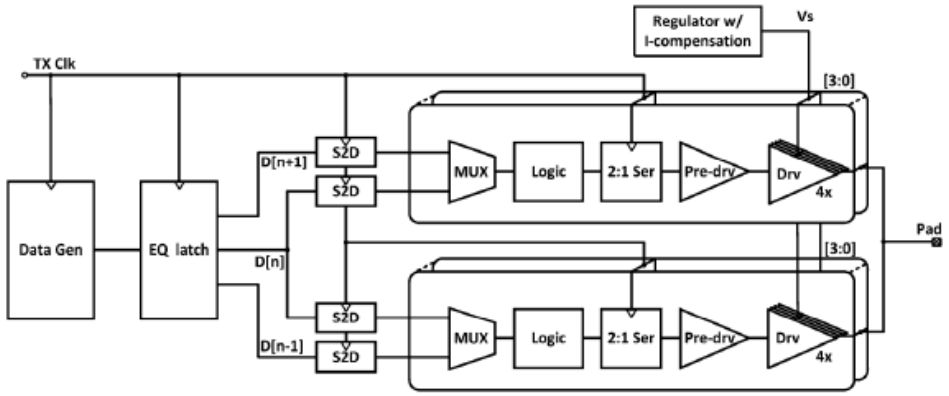
1032

Claim 7	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

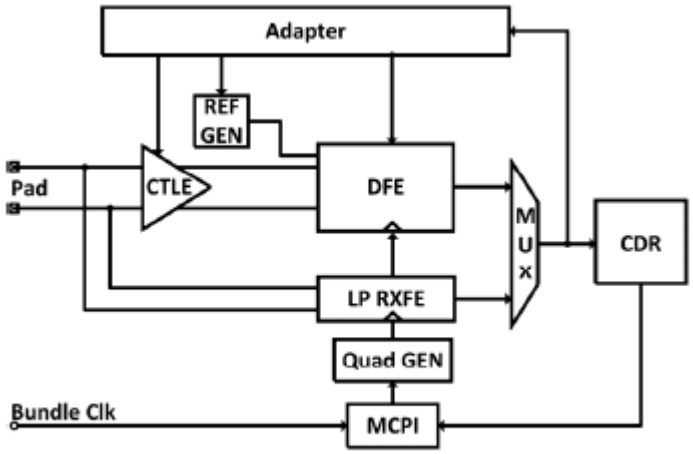
1033

Claim 7	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1034

Claim 7	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1035

Claim 7	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

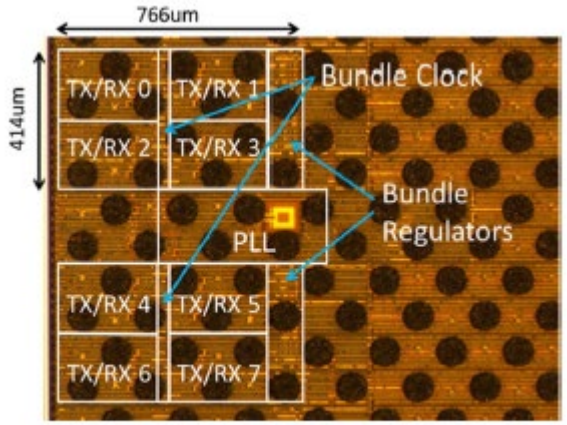
1036

Claim 7	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1037

Claim 7	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1038

Claim 7	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

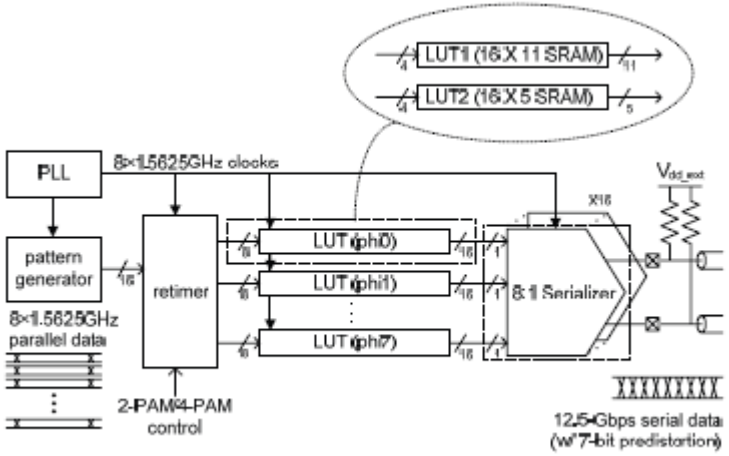
1039

Claim 7	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1040

Claim 7	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1041

Claim 7	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1042

Claim 7	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1043

Claim 7	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1044

Claim 7	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1045

Claim 7	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1046

Claim 7	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 7	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1048

Claim 7	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1049

Claim 7	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1050

Claim 7	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1051

Claim 7	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1052

Claim 7	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1053

Claim 7	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

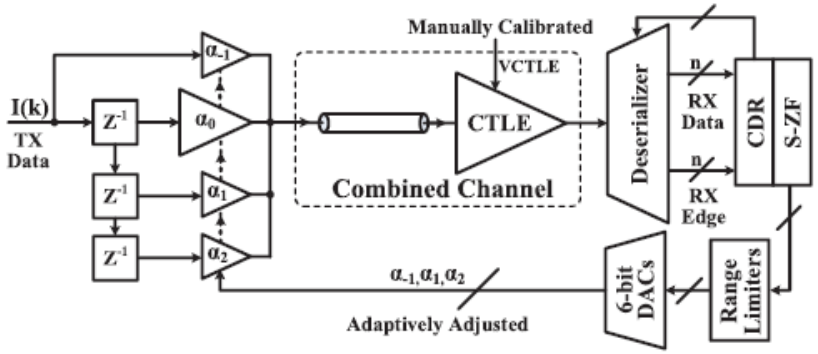


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1054

Claim 7	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

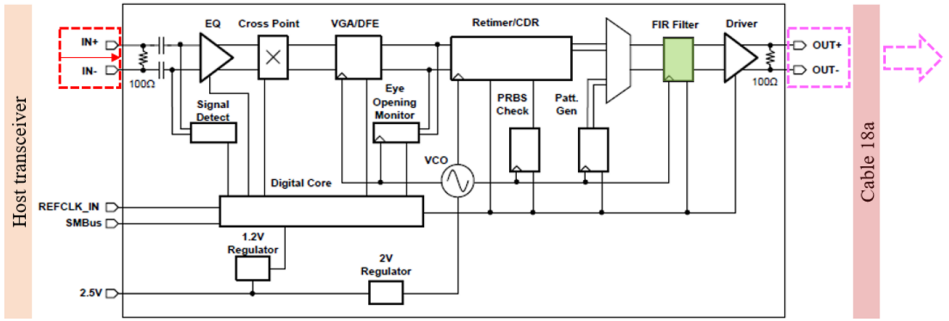
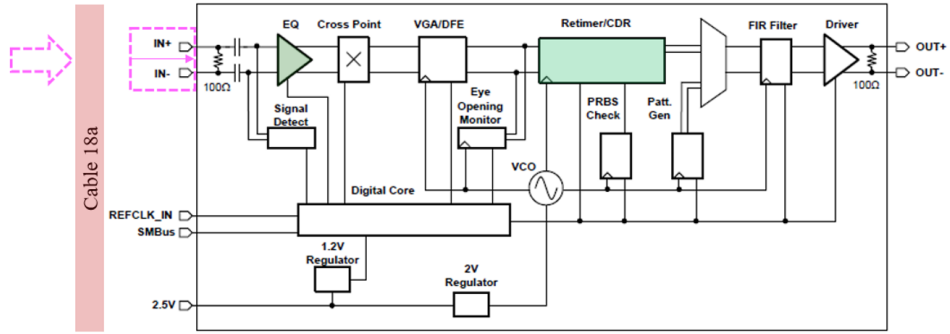
1055

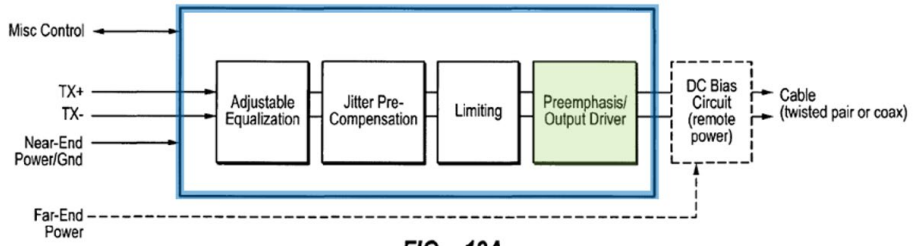
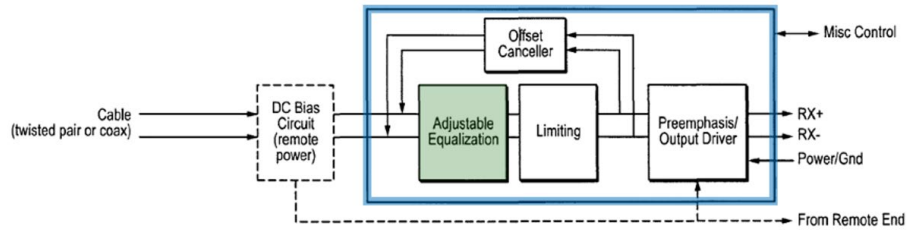
Claim 7	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The PCS 110 is configured to encode Ethernet frame data received from the MAC 112 into encoded data blocks that may be transmitted by the PMD/PMA 106, and to decode data received from the PMD/PMA 106 into decoded Ethernet frame data that may be processed by the MAC 112. The PCS 110 may encode and distribute the data blocks over one or more logical PCS lanes. The PCS 110 may encode data for transmission over the communication lanes 102 for example, to improve communication efficiency. For example, encoding the data may add timing or synchronization symbols, align the data, add state transitions to the encoded data to improve clock recovery, adjust the DC balance of the data signal, or otherwise prepare the encoded data for serial transmission. The PCS 110 may be capable of encoding or decoding the data using a 64b/66b line code in which 64-bit blocks of data are encoded into 66-bit blocks of encoded data, and vice versa." Lusted, paragraph 20.

1056

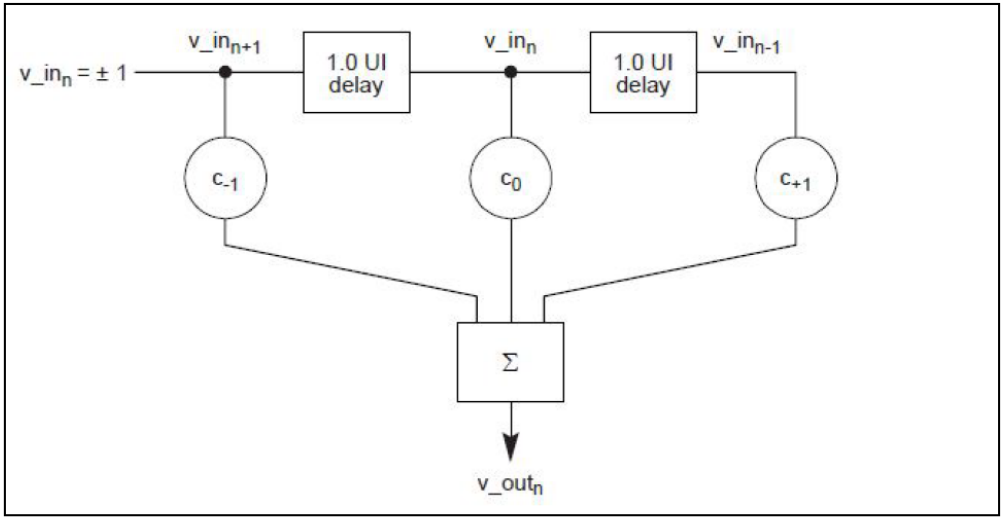
Claim 7	Disclosure
	"FIG. 9 is a block diagram of still another example of a configuration of the second embodiment of the active connector according to the present invention. An active connector 107 illustrated in FIG. 9 differs from the active connector 106 illustrated in FIG. 8 in that the connector base 20 of the active connector 107 further includes a clock recovery circuit (CDR circuit) that recovers the timing of the input signal and reclocks the signal with an optimal timing at each time slot. More specifically, as illustrated in FIG. 9, after an input signal P32 supplied from the coaxial cable transmission line to the external port 30 is supplied to the matching circuit 210, the signal from the matching circuit 210 is supplied to the equalizing circuit 216 and then to a clock recovery circuit 218. Next, the output from the clock recovery circuit 218 is supplied to the buffer circuit 214, and then the output signal P12 is supplied from the buffer circuit 214 through the matching circuit 212 to the internal port 10. The clock recovery circuit 218 is driven by direct-current voltage Vcc supplied from an electrical power source (not shown) to a power supply terminal 12." Ikeda, 10:20-38. *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit

1057

Claim 7	Disclosure
1058	signal.”  Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram *U.S. Patent No. 7,762,727 to Aronson* Aronson discloses cable-side equalization and pre-equalization that use cable-dependent parameters to address cable-side losses. *See, e.g.*, Aronson, 14:33-43, 15:1-8. Aronson’s transmitter IC’s line-side pre-emphasis/output driver overcomes “loss on the cable” using cable-dependent parameters that “match the particular length [and/or] the characteristics of the copper cabling.” *See id*. at 14:33-43, Fig. 12A (below).

Claim 7	Disclosure
	Similarly, Aronson's receiver IC equalizes received line-side signals to "compensat[e] for the cable" losses using cable-dependent parameters based on "cable length and characteristics." *See id*. at 15:1-8, Fig. 12B (below).  FIG. 12A  FIG. 12B

1059

Claim 7	Disclosure
	 **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.

1060

H. DEPENDENT CLAIM 8

Claim 8	Disclosure
8. The communication method of claim 7, wherein said cable-dependent equalization parameters are adaptively updated.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

1061

Claim 8	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1062

Claim 8	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

1063

Claim 8	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1064

Claim 8	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1065

Claim 8	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1066

Claim 8	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1067

Claim 8	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

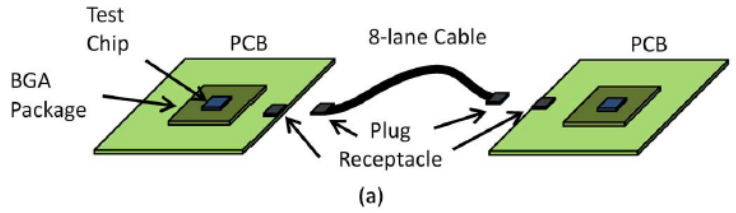
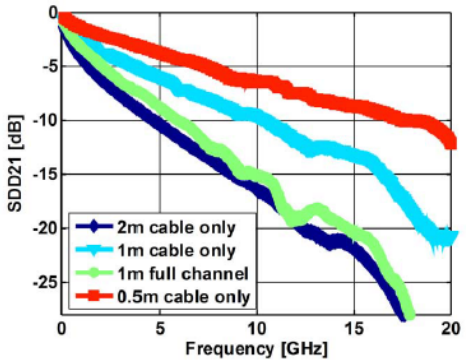
1068

Claim 8	Disclosure
	Abstract—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1069

Claim 8	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1070

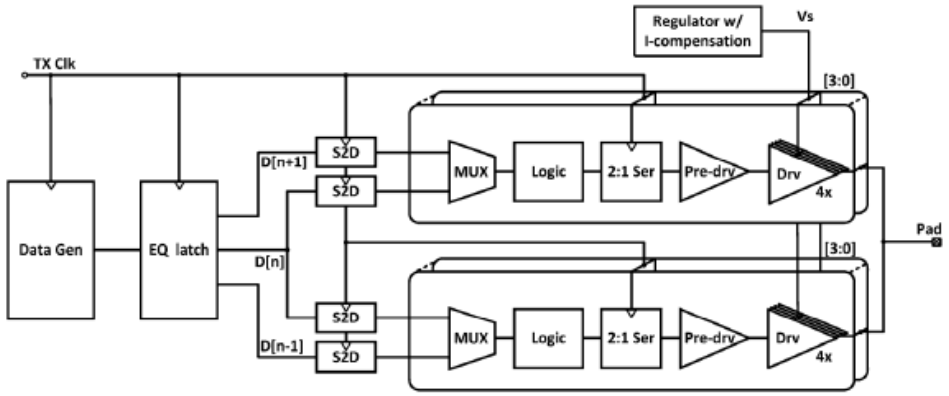
Claim 8	Disclosure
1071	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 8	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

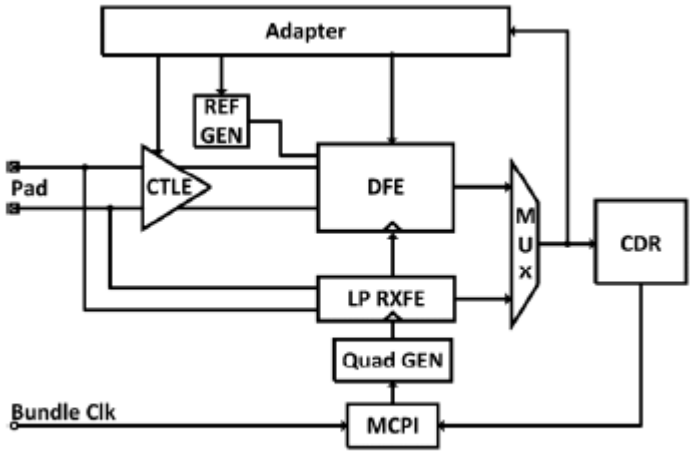
1072

Claim 8	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1073

Claim 8	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1074

Claim 8	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

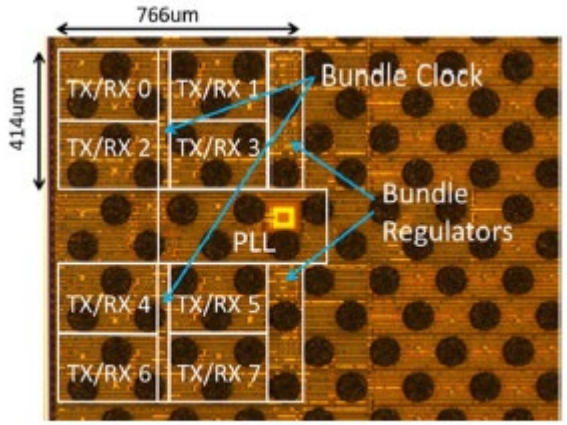
1075

Claim 8	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1076

Claim 8	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1077

Claim 8	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

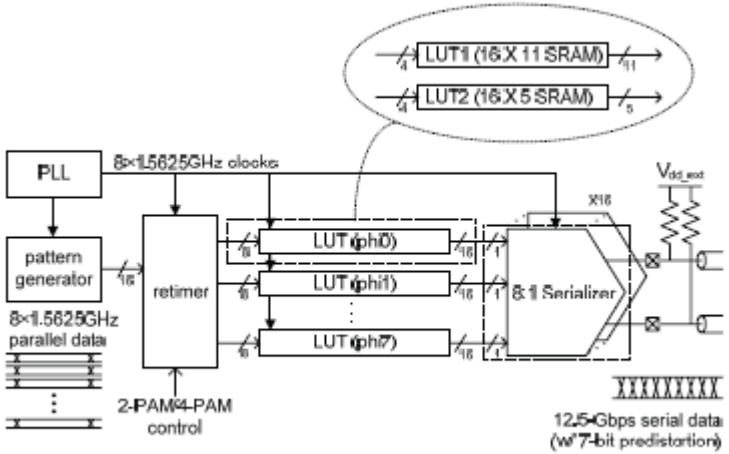
1078

Claim 8	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1079

Claim 8	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1080

Claim 8	Disclosure
1081	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 8	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1082

Claim 8	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1083

Claim 8	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1084

Claim 8	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1085

Claim 8	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

1086

Claim 8	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1087

Claim 8	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1088

Claim 8	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1089

Claim 8	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1090

Claim 8	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1091

Claim 8	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1092

Claim 8	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

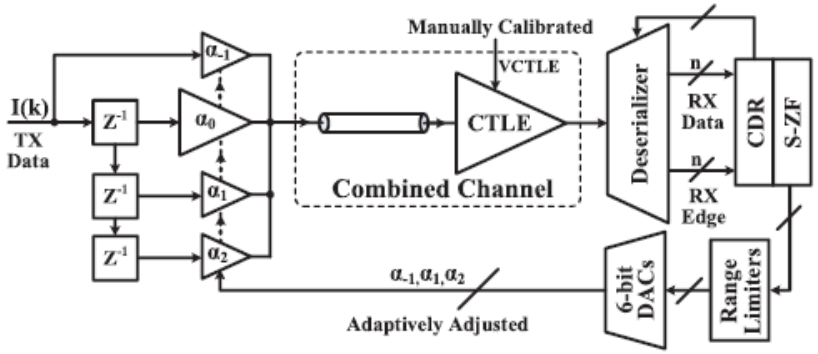


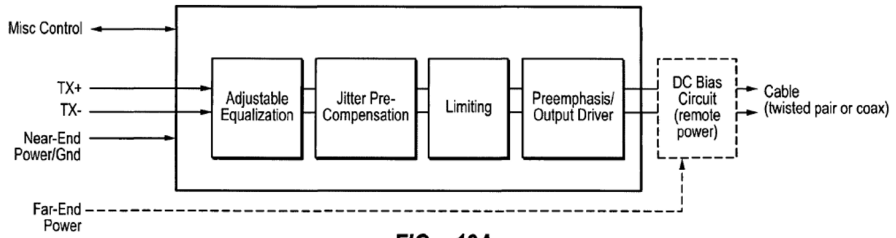
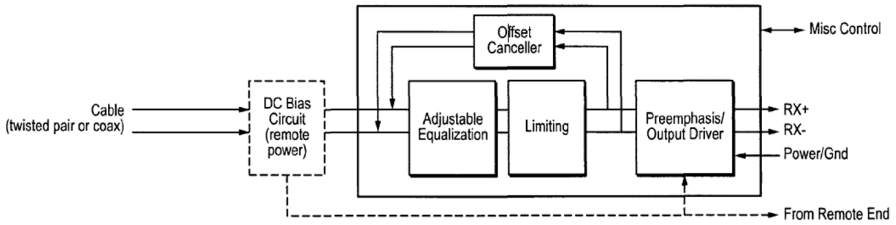
Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

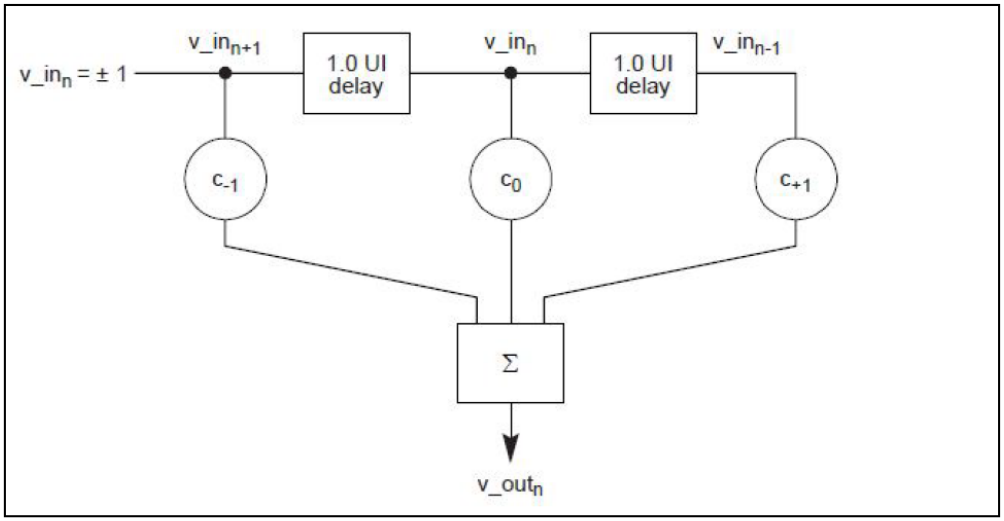
1093

Claim 8	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1094

Claim 8	Disclosure
1095	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. *U.S. Patent No. 7,762,727 to Aronson*

Claim 8	Disclosure
1096	Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id*. at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id*. at 14:16-15:24.  *FIG. 12A*  *FIG. 12B* Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23.

Claim 8	Disclosure
1097	 **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.

I. DEPENDENT CLAIM 9

Claim 9	Disclosure
9[a] The communication method of claim 7, wherein said cable-dependent equalization parameters are fixed during normal usage, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

1098

Claim 9	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1099

Claim 9	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

1100

Claim 9	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1101

Claim 9	Disclosure
1102	**1.7 Module Memory Map** This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module. CMIS, 44. **1.5.5 Signal Integrity Control Field Background** The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated. The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11). Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable Control / Tx Adaptive Input Eq Enable value to use this control Tx Input Eq Adaptation Freeze — 1 Tx Input Eq Adaptation Store — 1 Tx Adaptive Input Eq Recall — 1 Tx Input Eq control — 0 The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module. CMIS, 40.

Claim 9	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1103

Claim 9	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1104

Claim 9	Disclosure

1105

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

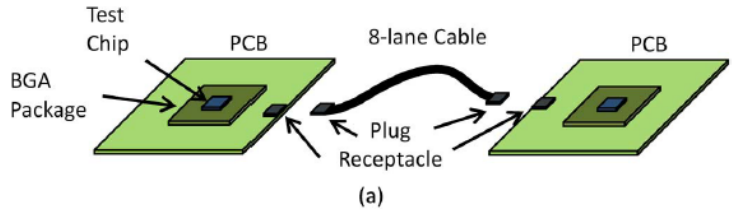
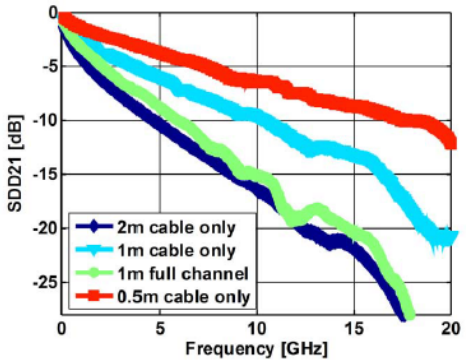
CMIS, 71.

Claim 9	Disclosure
	Abstract—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1106

Claim 9	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1107

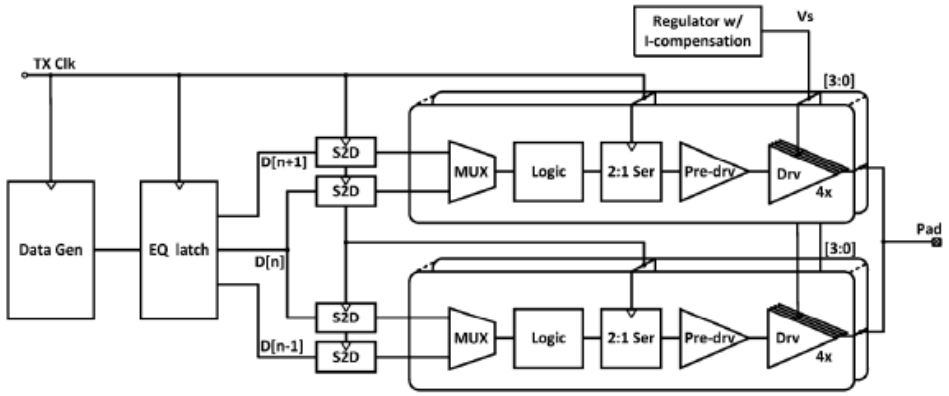
Claim 9	Disclosure
1108	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 9	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

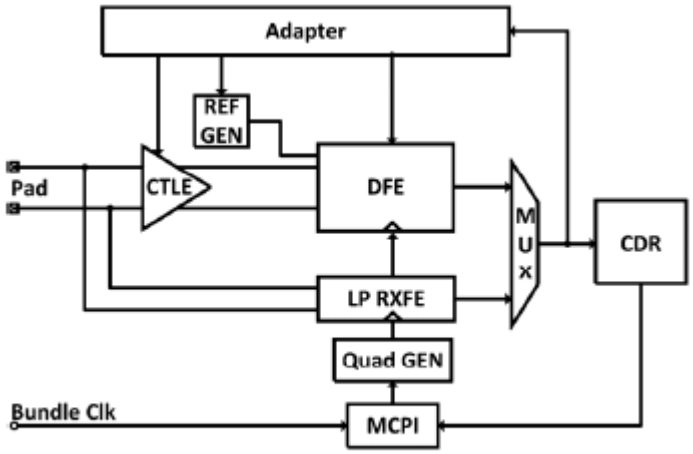
1109

Claim 9	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1110

Claim 9	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1111

Claim 9	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

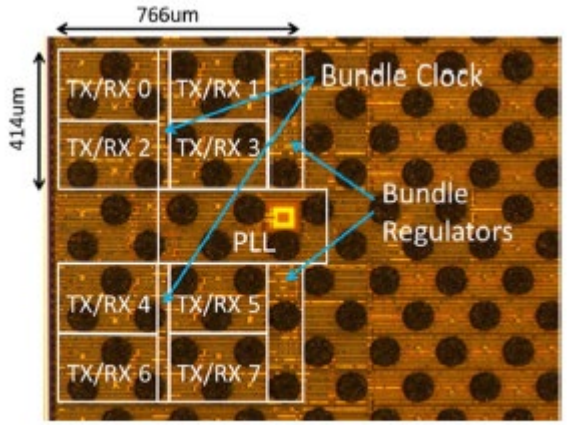
1112

Claim 9	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1113

Claim 9	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1114

Claim 9	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

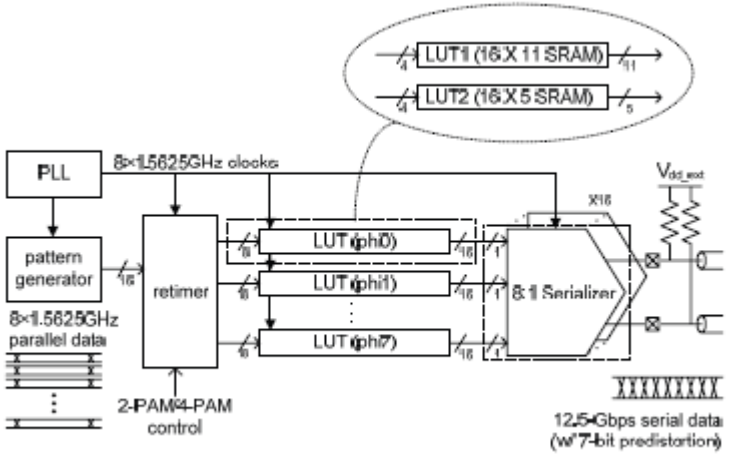
1115

Claim 9	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1116

Claim 9	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1117

Claim 9	Disclosure
1118	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 9	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1119

Claim 9	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1120

Claim 9	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1121

Claim 9	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1122

Claim 9	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

1123

Claim 9	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1124

Claim 9	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1125

Claim 9	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/ LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1126

Claim 9	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1127

Claim 9	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1128

Claim 9	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1129

Claim 9	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

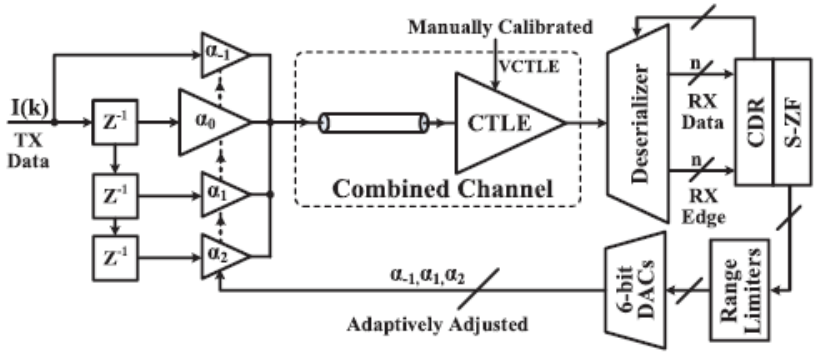
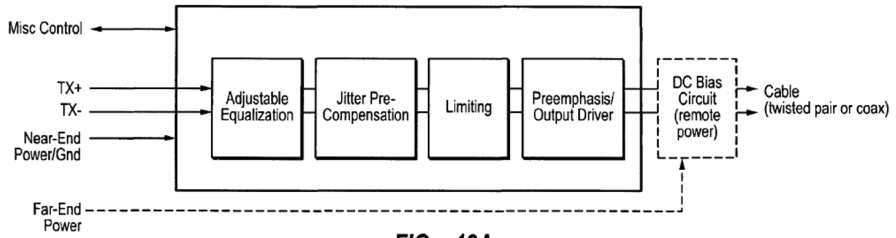
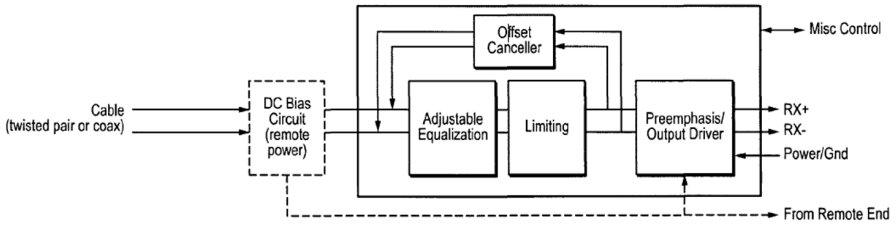


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1130

Claim 9	Disclosure
1131	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

Claim 9	Disclosure
1132	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. *U.S. Patent No. 7,762,727 to Aronson*

Claim 9	Disclosure
	Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id*. at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id*. at 14:16-15:24.  *FIG. 12A*  *FIG. 12B* Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23.

1133

Claim 9	Disclosure
9[b] wherein the method further comprises: determining said cable-dependent equalization parameters during manufacturer-testing of the network cable.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

1134

Claim 9	Disclosure
1135	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

Claim 9	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

1136

Claim 9	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1137

Claim 9	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1138

Claim 9	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1139

Claim 9	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1140

Claim 9	Disclosure

1141

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

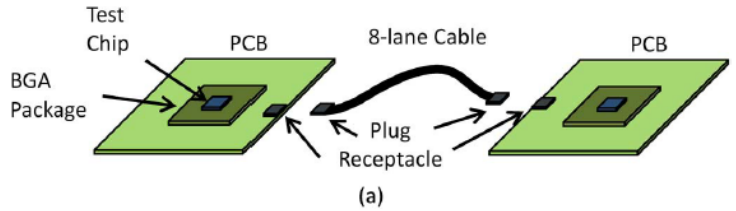
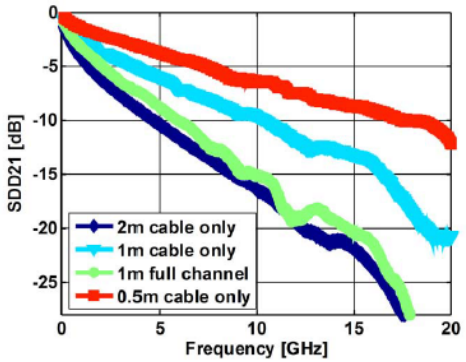
CMIS, 71.

Claim 9	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1142

Claim 9	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1143

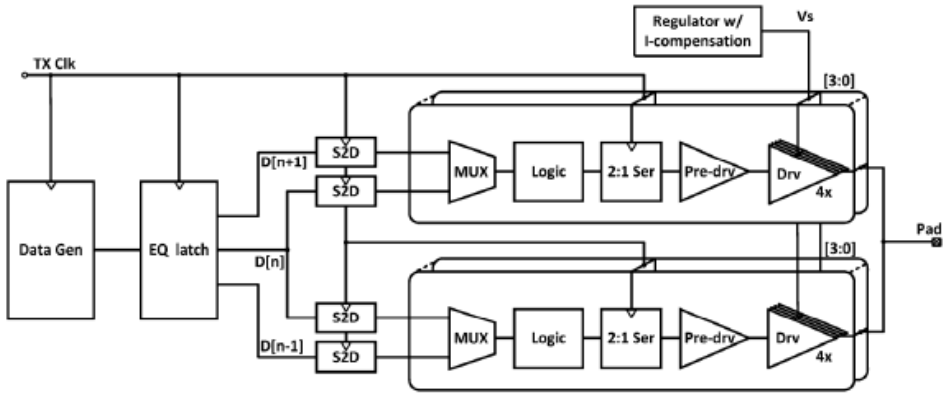
Claim 9	Disclosure
1144	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 9	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

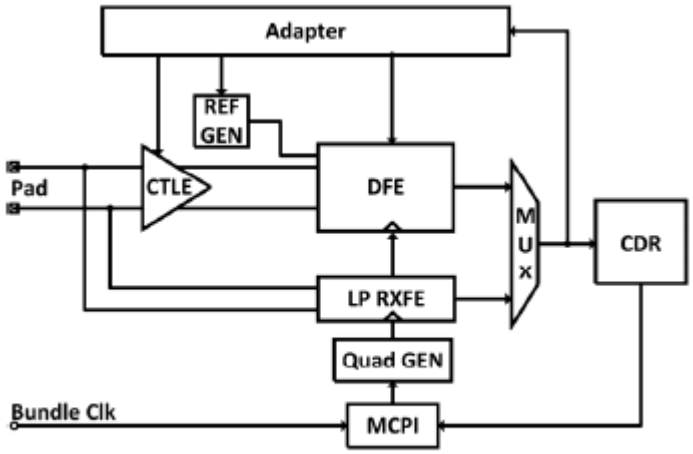
1145

Claim 9	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1146

Claim 9	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1147

Claim 9	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

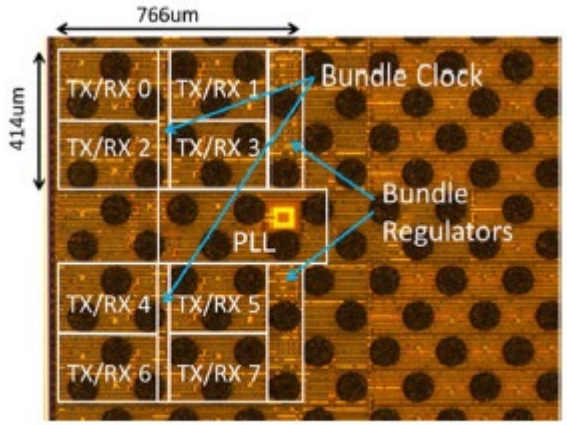
1148

Claim 9	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1149

Claim 9	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1150

Claim 9	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

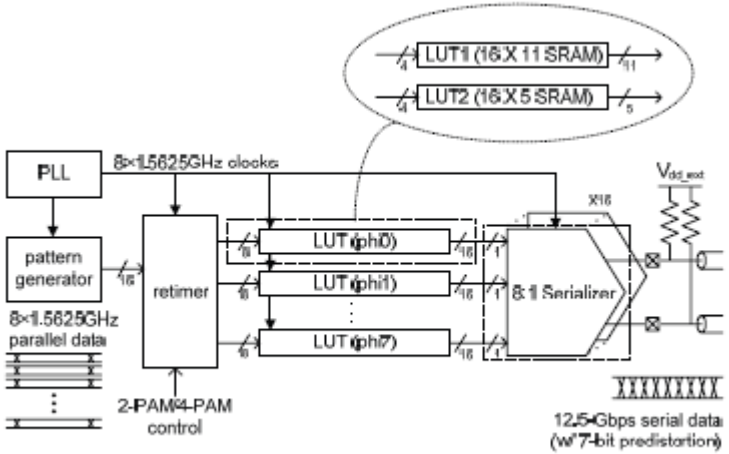
1151

Claim 9	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1152

Claim 9	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1153

Claim 9	Disclosure
1154	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 9	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1155

Claim 9	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1156

Claim 9	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1157

Claim 9	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

Claim 9	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 9	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1160

Claim 9	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1161

Claim 9	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/ LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1162

Claim 9	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1163

Claim 9	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1164

Claim 9	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1165

Claim 9	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

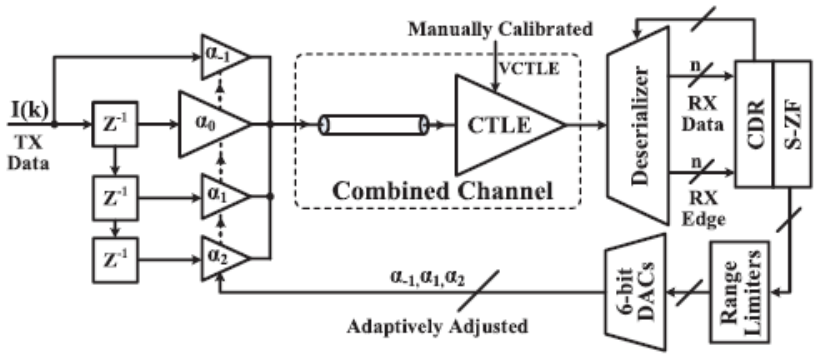


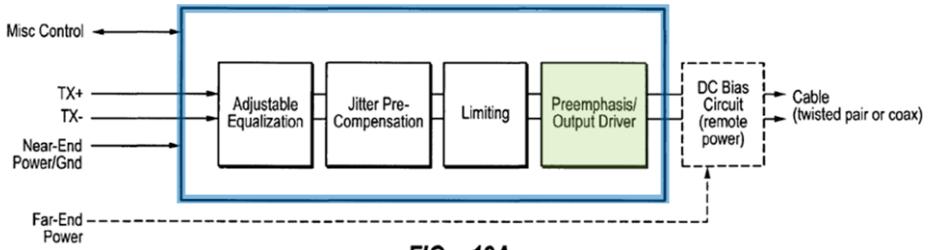
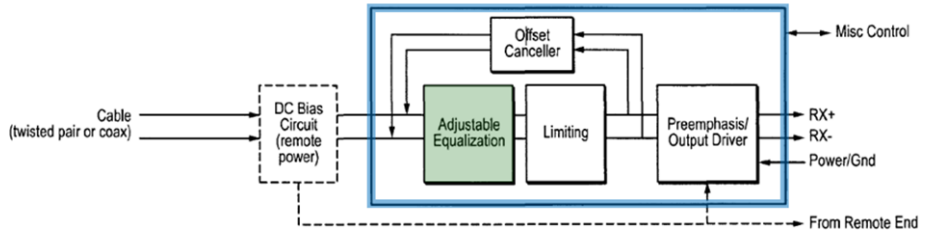
Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1166

Claim 9	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1167

Claim 9	Disclosure
1168	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55.

Claim 9	Disclosure
1169	*U.S. Patent No. 7,762,727 to Aronson* One obvious way to implement the equalization coefficients used by a FIR filter and an equalizer is during manufacturing-testing. *See, e.g.*, Aronson at 14:35-39 (Fig. 12A's pre-equalization parameters for transit over the line-side are "adjusted at factory setup"), 15:1-23 (Fig. 12B's equalization may be "fixed" or "adjustable but factory set" "at the time of the cable manufacture").  *FIG. 12A*  *FIG. 12B*

J. DEPENDENT CLAIM 10

Claim 10	Disclosure
10. The communication method of claim 9, wherein the first inbound data stream has a per-lane symbol rate in excess of 50 GBd.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** ***Index Terms*—Active inductor CTLE, bidirectional link, collaborative CDR, conditional phase detection, decision-feedback equalizer (DFE), phase error decimation, regulated CMOS clocking, source-series terminated (SST) driver.** Musah, 3079.

1170

Claim 10	Disclosure
	II. LINK ARCHITECTURE The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1171

Claim 10	Disclosure
	IV. MEASUREMENTS Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. The measured TX output is shown in Fig. 16 at 32 Gb/s operation over a short channel. Fig. 17 shows the on-die 10^{-10} BER eye at 32 Gb/s over a 0.5 m cable using PRBS10 data on a single lane. Amortizing the PLL and clock distribution power over 4 lanes, the link power consumption is 205 mW at 32 Gb/s. Transceiver power efficiency given for multiple data rates from 4 to 32 Gb/s is shown in Fig. 18. Excluding the PLL power, the transceiver power efficiency ranges from 3.4 to 5.7 pJ/b, while equalizing 24 dB and 16 dB of Nyquist rate channel loss at 16 Gb/s and 32 Gb/s, respectively. Due to circuit BW limitations, the equalized channel loss tolerance is 24 dB up to 27 Gb/s and reduces to 16 dB at the highest data rate. When the HP RX is disabled and the LP RX is enabled, the measured power efficiencies are 1.0 pJ/b at 4 Gb/s, 1.5 pJ/b at 8 Gb/s and 2.2 pJ/b at 12 Gb/s with 4.5 dB, 8 dB and 10 dB of channel loss, respectively. A breakdown of the link power is shown in Table I, with PLL power amortized over 4 lanes. Musah, 3086.

1172

Claim 10	Disclosure

TABLE II
COMPARISON TABLE

	[1]	[2]	[3]	[4]	This Work	
CMOS technology	32nm	32nm	28nm	40nm	22nm	
Supply voltage [V]	1.2/1.05	1.00	1.5/1.05	0.90	0.90	1.07
Data rate [Gb/s]	28	32	28	28	16	32
Channel loss at Nyquist [dB]	35	30	34	20	24	16
TX FFE	4-tap	4-tap	3-tap	3-tap	3-tap	3-tap
RX DFE	15-tap	15-tap	14-tap	N/A	6-tap	6-tap
RX peaking amplifier	2-stage	2-stage	1-stage	1-stage	1-stage	1-stage
TX output swing [$V_{pp\text{-}diff}$]	0.950			0.6	0.515	0.630
TX output jitter [ps_{rms}]	0.45		0.25	0.2		<0.54
RX input sensitivity [mV]	15		15	27		35
Area density [mm^2/Gb/s]	0.0289	0.0219	0.0298	0.0506	0.0049	0.0025
Transceiver power efficiency per lane [pJ/b]	24.8	21	20	6.96	4.6	6.4

Musah, 3087.

V. CONCLUSION

We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane.

Musah, 3088.

1173

Claim 10	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1174

Claim 10	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1175

Claim 10	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

1176

Claim 10	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1177

Claim 10	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1178

Claim 10	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

Claim 10	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963. *A. Transmitter Chip Measurement* The TX output is measured after a channel consisting of a doubled bonding wire, a 4-cm PCB trace, and a 0.5-m

1180

Claim 10	Disclosure
	connection cable. Fig. 23(a) shows the over-equalized eye-diagram at 5 Gb/s, where the four sub-levels are contributed by the four FFE taps. Fig. 23(b) and (c) gives the output eye-diagrams at 40 Gb/s before and after applying the four-tap FFE. Obviously, the FFE can significantly improve the eye opening. The eye height and eye width are optimized from 140 mV and 0.45 UI to 180 mV and 0.68 UI, respectively. Meanwhile, the thickness of the eyelid is dramatically reduced from around 330 to 140 mV. Fig. 23(d) shows the properly compensated eye-diagram at the maximum operation speed of 50 Gb/s. Its eye height and eye width are 50 mV and 0.38 UI. Clearly, a wide operation range from 5 to 50 Gb/s is achieved, which is mainly attributed to the multi-MUX-based FFE implementation. Fig. 24 further shows the TX output with four separate eyes. It can be seen that the horizontal eye widths for both fixed clock and PRBS patterns are almost identical, thus proving that the four sampling phases are properly aligned. Zheng, 2972-2973. *B. Receiver Chip Measurement* The RX standalone measurement results are presented in this part. Fig. 25(a) shows the eye-diagram of the 40-Gb/s input data generated by an Anritsu MP1812A through combining four 10-Gb/s PRBS7 sequences, where the single-end eye height and eye width are around 360 mV and 0.71 UI. Fig. 25(b) shows the eye-diagram of the 10-Gb/s recovered

1181

Claim 10	Disclosure
	data with a total jitter of 12.73 ps. The eye-diagrams of the recovered clocks (divided by 2) for the edge sampling and data sampling are shown in Fig. 25(c) and (d), which reveal that the introduced LPFs can optimize the total jitter from 11.48 to 7.66 ps. To demonstrate the effect of the LPFs with adaptively adjusting bandwidth, the jitter transfer (JTRAN) and JTOL curves are measured using a Tektronix BSA286C with a CDR block. The input peak-to-peak swing is tuned to 800 mV and the control voltage of the CTLE is manually set to 710 mV. The JTRAN curves in Fig. 26 illustrate that the bandwidth of the data-sampling path depending on the LPFs is 4 MHz, which is much smaller than 18 MHz for the edge-sampling path determined by the loop parameters. The measured JTOL in Fig. 26 indicates that the embedded LPFs result in a significant dip attenuation around the corner frequency and improve the JTOL amplitudes apparently at high jitter frequencies. Meanwhile, the adaptively adjusting bandwidth of the LPFs makes them exhibit little effect on the phase-tracking slew rate at low jitter frequencies. Additionally, the corner frequency of the JTOL is about 20 MHz, which is much larger than the JTRAN bandwidth of 4 MHz. Zheng, 2973-2974.

1182

Claim 10	Disclosure

1183

Transmitter Chip				
	[4]	[16]	[20]	This work
Technology (nm)	65	14	65	65
Supply (V)	1.2	N/A	1.2	1.2
Data Rate (Gb/s)	50-64	16-40	38.8-42	5-50
Area (mm^2)	1.2 × 1.0	0.215 × 0.13	0.9 × 0.7	1.2 × 0.5
TX Equalization	4-tap FFE	4-tap FFE	5-tap FFE	4-tap FFE
1UI-Delay Gen.	LC-Delay	Multi-MUX	CML-Delay	Multi-MUX
MUX Type	4:1	4:1	N/A	4:1
Data Jitter RJ (ps_{rms})	N/A	0.33@28Gb/s 0.51@40Gb/s	N/A	0.23@40Gb/s 0.18@50Gb/s
Data Jitter (ps) TJ (BER=1e-12)	N/A	10.72@28Gb/s 12.89@40Gb/s	15.1@40Gb/s	9.90@40Gb/s 10.58@50Gb/s
Power (mW)	199*	518	135**	145*
Energy Efficiency (pJ/bit)	3.1	12.9	3.4	3.6

*Exclude PLL power consumption, ** Standalone FFE without data serializer.

Zheng, 2974.

"FIG. 5 illustrates an EYE waveform of a signal obtained after passing a 56 Gbps PAM-4 signal through a raised cosine filter with a bandwidth of 22 GHz and a roll-off of 0.5. In FIG. 5, when the signal has a value higher than a determination threshold 501, the signal is determined to be +3, when the signal has a value between the determination threshold 501 and a determination threshold 502, the signal is determined to be +1, when the signal has a value between the determination threshold 502 and a determination threshold 503, the signal is determined to be −1, and when the signal has a value lower than the determination threshold 503, the signal is determined to be −3. Ranges above the determination thresholds 501, 502, and 503 through which the signal does not pass is referred to as the EYE widths. As seen in FIG. 5, the EYE width is defined by the upper EYE width 505 and the lower the EYE width 506, and the EYE widths 505 and 506 are defined by the data transition time, −3→+3 and +3→−3. The EYE width is narrowed by approximately 15 ps even under the condition that the bandwidth is limited by a frequency higher than the

Claim 10	Disclosure
	Nyquist frequency and there is almost no influence of ISI. In this manner, the EYE width of the PAM multilevel signal is narrowed relative to that of the NRZ signal due to data pattern dependence." Norimatsu, 4: 27-50.

K. INDEPENDENT CLAIM 11

Claim 11	Disclosure
11[pre] A cable manufacturing method that comprises:	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: **1.1 Introduction** This specification defines the Common Management Interface Specification (CMIS) for QSFP-DD, OSFP and COBO in order to enable flexible use of modules and cable assemblies (hereafter referred to as modules unless cable assemblies are specifically mentioned) by the user. It shares some similarities with the management interfaces used in other form factors such as QSFP, SFP, and CFP. This specification supports a module with 8 host electrical lanes. Support for modules with more host lanes is possible (in increments of up to 8 lanes) by a bank selection feature (which allows several instances of an adressable page to exist in the paged register map). Due to this new mechanism this specification is not backwards compatible with SFF-8636. Some host ports may support both 4-channel and 8-channel devices; when a 4-channel device is inserted into such a port the appropriate memory map (e.g. SFF-8636) must be used. This case is outside the scope of this document. The host-to-module interactions needed to initialize active modules are presented in two state machines. A module-level state machine defines the initialization of the management interface as well as control of the module-wide power mode. A data path-specific state machine defines the host and module interactions and behaviors needed for power up and initialization of the data path(s) in the module. The CMIS is intended to cover a wide range of possible module functionalities and applications, ranging from cable assemblies to coherent DWDM modules. The CMIS is designed with the intent not to burden the simple cases but to provide a common set of management interactions and a common register map. In some modules, multiplexing of the high speed signals between the host electrical interface and the module media interface may occur in the module. In this specification, all references to lane numbers are based on the electrical connector interface lanes, unless otherwise indicated. In cases where a status or control aspect is applicable only to lanes after multiplexing or demultiplexing has occurred, the status or control is intended to apply to all lanes in the data path, unless otherwise indicated. Symmetry is expected between the Tx and the Rx hardware structure; for example Tx lanes that are multiplexed in the Tx are demultiplexed in the Rx. Each Application selected by the host is applied to the same Tx and Rx host electrical lanes. CMIS, 5-6.

1185

Claim 11	Disclosure
	Abstract: This document defines the Common Management Interface Specification (CMIS)for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. Table 1 and Table 2 provide a summary of the high-level behaviors and properties of each module state for passive copper cable assembly and active module implementations, respectively. Implementers of passive copper cable assemblies should note that this specification was designed to allow for flat-memory EEPROM implementations for that cable type. Active modules and active cable assemblies shall use the paged memory model. For all module types, refer to sections 1.4.3-1.4.9 for detailed requirements for each state. CMIS, 18.

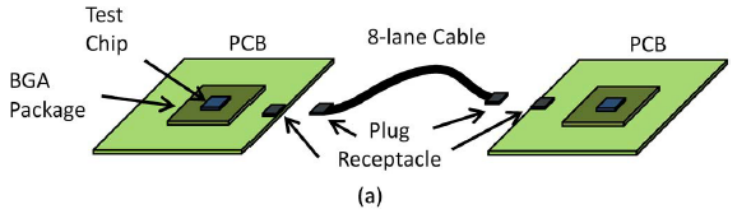
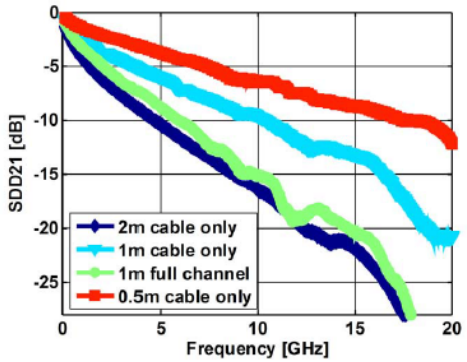
1186

Claim 11	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1187

Claim 11	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1188

Claim 11	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

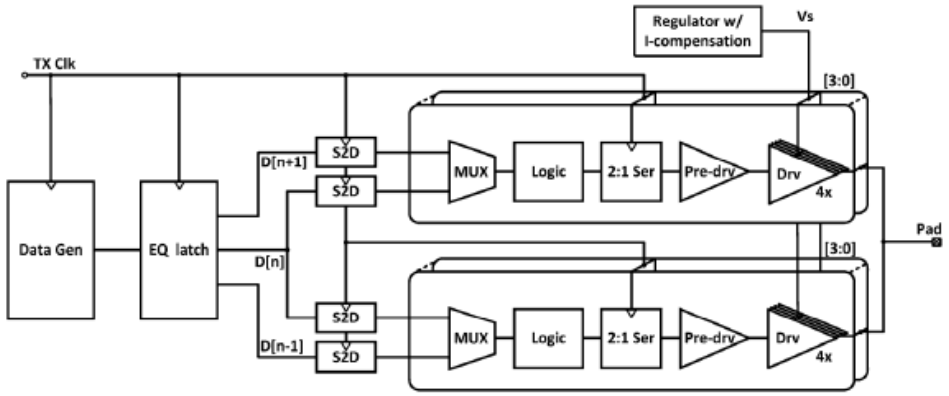
1189

Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

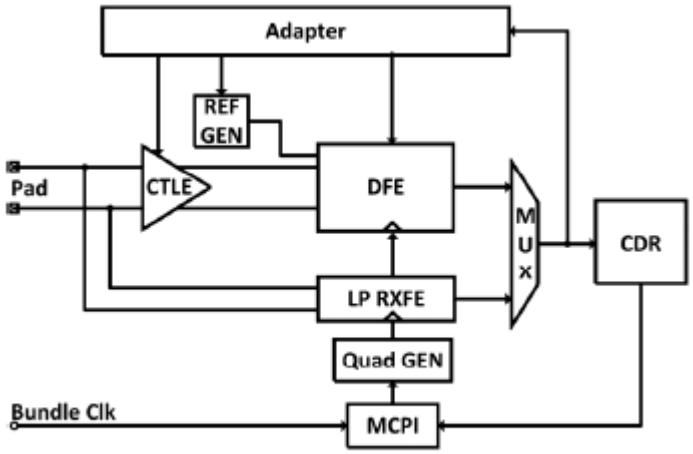
1190

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1191

Claim 11	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1192

Claim 11	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

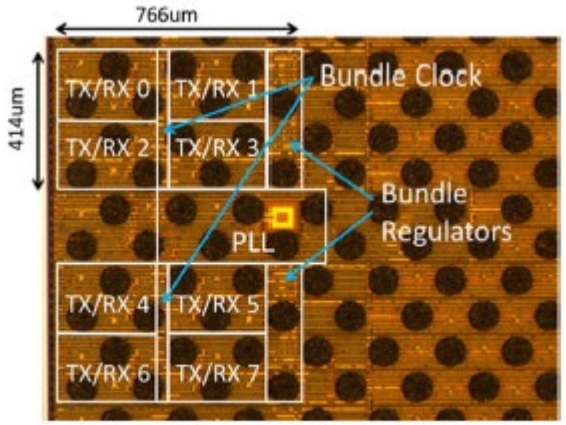
1193

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1194

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

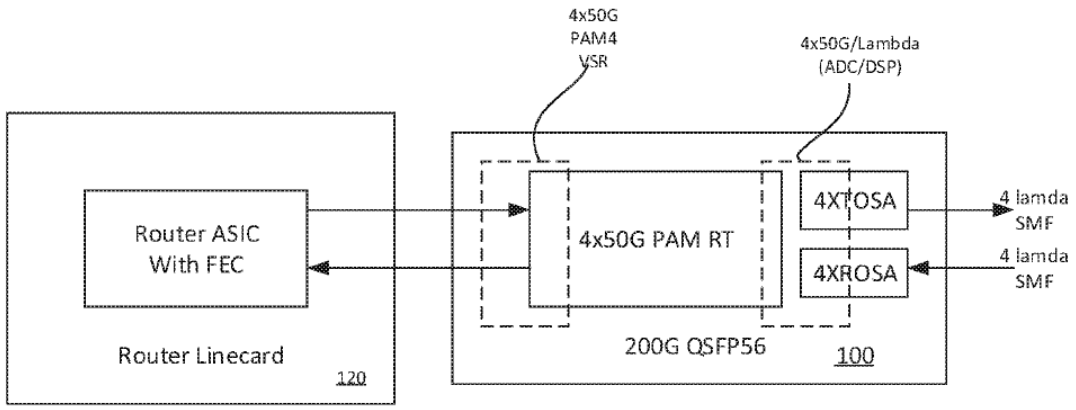
1195

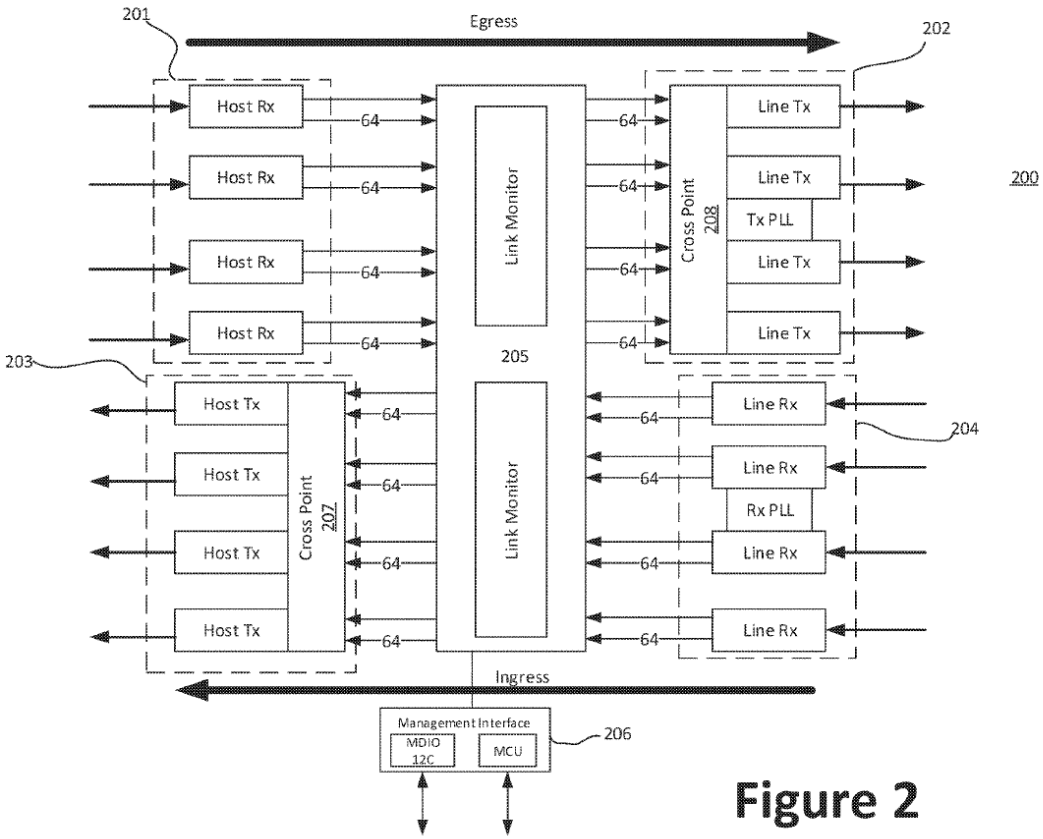
Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

1196

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free $(BER < 10^{-12})$ transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088.

1197

Claim 11	Disclosure
1198	 **Figure 1** Riani, Figure 1.

Claim 11	Disclosure
1199	 Figure 2 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1200

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

1201

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1202

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1203

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1204

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1205

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1206

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1207

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

1208

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1209

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1210

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1211

Claim 11	Disclosure
	THE exponential growth of cloud computing, social networking, and multimedia sharing has led to an explosive bandwidth demand on data communication in both telecommunication equipment and inter/intra data center [1], [2]. To accommodate to this requirement, the data rate of the wireline serializer/deserializer (SerDes) transceiver has been continuously increased [3]–[5]. Currently, 25–28 Gb/s serial links approved by InfiniBand EDR, 32GFC, and CEI-28G have stepped into the period of industrial deployment [1], [3], [6]. The 38–64 Gb/s transceivers, which will play a key role in the next-generation data rate supported by Ethernet 400GbE, InfiniBand HDR, and CEI-56G, have attracted increasing attentions in the industry and the academia [2], [4], [5], [7]–[11]. The main challenges in designing such high-speed transceivers originate from the ever decreased unit interval (UI) period, which not only poses high bandwidth requests on the blocks located at the critical path, but also makes the link timing budget extremely tight. Moreover, advanced processes cannot completely solve these problems, since the parasitic capacitances/resistances at the high-speed outputs usually do not scale well with the technology due to the bonding and/or electro-static discharge (ESD) protection requirements. Zheng, 2963.

1212

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1213

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

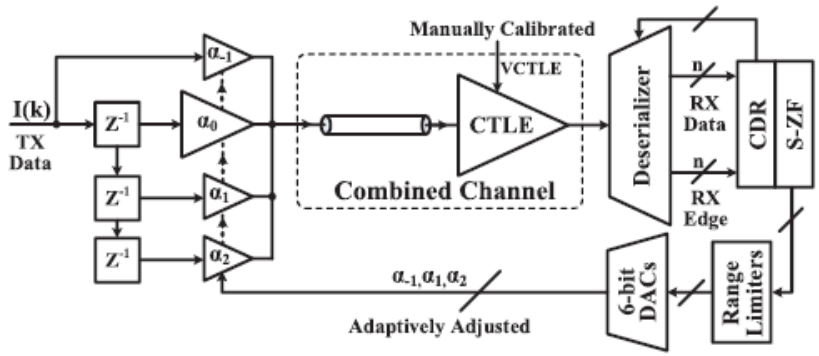
1214

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1215

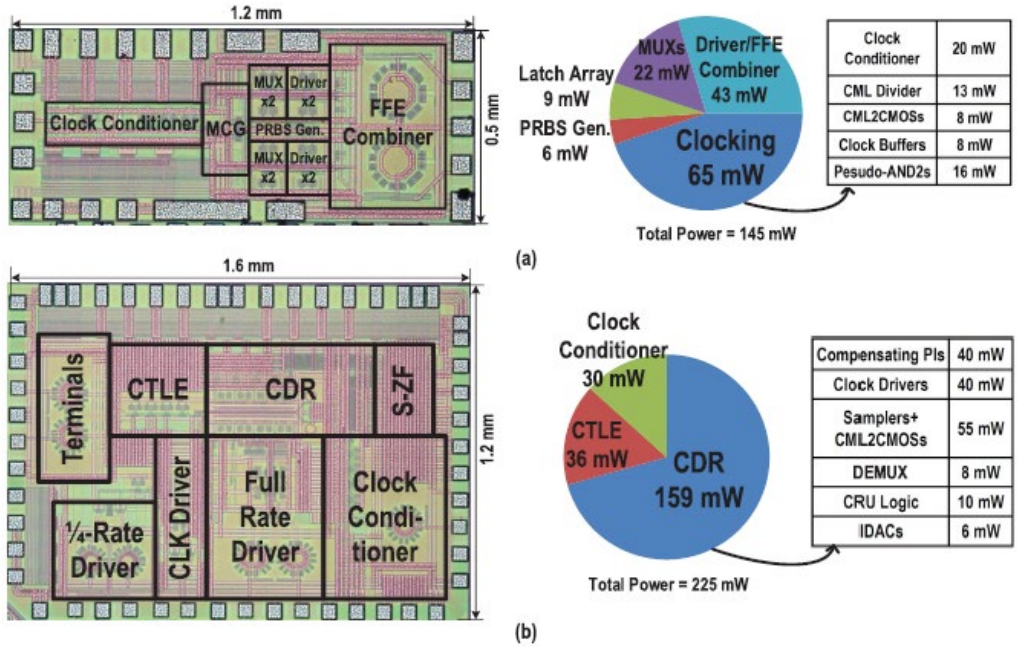
Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1216

Claim 11	Disclosure
1217	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1218

Claim 11	Disclosure
1219	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

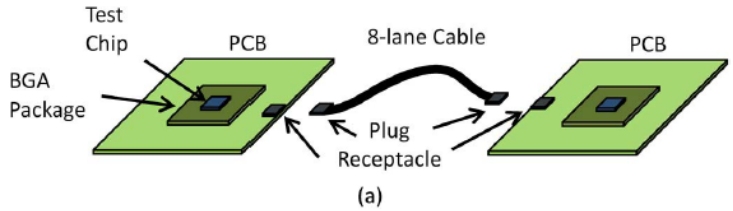
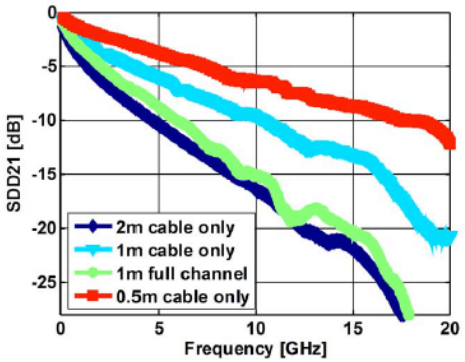
Claim 11	Disclosure
1220	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 1, the data transmission system 1 is in accordance with DP1.4, for example, and includes an upstream device 2, a downstream device 4, and an active cable 3 connecting the upstream device 2 and downstream device 4." Kobayashi, 2:41-45. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream

Claim 11	Disclosure
	transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK" Kobayashi, 4:4-22.
11[a] connecting a first end and a second end of a set of conductor pairs to a first transceiver and a second transceiver, respectively, to transport a first electrical transit signal from the first transceiver to the second transceiver and a second electrical transit signal from the second transceiver to the first transceiver;	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.**

1221

Claim 11	Disclosure
	Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1222

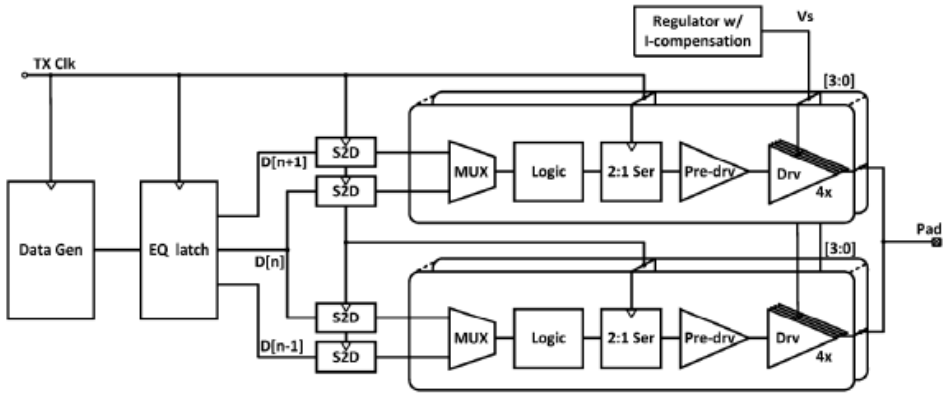
Claim 11	Disclosure
1223	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

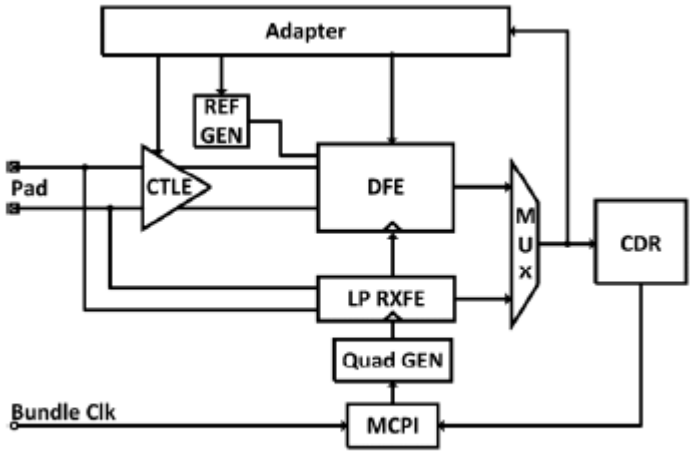
1224

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1225

Claim 11	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1226

Claim 11	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

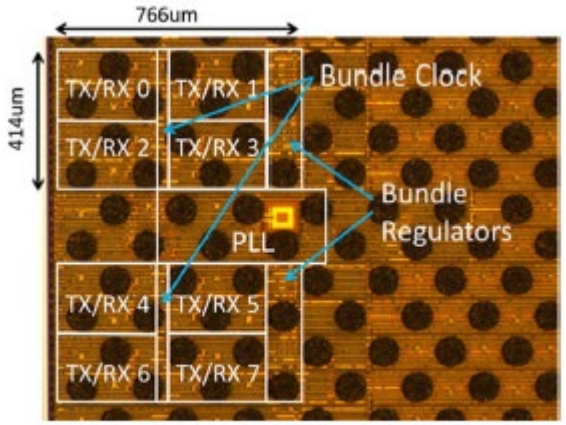
1227

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1228

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1229

Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

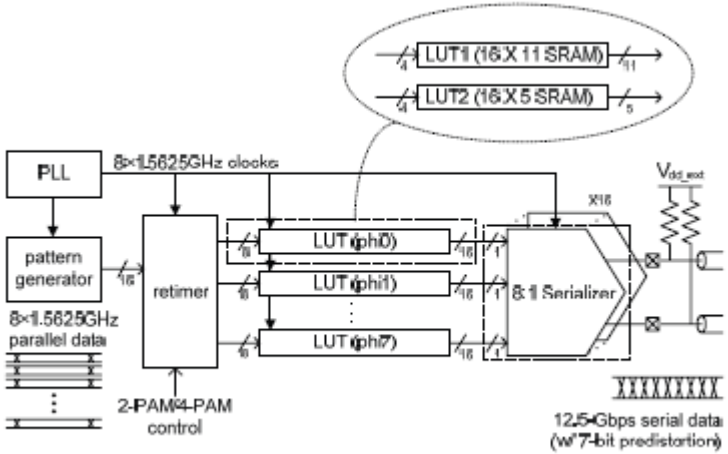
1230

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1231

Claim 11	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1232

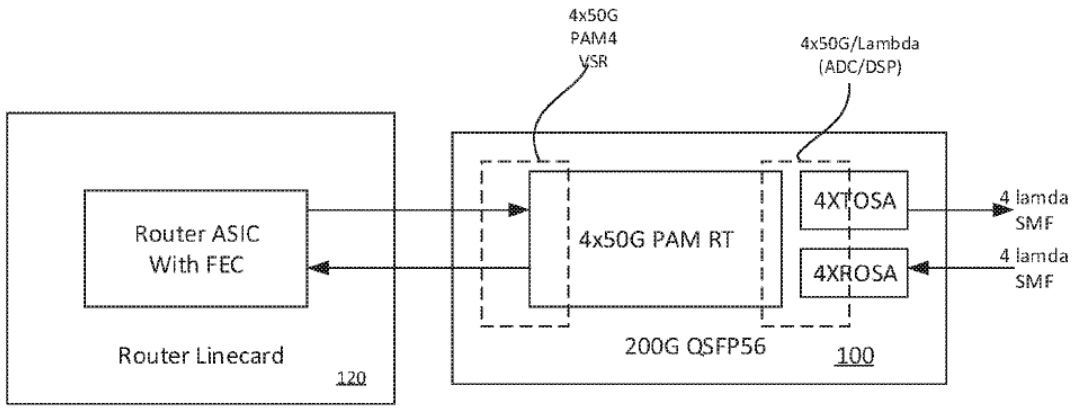
Claim 11	Disclosure
1233	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

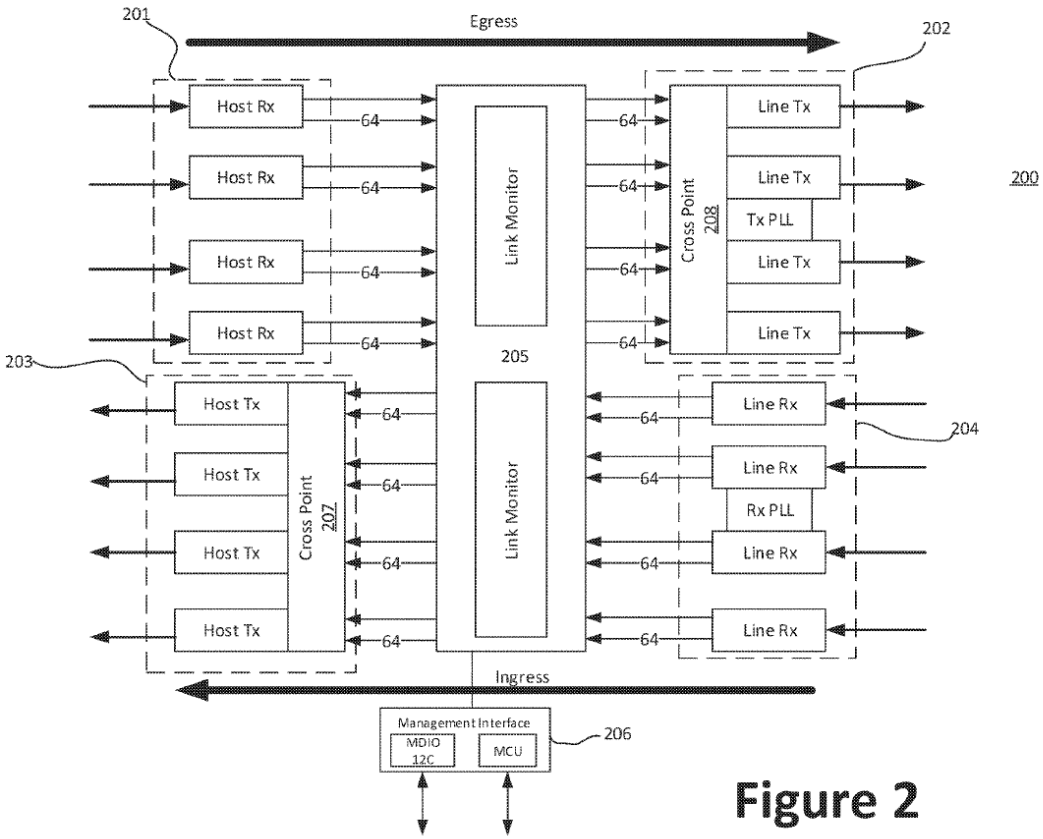
Claim 11	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1234

Claim 11	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1235

Claim 11	Disclosure
1236	 **Figure 1** Riani, Figure 1.

Claim 11	Disclosure
1237	 Figure 2 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1238

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

1239

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1240

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1241

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1242

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1243

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1244

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1245

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1247

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1248

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1249

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1250

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

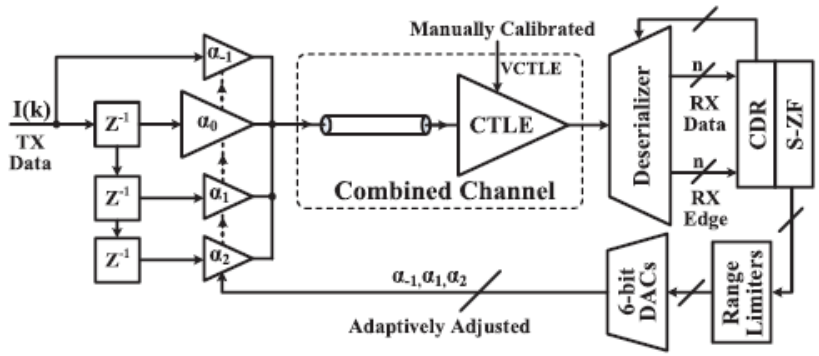
1251

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1252

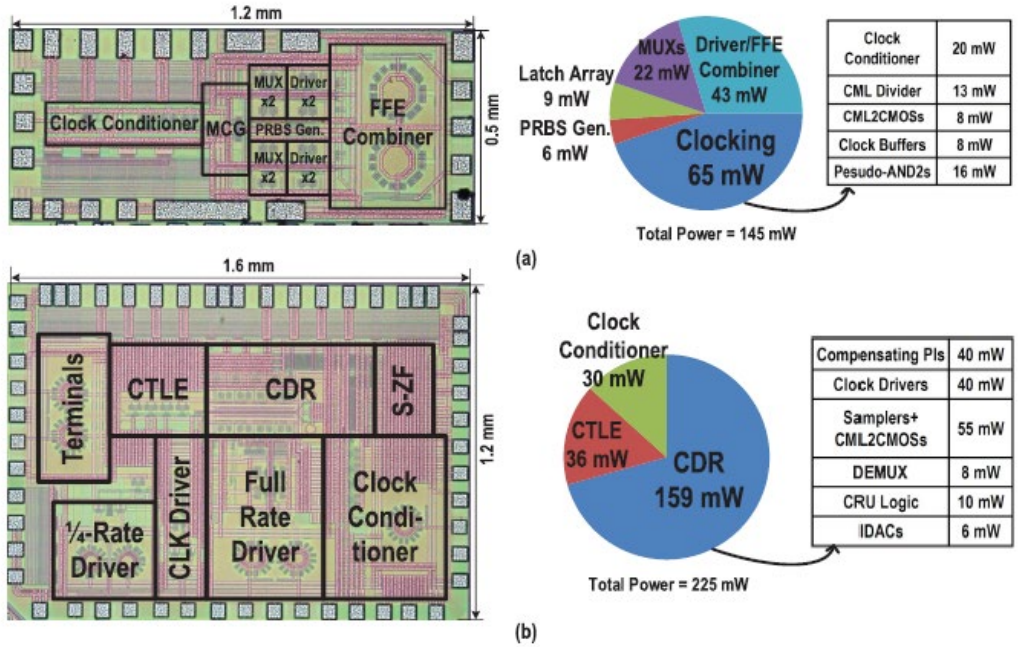
Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1253

Claim 11	Disclosure
1254	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1255

Claim 11	Disclosure
1256	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

Claim 11	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 1, the data transmission system 1 is in accordance with DP1.4, for example, and includes an upstream device 2, a downstream device 4, and an active cable 3 connecting the upstream device 2 and downstream device 4." Kobayashi, 2:41-45. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream

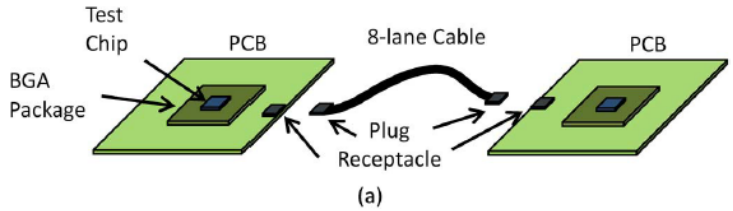
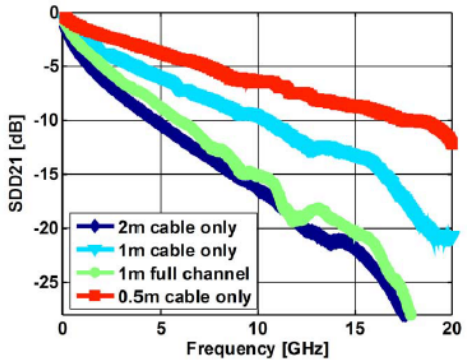
1257

Claim 11	Disclosure
	transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK" Kobayashi, 4:4-22.
11[b] packaging the first transceiver into a first connector configured to couple a first electrical input signal from a network interface port of a first host device to the first transceiver and a first electrical output signal from the first transceiver to the network interface port of the first host device; and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.**

1258

Claim 11	Disclosure
	Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1259

Claim 11	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

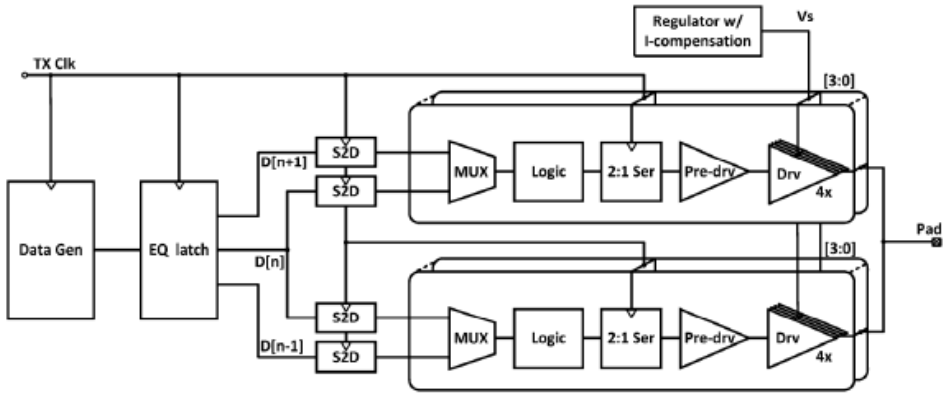
1260

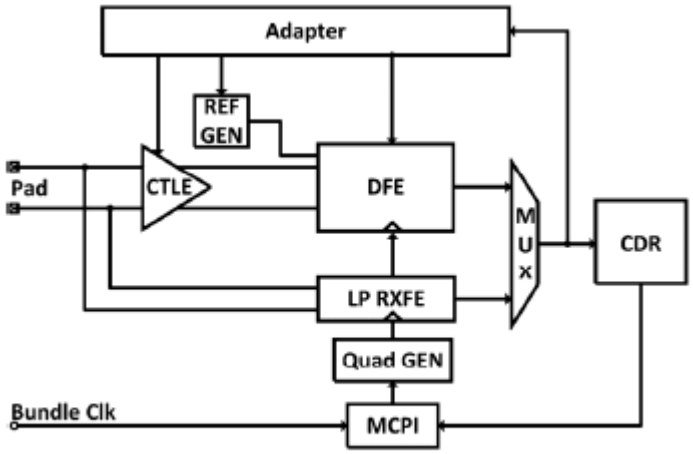
Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

1261

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1262

Claim 11	Disclosure
1263	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 11	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

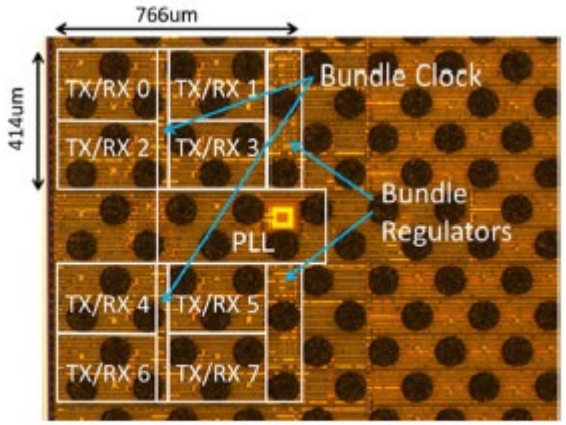
1264

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1265

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1266

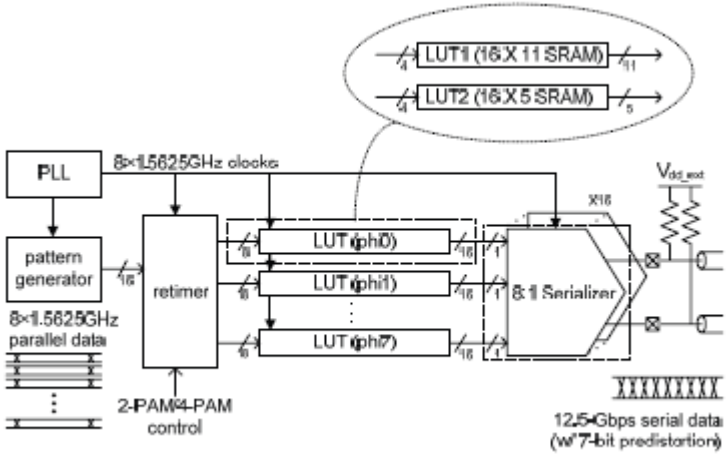
Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

1267

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1268

Claim 11	Disclosure
1269	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

Claim 11	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

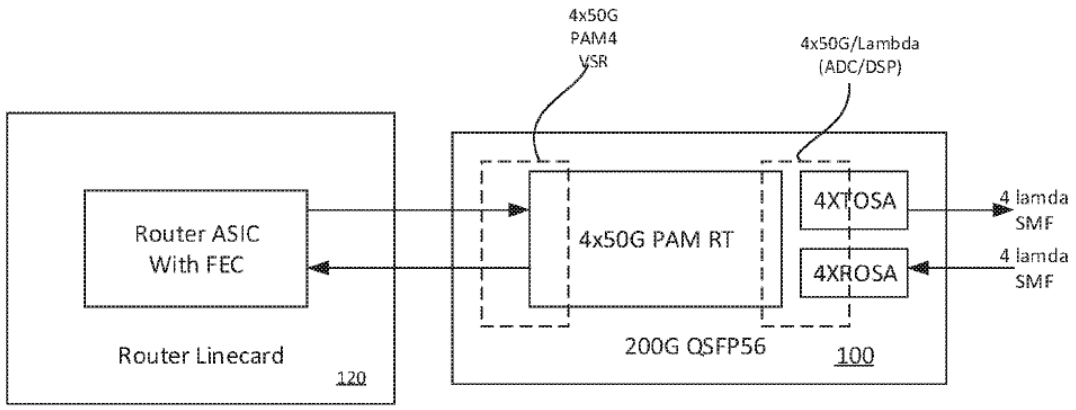
1270

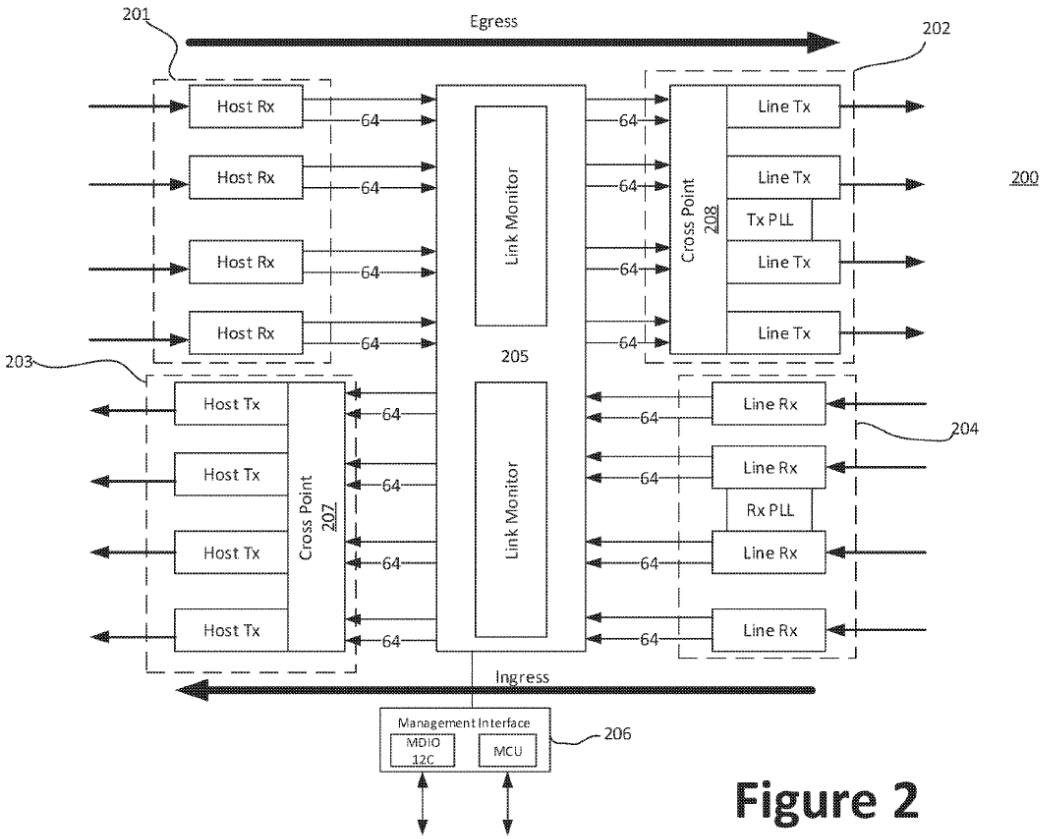
Claim 11	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1271

Claim 11	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1272

Claim 11	Disclosure
1273	 **Figure 1** Riani, Figure 1.

Claim 11	Disclosure
1274	 Figure 2 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1275

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1277

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1278

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1279

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1280

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1281

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1282

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

1283

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1284

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1285

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1286

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1287

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

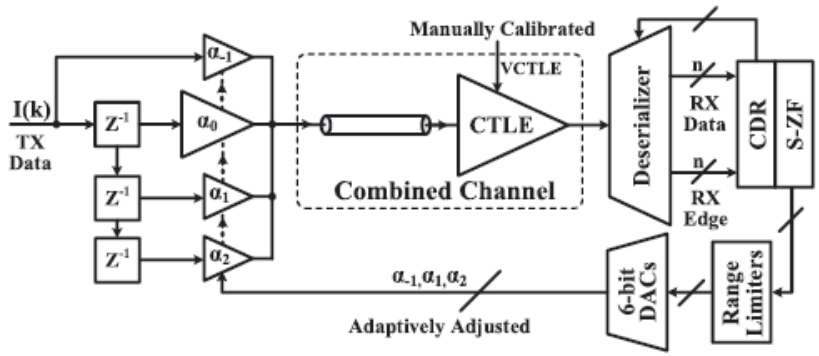
1288

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1289

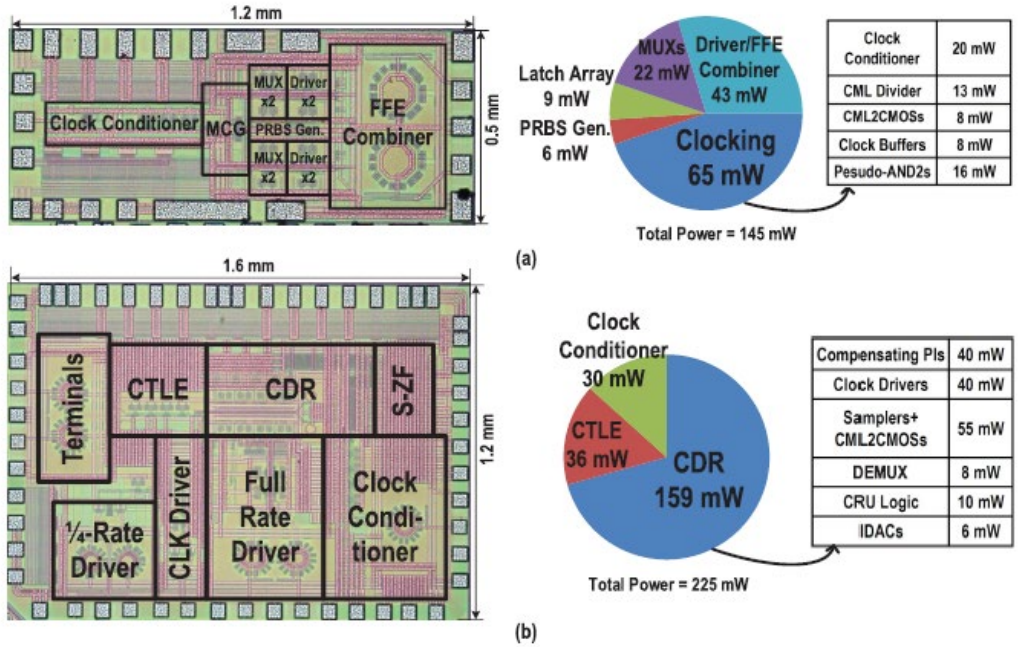
Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1290

Claim 11	Disclosure
1291	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1292

Claim 11	Disclosure
1293	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

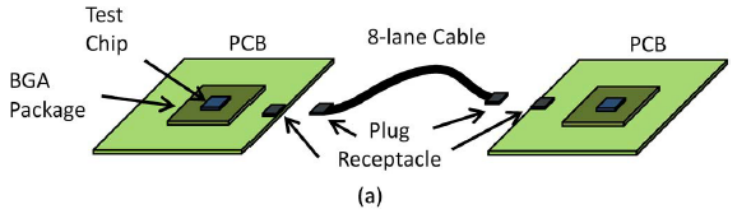
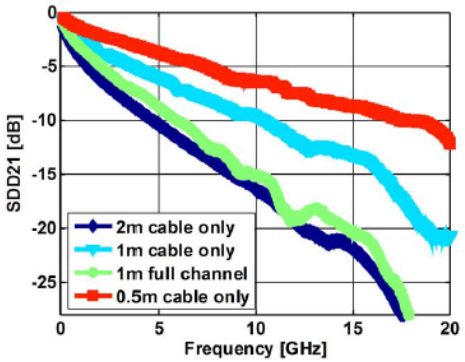
Claim 11	Disclosure
1294	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the

Claim 11	Disclosure
	IN_LT_SEQ_MARK" Kobayashi, 4:4-22.
11[c] packaging the second transceiver into a second connector configured to couple a second electrical input signal from a network interface port of a second host device to the second transceiver and a second electrical output signal from the second transceiver to the network interface port of the second host device,	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1295

Claim 11	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1296

Claim 11	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

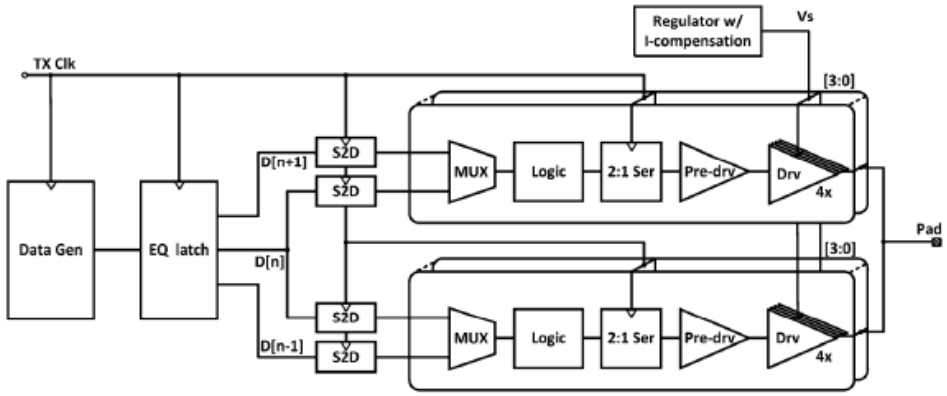
1297

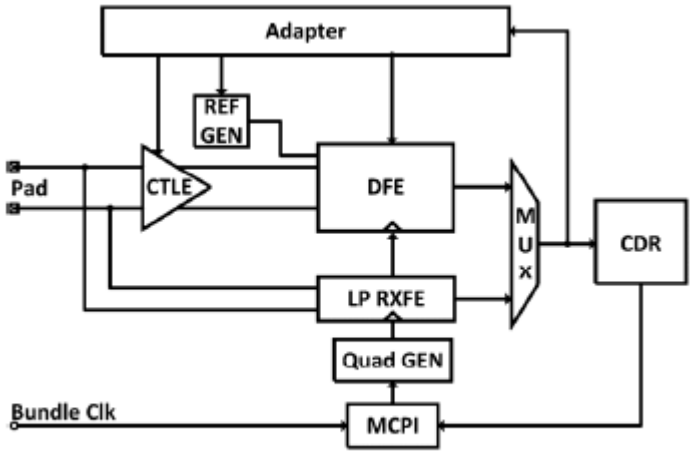
Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

1298

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1299

Claim 11	Disclosure
1300	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

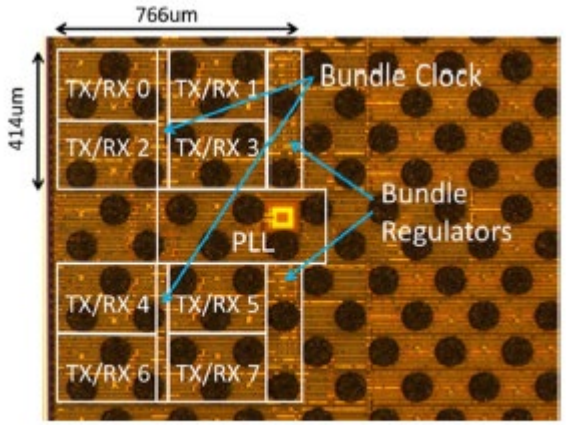
Claim 11	Disclosure
1301	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1302

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1303

Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

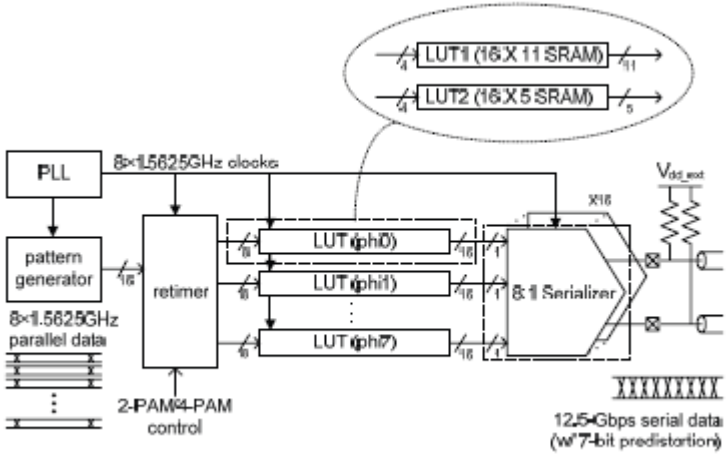
1304

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1305

Claim 11	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1306

Claim 11	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

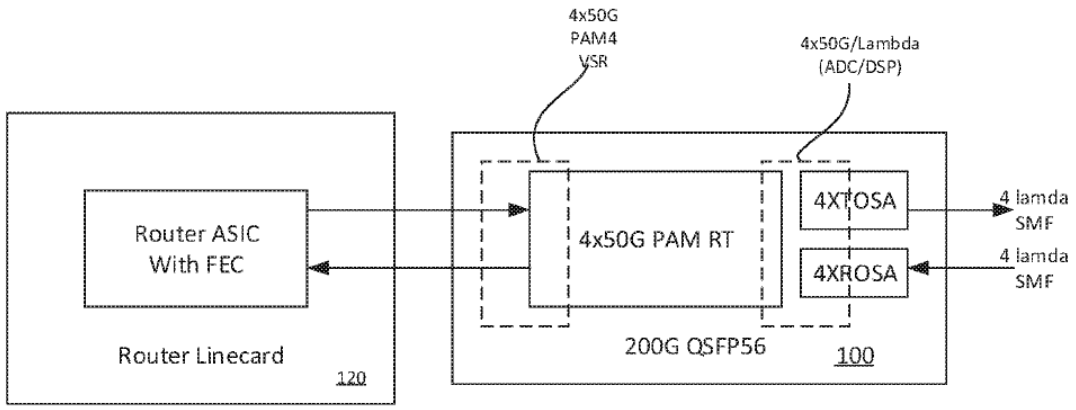
1307

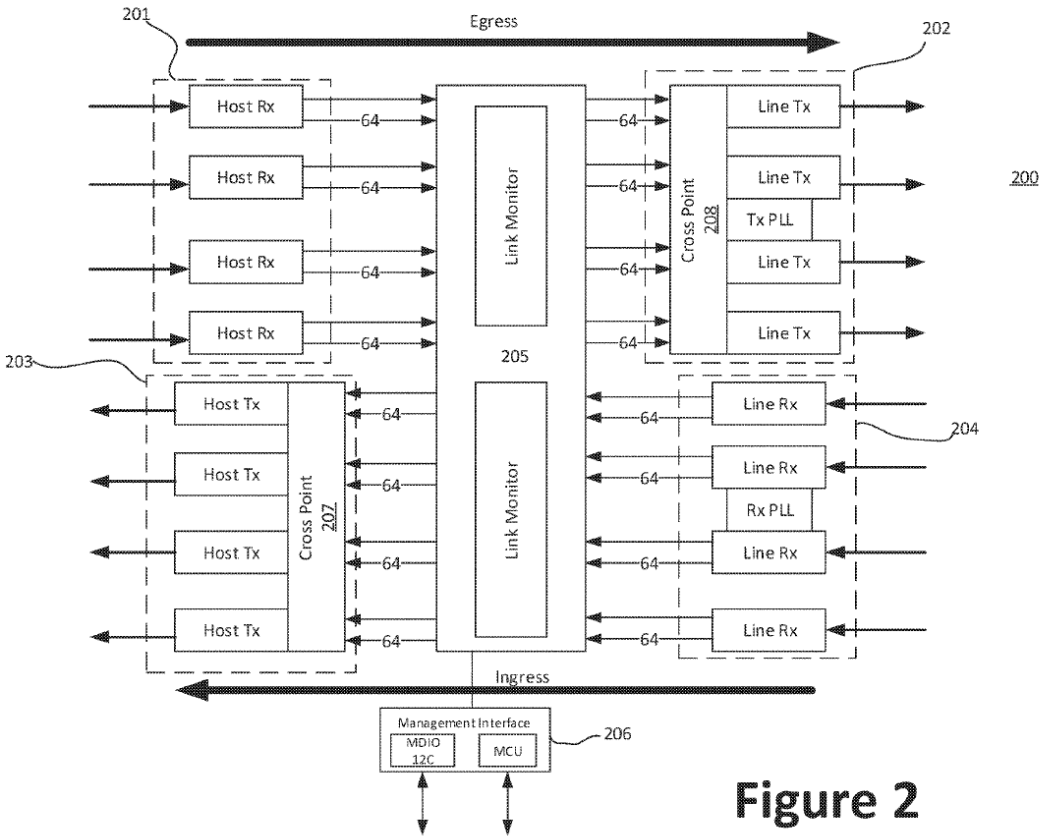
Claim 11	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1308

Claim 11	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1309

Claim 11	Disclosure
1310	 **Figure 1** Riani, Figure 1.

Claim 11	Disclosure
1311	 Figure 2 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1312

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1314

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1315

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1316

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1317

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1318

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1319

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

1320

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1321

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1322

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1323

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1324

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

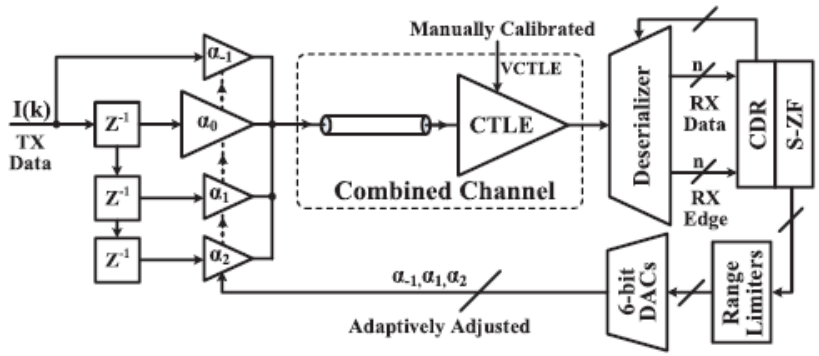
1325

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1326

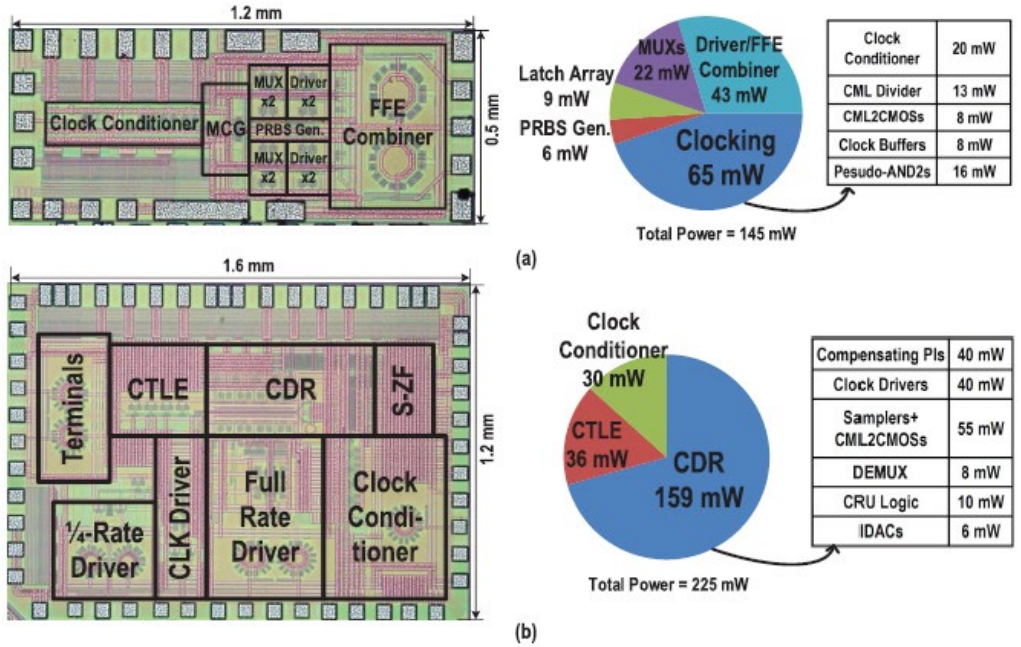
Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1327

Claim 11	Disclosure
1328	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1329

Claim 11	Disclosure
1330	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

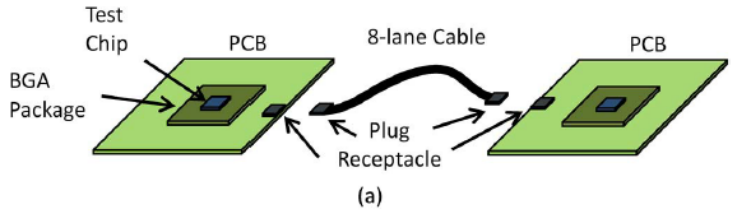
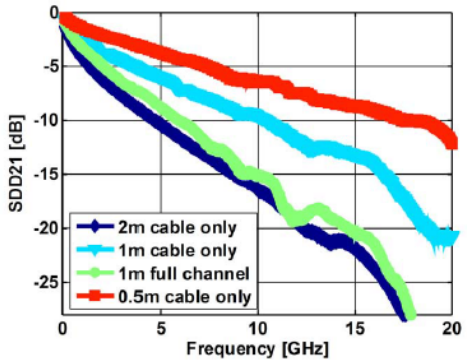
Claim 11	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The downstream device 4 receives the serial bit stream transmitted from the active cable 3 in the upstream facing RX 40a in the retimer 40. The retimer 40 performs a waveform shaping on the serial bit stream received by the upstream facing RX 40a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 40b to a sink 41. The sink 41 displays a video, for example, based on the received serial bit data." Kobayashi, 3:20-28.
11[d] the first and second transceivers being configured to perform clock and data recovery on the first and second	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations:

1331

Claim 11	Disclosure
electrical input signal to extract and re-modulate the first and second inbound data streams respectively as the first and second electrical transit signals conveying first and second transit data streams,	*Abstract*—**This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1332

Claim 11	Disclosure
1333	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

Claim 11	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

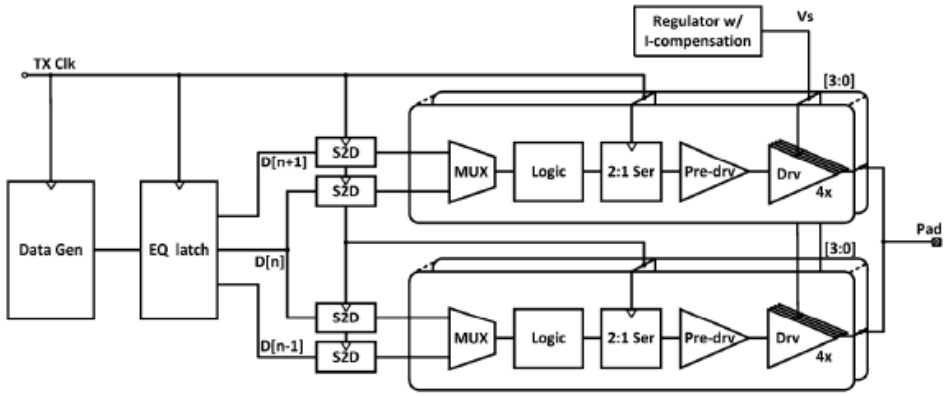
1334

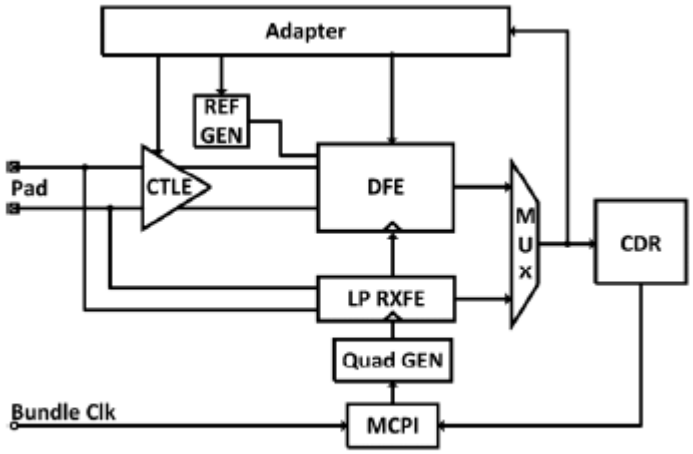
Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

1335

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1336

Claim 11	Disclosure
1337	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

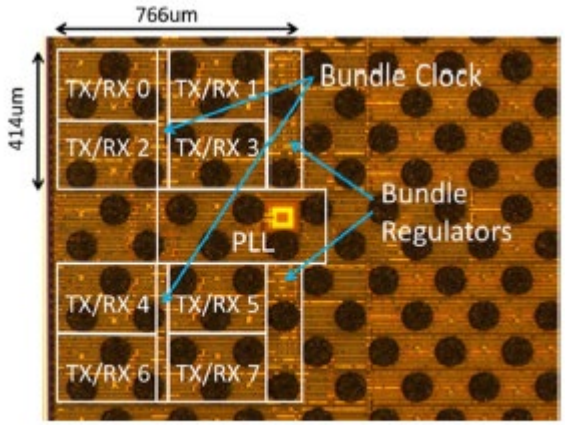
Claim 11	Disclosure
1338	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1339

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1340

Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

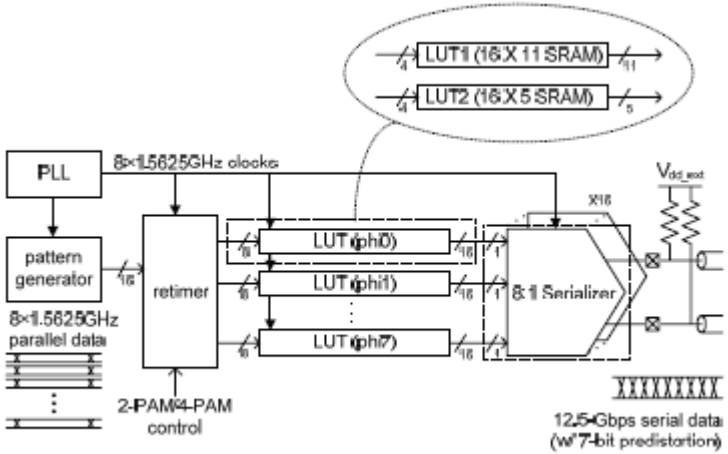
1341

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1342

Claim 11	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1343

Claim 11	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

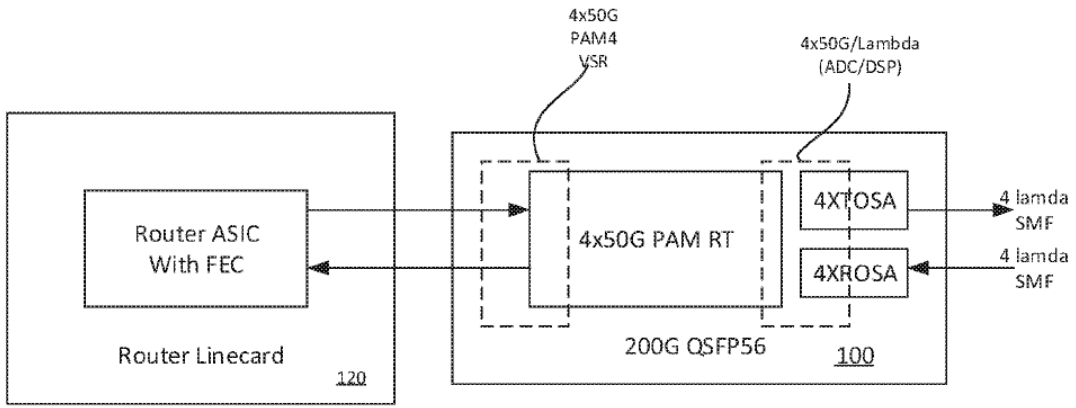
1344

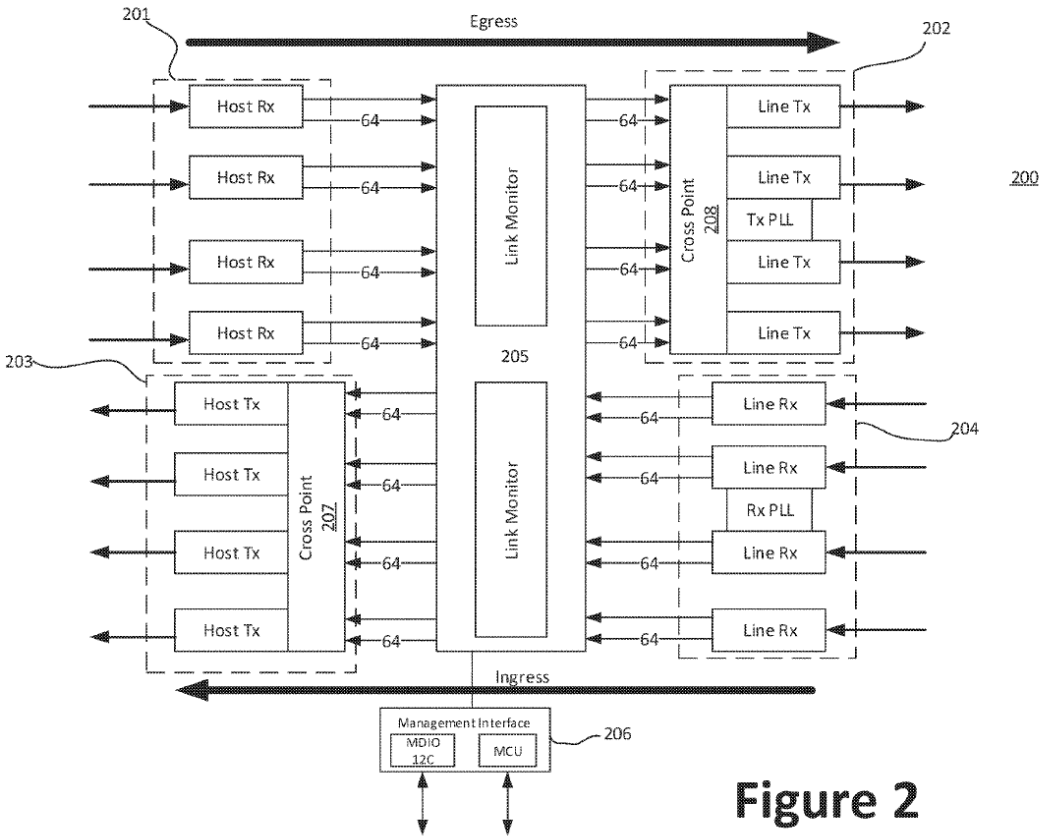
Claim 11	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1345

Claim 11	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1346

Claim 11	Disclosure
1347	 **Figure 1** Riani, Figure 1.

Claim 11	Disclosure
1348	 Figure 2 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1349

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

1350

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1351

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1352

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1353

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1354

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1355

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1356

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

1357

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1358

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1359

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1360

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1361

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

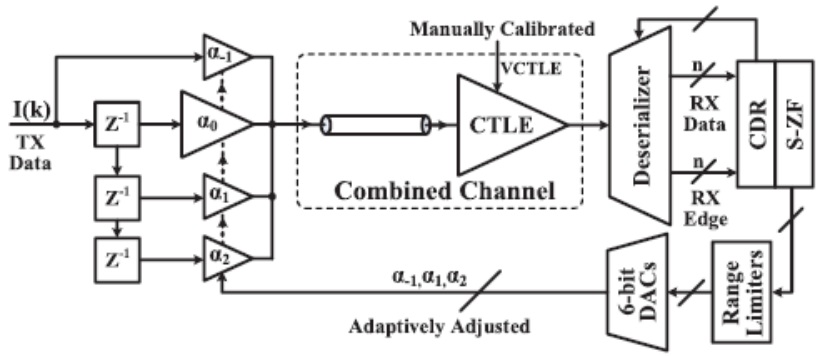
1362

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1363

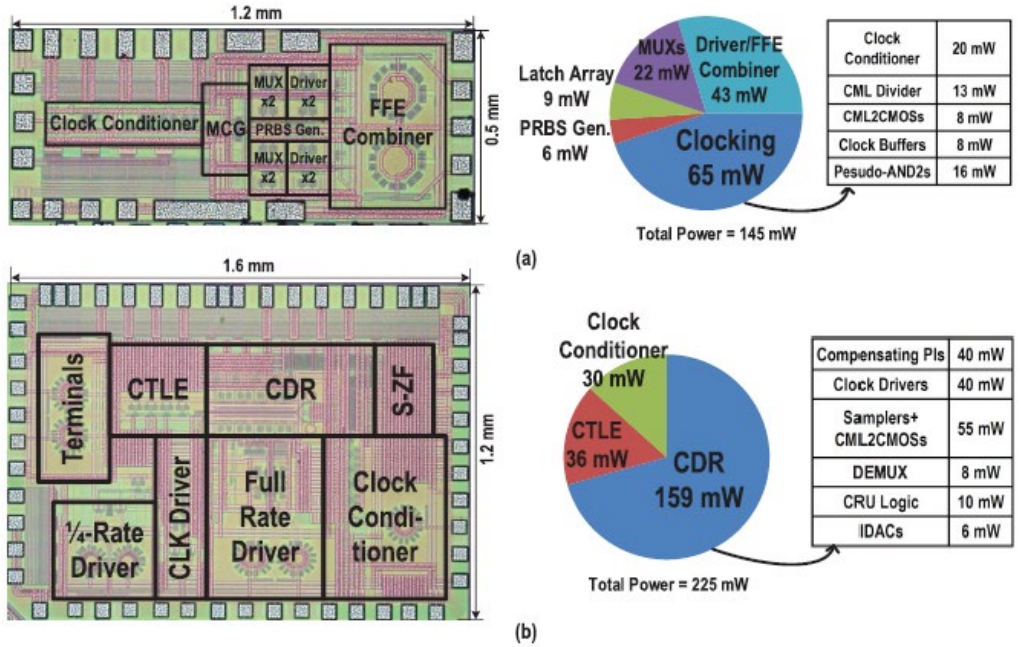
Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1364

Claim 11	Disclosure
1365	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1366

Claim 11	Disclosure
1367	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

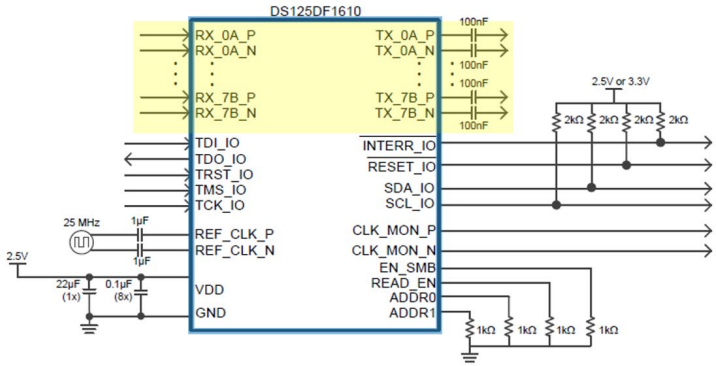
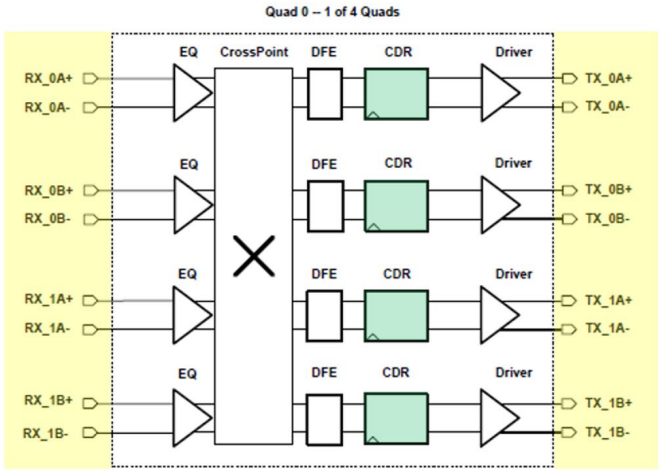
Claim 11	Disclosure
1368	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The active cable 3 receives the serial bit stream transmitted from the upstream device 2 in the upstream facing RX 30a in one retimer 30. The retimer 30 performs a waveform shaping on the serial bit stream received by the upstream facing RX 30a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 30b to the other retime 31. The retimer 31 receives the serial bit stream transmitted from the retimer 30 in the upstream facing RX 31a. The retimer 31 performs a waveform shaping on the serial bit stream received by the upstream facing RX 31a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 31b to the downstream device 4." Kobayashi, 3:5-19. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a

Claim 11	Disclosure
	TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver." Kobayashi, 4:4-9. "The exchange module 202 is configured to communicate capabilities information with a remote link partner including host loss information, RS-FEC ability information, and RS-FEC request information. The exchange module 202 may transmit local host loss information to the link partner and receive remote host loss information from the link partner. In some embodiments, the exchange module 202 may be included in or otherwise established by the autonegotiation module 104 and thus the capabilities information may be exchanged during execution of an autonegotiation protocol. In some embodiments, the exchange module 202 may be included in or otherwise established by the PMD/PMA module 106 and thus the capabilities information may be exchanged during a startup protocol performed by the PMD/PMA module 106." Lusted, paragraph 23. "The autonegotiation module 104 is configured to auto-negotiate line transmission speed, mode of operation, and other communication parameters with a link partner when the communication lane 102 is brought up. Additionally, in some embodiments the autonegotiation module 104 may be configured to exchange capabilities information with the link partner including host insertion loss information, RS-FEC ability information, and RS-FEC request information. The autonegotiation module 104 may encode the capabilities information, for example, in an autonegotiation base page or in a next page. Although illustrated in FIG. 1 as coupled between the PMD/PMA 106 and the communication lanes 102, in other embodiments the autonegotiation module 104 may be included at different positions in the Ethernet port logic 100. For example, in some embodiments the autonegotiation module 104 may be coupled between the PMD/PMA 106 and the FEC 108." Lusted, paragraph 17. "In view of the above transmission characteristics of twisted-pair, current Ethernet chips are usually equipped with preemphasis technique of a transmitting end and equalization technique of a receiving end

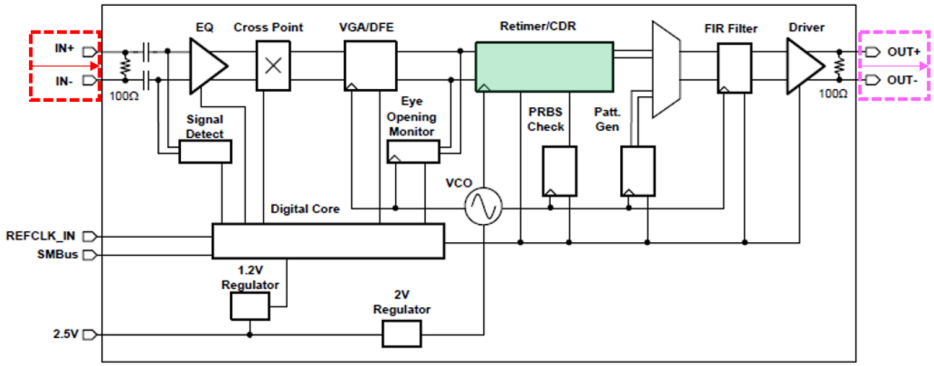
1369

Claim 11	Disclosure
	to compensate the difference in transmission attenuation of physical layer signals on different frequencies on the same medium during Ethernet baseband transmission. Main functions of the preemphasis technique of the transmitting end and the equalization technique of the receiving end include: amplifying signals on different frequencies differently at the transmitting end to compensate different transmission attenuation in the twisted pair; applying the equalization technique to different frequency components in signals received at the receiving end to compensate the attenuation." Yu, paragraph 34. "An embodiment of the present invention provides an equalizer setting apparatus, which uses a filter circuit 101, an analog-to-digital converter 102, a digital memory 103, a digital comparator 104 and a first digital accumulator 105, the filter circuit 101, the analog-to-digital converter 102, the digital memory 103, the digital comparator 104 and the first digital accumulator 105 are connected in sequence, the output end of the equalizer 201 is connected with the input end of the filter circuit 101, the output end of the first digital accumulator 105 is connected with the set value input end of the equalizer 201, the output end of the analog-to-digital converter 102 is connected with the input end of the digital comparator 104, so that the problems that the setting value of the equalizer is accurately set in the prior art, which is time-consuming and labor-consuming are solved, the method has the advantages of achieving accurate setting of the set value of the equalizer in a single time, fast, simple, convenient, accurate and low-cost manner and saving labor force to a certain extent." Yin, paragraph 26 (English translation). *DS125DF1610 Retimer Publication* The DS125DF1610 circuit is illustrated in the Simplified Schematic and Figure 1 below. The highlighted pins illustrate sixteen channels of data, with each channel having two differential (Positive and Negative) input pins (RX_0A_P/N through RX_7B_P/N) to receive the channel of data and corresponding output pins (TX_0A_P/N through TX_7B_P/N) to transmit the channel of data after processing. *See, e.g.*, DS125DF1610 at §§ 3, 7.2. Thus, DS125DF1610's retimer is a transceiver that, when used in Tang's

1370

Claim 11	Disclosure
	cable, receives one or more inbound data streams (on RX_0A_P/N through RX_7B_P/N), and transmits one or more transit data streams (via pins TX_0A_P/N through TX_7B_P/N). **Simplified Schematic**   **Figure 1. DS125DF1610 Simplified Cross Point Diagram** Figure 1 above shows four channels and illustrates that DS125DF1610's retimer processes each data

1371

Claim 11	Disclosure
	channel by, inter alia, performing "clock and data recovery (CDR)." *See id.* at §§ 3, 7.1. A channel's data path is shown in more detail in DS125DF1610 Figure 2 (below), which illustrates the input signal conveying an inbound data stream and the transit signal conveying a transit data stream in Tang's cable when DS125DF1610's retimer implements Tang's signal driver 44.  **Figure 2. DS125DF1610 Simplified Data Path Diagram** The Retimer/CDR circuit performs clock and data recovery on the electrical input signal to extract and re-modulate the inbound data stream for transit via the electrical conductors as a respective electrical transit signal conveying a transit data stream. *See id.* at § 7.3.6. The '252 Patent specification makes clear that the input signal from the host can be pre-processed before CDR is performed on it. An ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data for output as the transit data stream because after extracting the input signal's timing and using that timing to sample and extract the input signal's data in an analog-to-digital CDR process, the retimer generates a new data stream for equalization by the FIR Filter (a digital-to-analog process) and to be driven for output by the differential driver (which produces an analog signal with configurable output voltage levels) to convey the output data stream (transit data stream when outputting over the cable) as an analog signal over the cable conductors, where the new signal is a "re-modulation" of the data. *See id.* at §§ 7.3.6-7.3.9.

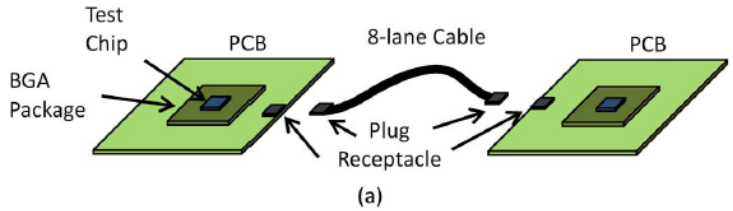
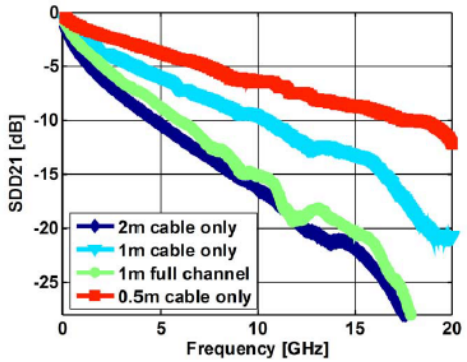
1372

Claim 11	Disclosure
11[e] configured to perform clock and data recovery on the second and first electrical transit signals to extract and re-modulate the second and first transit data streams as first and second outbound data streams conveyed by the first and second electrical output signals from the cable, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: ***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1373

Claim 11	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1374

Claim 11	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

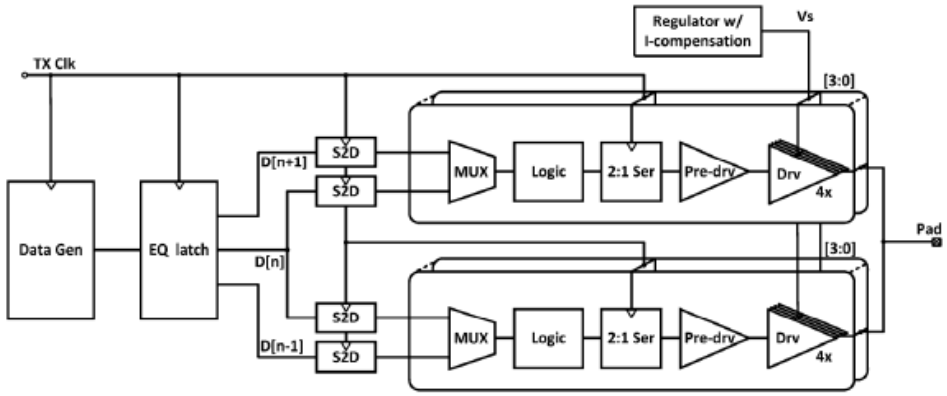
1375

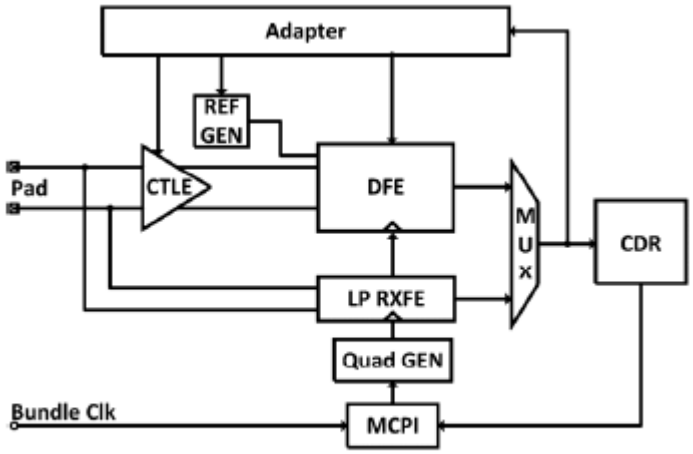
Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

1376

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1377

Claim 11	Disclosure
1378	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

Claim 11	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

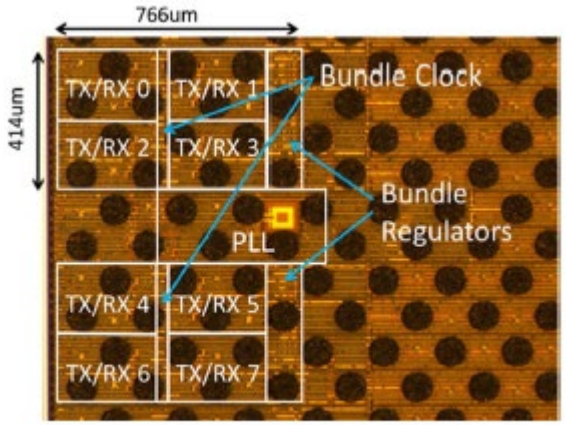
1379

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1380

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1381

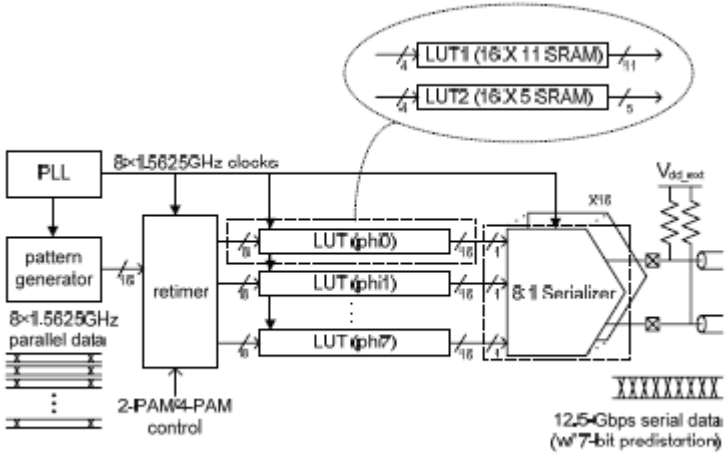
Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1383

Claim 11	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1384

Claim 11	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

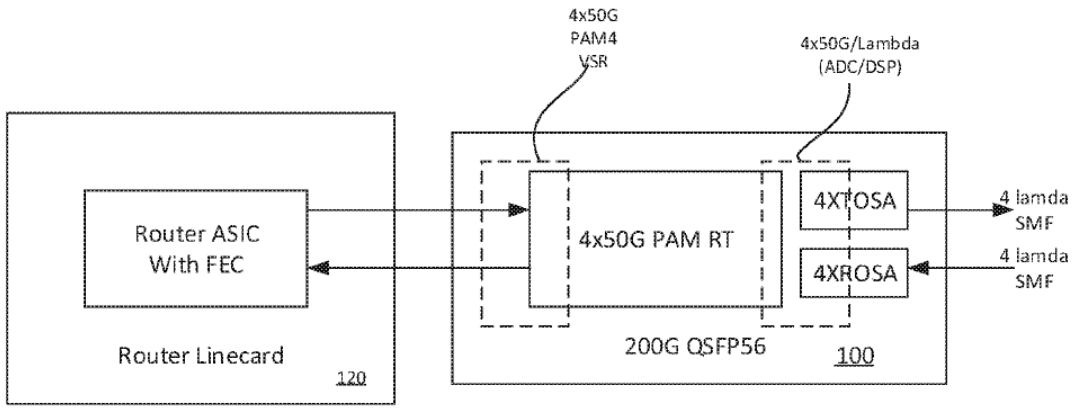
1385

Claim 11	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

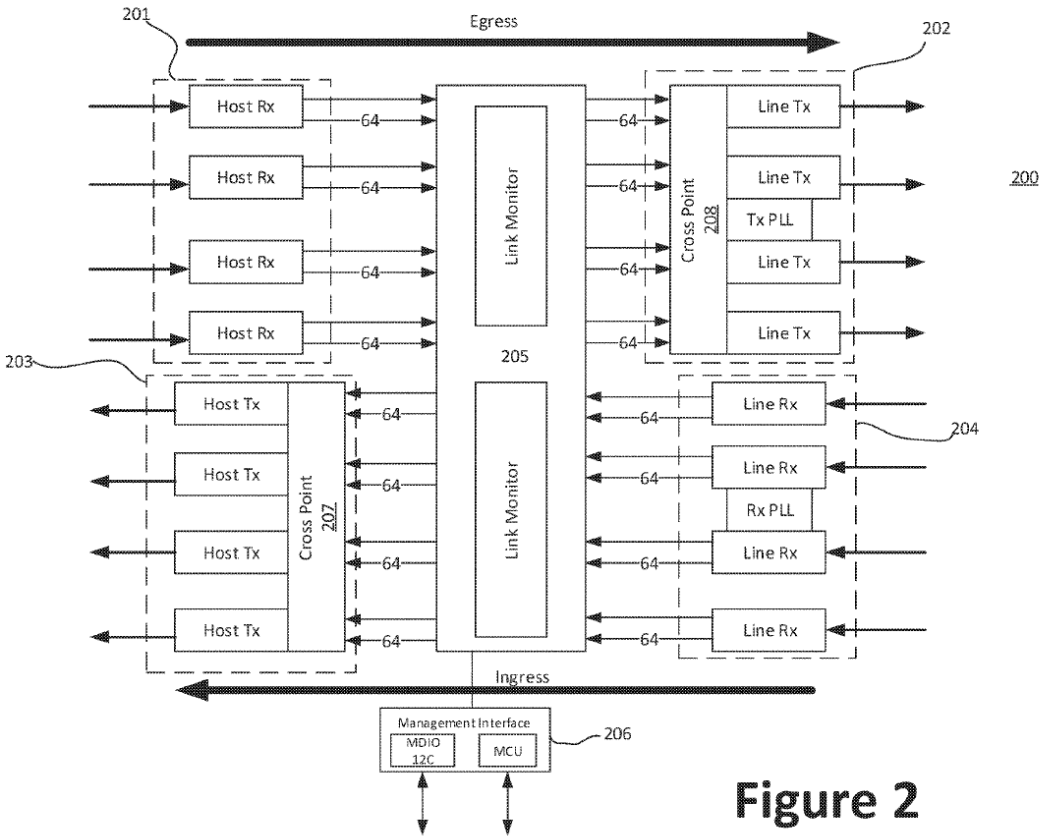
1386

Claim 11	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1387

Claim 11	Disclosure
	 **Figure 1** Riani, Figure 1.

1388

Claim 11	Disclosure
1389	 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1390

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1392

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1393

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1394

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1395

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1396

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1397

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

1398

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1399

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1400

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1401

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1402

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1403

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1404

Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1405

Claim 11	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

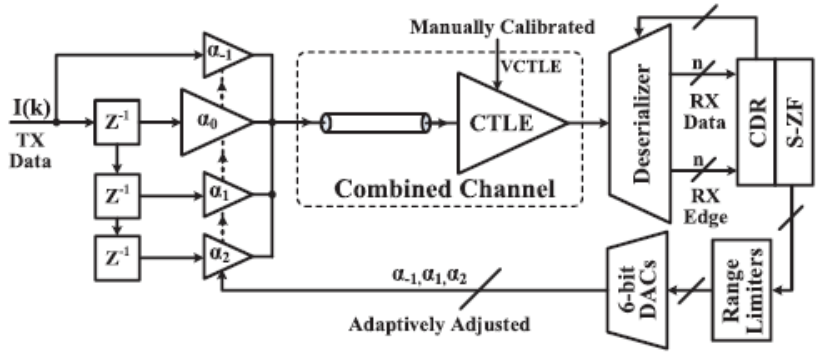
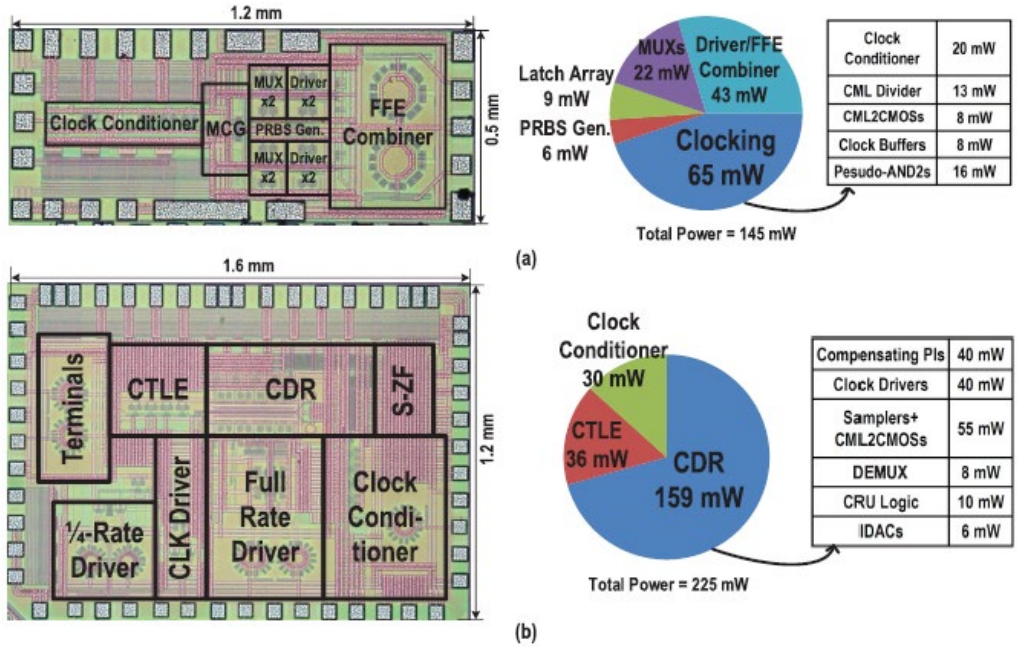


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1406

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1407

Claim 11	Disclosure
1408	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

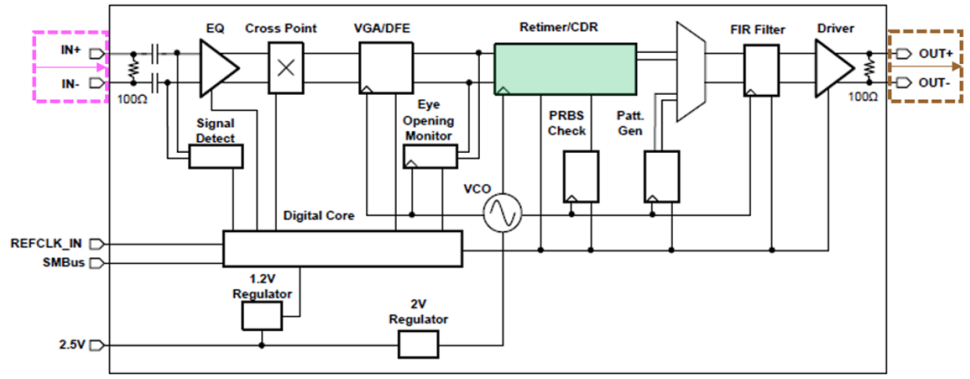
Claim 11	Disclosure
1409	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The active cable 3 receives the serial bit stream transmitted from the upstream device 2 in the upstream facing RX 30a in one retimer 30. The retimer 30 performs a waveform shaping on the serial bit stream received by the upstream facing RX 30a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 30b to the other retime 31. The retimer 31 receives the serial bit stream transmitted from the retimer 30 in the upstream facing RX 31a. The retimer 31 performs a waveform shaping on the serial bit stream received by the upstream facing RX 31a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 31b to the downstream device 4." Kobayashi, 3:5-19. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a

Claim 11	Disclosure
	TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver." Kobayashi, 4:4-9. "The exchange module 202 is configured to communicate capabilities information with a remote link partner including host loss information, RS-FEC ability information, and RS-FEC request information. The exchange module 202 may transmit local host loss information to the link partner and receive remote host loss information from the link partner. In some embodiments, the exchange module 202 may be included in or otherwise established by the autonegotiation module 104 and thus the capabilities information may be exchanged during execution of an autonegotiation protocol. In some embodiments, the exchange module 202 may be included in or otherwise established by the PMD/PMA module 106 and thus the capabilities information may be exchanged during a startup protocol performed by the PMD/PMA module 106." Lusted, paragraph 23. "The autonegotiation module 104 is configured to auto-negotiate line transmission speed, mode of operation, and other communication parameters with a link partner when the communication lane 102 is brought up. Additionally, in some embodiments the autonegotiation module 104 may be configured to exchange capabilities information with the link partner including host insertion loss information, RS-FEC ability information, and RS-FEC request information. The autonegotiation module 104 may encode the capabilities information, for example, in an autonegotiation base page or in a next page. Although illustrated in FIG. 1 as coupled between the PMD/PMA 106 and the communication lanes 102, in other embodiments the autonegotiation module 104 may be included at different positions in the Ethernet port logic 100. For example, in some embodiments the autonegotiation module 104 may be coupled between the PMD/PMA 106 and the FEC 108." Lusted, paragraph 17. "In view of the above transmission characteristics of twisted-pair, current Ethernet chips are usually equipped with preemphasis technique of a transmitting end and equalization technique of a receiving end

1410

Claim 11	Disclosure
	to compensate the difference in transmission attenuation of physical layer signals on different frequencies on the same medium during Ethernet baseband transmission. Main functions of the preemphasis technique of the transmitting end and the equalization technique of the receiving end include: amplifying signals on different frequencies differently at the transmitting end to compensate different transmission attenuation in the twisted pair; applying the equalization technique to different frequency components in signals received at the receiving end to compensate the attenuation." Yu, paragraph 34. "An embodiment of the present invention provides an equalizer setting apparatus, which uses a filter circuit 101, an analog-to-digital converter 102, a digital memory 103, a digital comparator 104 and a first digital accumulator 105, the filter circuit 101, the analog-to-digital converter 102, the digital memory 103, the digital comparator 104 and the first digital accumulator 105 are connected in sequence, the output end of the equalizer 201 is connected with the input end of the filter circuit 101, the output end of the first digital accumulator 105 is connected with the set value input end of the equalizer 201, the output end of the analog-to-digital converter 102 is connected with the input end of the digital comparator 104, so that the problems that the setting value of the equalizer is accurately set in the prior art, which is time-consuming and labor-consuming are solved, the method has the advantages of achieving accurate setting of the set value of the equalizer in a single time, fast, simple, convenient, accurate and low-cost manner and saving labor force to a certain extent." Yin, paragraph 26 (English translation). *DS125DF1610 Retimer Publication* The DS125DF1610 retimer performs CDR for respective transit signals conveying transit data streams received from electrical conductors to generate output signals conveying outbound data streams that exit cable 10 to respective hosts, as shown in DS125DF1610 Fig. 2 below. For similar reasons to those discussed in claim limitation 1[c], an ordinarily skilled artisan understood that DS125DF1610's retimer remodulates the data as the outbound data after extracting the input signal's data and generating a newly-

1411

Claim 11	Disclosure
	timed signal that is a "re-modulation" of the data.  **Figure 2. DS125DF1610 Simplified Data Path Diagram**
11[f] each configured to employ fixed, cable-independent, equalization parameters for clock and data recovery on the respective electrical input signals and for generating the respective electrical output signals.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: `Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers.` CMIS, 1.

1412

Claim 11	Disclosure
	1.4.4 MgmtInit State The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

1413

Claim 11	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1414

Claim 11	Disclosure
1415	Table 10- Control field dependency on Explicit Control bit Control field | Explicit Control value Tx Adaptive Input Eq Enable | 1 Tx Adaptive Input Eq Recall | 1 or 0 Tx Input Eq control | 1 Tx CDR control | 1 Rx CDR control | 1 Rx Output Eq control, pre-cursor | 1 Rx Output Eq control, post-cursor | 1 Rx Output Amplitude control | 1 CMIS, 38. 1.5.4.1 Control Set Usage There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

Claim 11	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1416

Claim 11	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1417

Claim 11	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))

The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1418

Claim 11	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: **Table 14- Rx Output Amplitude Codes** Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1419

Claim 11	Disclosure
1420	1.7.4.10 Implemented Signal Integrity Controls Table 46 describes the implemented signal integrity controls. Table 46- Implemented Signal Integrity Controls (Page 01h) CMIS, 71.

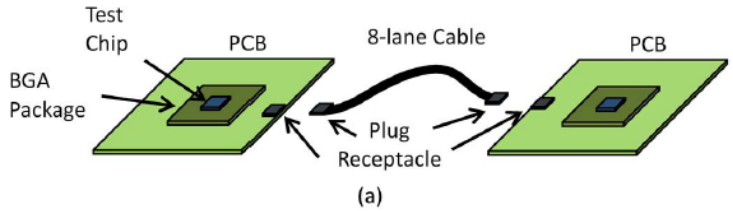
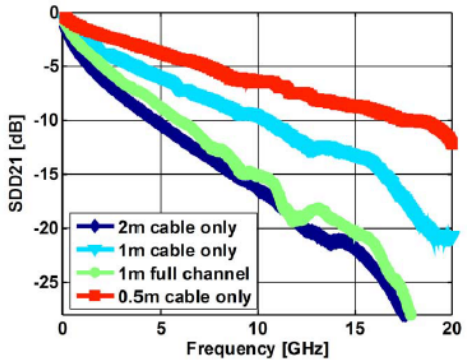
Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

Claim 11	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1421

Claim 11	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1422

Claim 11	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

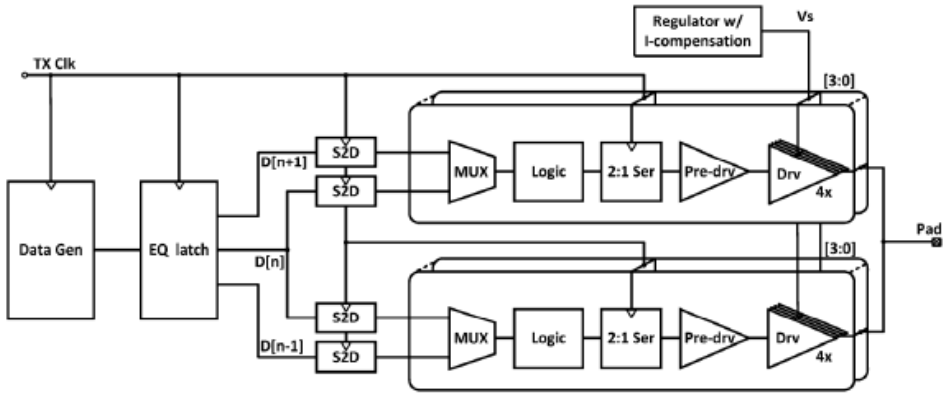
1423

Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

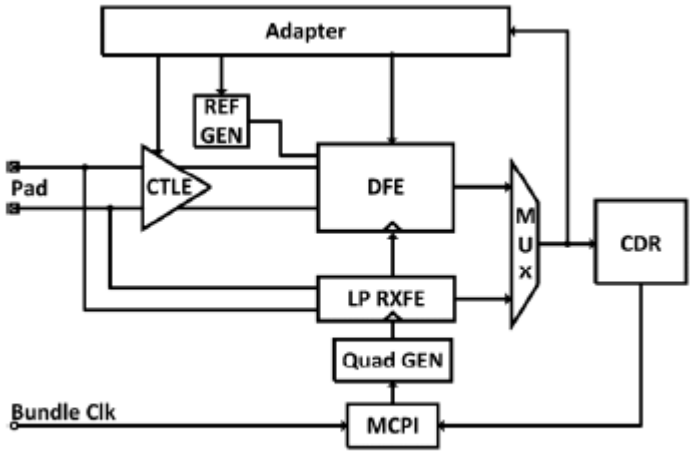
1424

Claim 11	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1425

Claim 11	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1426

Claim 11	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

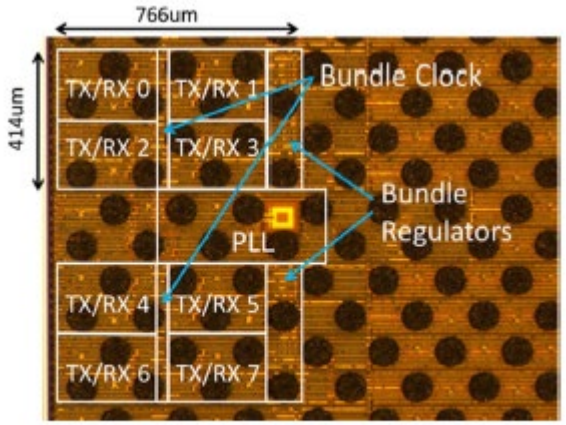
1427

Claim 11	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1428

Claim 11	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1429

Claim 11	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

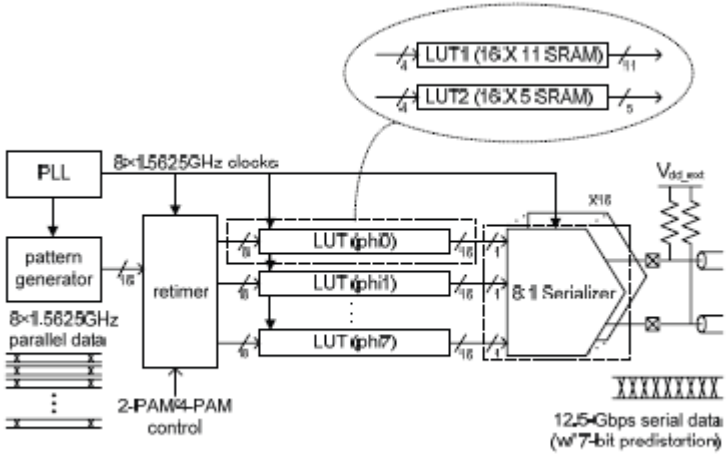
1430

Claim 11	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1431

Claim 11	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1432

Claim 11	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

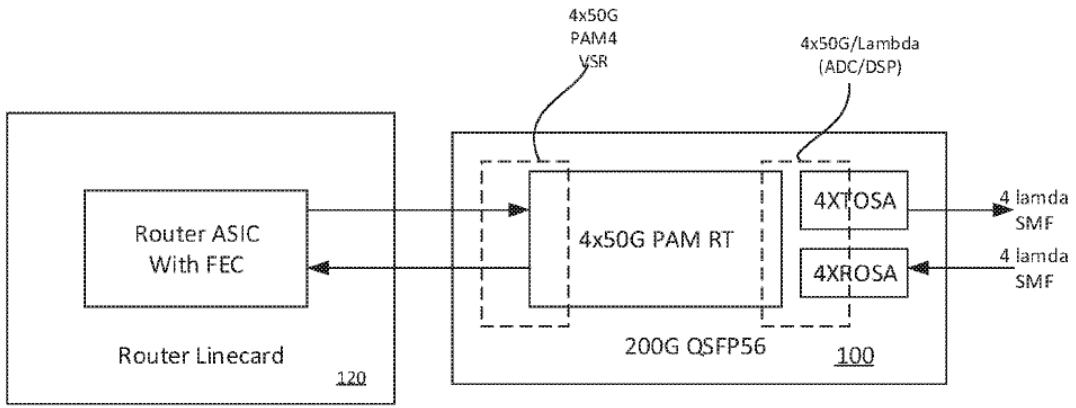
1433

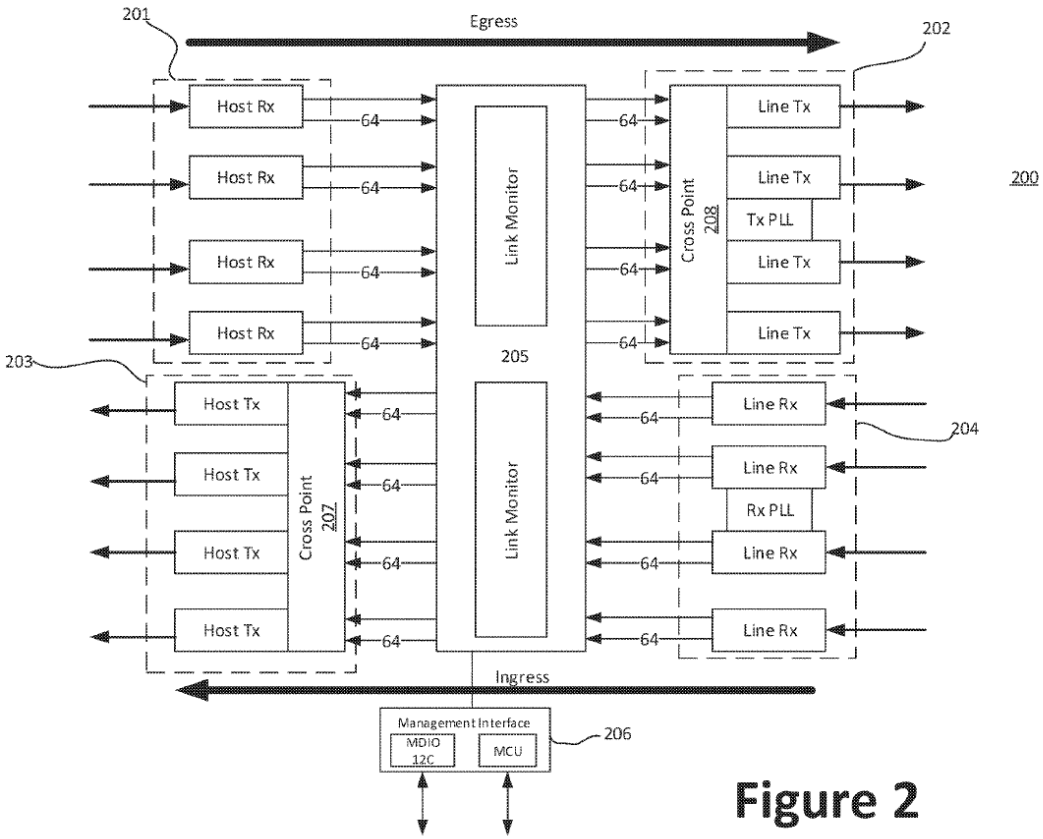
Claim 11	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1434

Claim 11	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1435

Claim 11	Disclosure
1436	 **Figure 1** Riani, Figure 1.

Claim 11	Disclosure
1437	 Figure 2 Riani, Figure 2.

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host system and a line system. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 1:45-52. It is to be appreciated that embodiments of the present invention provide many advantages over conventional techniques. Among other things, the re-timer module allows flexible implementation of high-speed communication network systems. For example, various fast communication protocols such as PAM-4 communication techniques is supported for optical communication networks. Re-timer modules according to embodiments of the present invention can be readily adopted into data centers and provide performance improvement and cost reduction. Riani, 3:3-13.

1438

Claim 11	Disclosure
	The present invention is directed to data communication systems and techniques thereof. More specifically, embodiments of the present invention provide a retimer module that includes plurality of communication lanes for interfacing with a host and a line. The retimer module includes a link monitor and cross point sections. The retimer also includes a management interface module. There are other embodiments as well. Riani, 3:60-67.

Claim 11	Disclosure
	FIG. **1** is a simplified diagram illustrating a communication module **100** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. Among other features, the communication module can be readily deployed at data centers that require 100 Gb/s or 200 Gb/s Ethernet. For data center applications, it is desirable to have optical modules that support 200 Gb/s and comply to the QSFP56 form factor, with lower cost and lower power compared to existing modules. To reduce the cost of the modules, efforts have been made to reduce the number of optical components by decreasing the number of wavelengths used to transmit data at 200 Gb/s, which requires advanced modulation techniques. Pulse amplitude modulation (PAM) is emerging as a preferred modulation technique for such application. For example, the "IN015050-CC02" module is a four-port duplex re-timer for PAM-4 signals between the line and host interfaces. The exemplary communication module illustrated in FIG. **1** is an application diagram of a QSFP56 module using the IN015050-CC02. The communication module **100** includes a transmitter interface (e.g., "4XTOSA") for transmitting four 50G data lanes and a receiver interface (e.g., "4XROSA") for receiving four 50G data lanes. Various DSP processes are performed for transmitted and received data. Communication module **100** interfaces with router line card **120** using four 50G PAM4 communication lanes. More detailed descriptions are provided below.

1440

Claim 11	Disclosure
	Riani, 4:43-5:5. FIG. **2** is a simplified diagram illustrating a re-timer module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. As shown in FIG. **2**, retimer module **200** includes host Rx section **201** and host Tx section **203** that interface with a host system (not shown). As respective names and arrows imply, host Rx section **201** receives data from the host system, and the host Tx section **203** transmits data to the host system. Retimer module **200** also includes line Tx section **202** and line Rx section **204**. For data transmitted from host system to line, a data path (i.e., egress data path) is formed by host Rx section **201** and line Tx section **202**, with link monitor section **205** in the middle. Riani, 5:6-20.

1441

Claim 11	Disclosure
	It is to be appreciated that the retimer module **200** can be configured to operate in various modes and can support a wide variety of standards. For example, the following data rates (e.g., from 10.3 Gbps to 28.125 Gbps) are supported: 20.625 GBaud, 25.78125 GBaud, and 26.5625 GBaud PAM-4 10.3125 Gbit/s, 25.78125 Gbit/s or 26.5625 Gbit/s NRZ data 28.125 Gbps NRZ or 28.125 GBaud PAM-4 data The operation of retimer module **200** is in large part controlled by management interface **206**, which is coupled to link monitor **205**. As explained above, link monitor **205** is positioned between the host side sections (**201** and **203**) and line side sections (**202** and **204**), and it is configured to monitor the data traffic going through retimer module **200** to ensure proper operation. In various embodiments, link monitor **205** is configured to provide diagnostic functions and analyzes a specific type of diagnostic data. Management interface **206** comprises a management data input/output (MIDIO) interface and a microcontroller unit (MCU) as shown. In a specific implementation, retimer module **200** has 0.75, 1V and 1.95V (1.6V in non-EML mode) power supply rails. Riani, 5:57-6:12.

1442

Claim 11	Disclosure
	In various embodiments, retimer module **200** works in conjunction with a storage unit, which may be implemented with an EEPROM device that is connected through an SPI interface. The EEPROM device stores the startup and DSP sequencing codes that are to be read by retimer module **200** at startup. This enables power up of the device independent from the power up of the entire module. It is to be appreciated that the initiation sequence automatically enables all the internal design blocks to their proper operating states. For example, many registers associated with the internal blocks do not have to be written externally through the MDIO/I2C interface. Riani, 6:25-36. Re-timer modules, as illustrated in FIG. **2**, include line side and host side. The line side includes line receivers. For example, the receiver module **204** in FIG. **2** includes four line side receivers. For example, four ADC/DSP-based line side receivers each operates at 28 Gbaud/s data rate. Each of the line receivers includes one or more equalizers. For example, an equalizer can be implemented as a combination of digital FFE and DFE. Line receivers may additionally incorporate automatic gain control loop to load the receiver ADC appropriately. Line receivers may also support eye monitor function on the receiver and slicer threshold adjust for each PAM4 level. Riani, 6:63-7:7.

1443

Claim 11	Disclosure
	FIG. **3** is a simplified diagram illustrating a line receiver module according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The line receiver module **300** is implemented as a part of retimer module **200**, and it is configured on the line side for receiving optical signals. Receiver module **300** includes four receivers (shown as stacked receiver slices) for supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL is configured to serve four lanes. As an example, the high-speed analog front end (AFE) section of the data path consists of four lanes that each handles up to 56 Gbps of PAM-4 data. In various implementations, the AFE section comprises equalizer and variable gain amplifier (VGA), which are not shown in FIG. **3**. In a specific embodiment, the VGA is AC-coupled to the line signal via external capacitors and has about 50 ohms of internal resistance. The AFE includes a continuous time linear equalizer (CTLE) that is designed to provide high performance boost. For example, the auto gain control and CTLE functions can be configured to be auto-adaptive. An on-chip low dropout regulator is used to provide a regulated voltage to function blocks of the AFE. Riani, 7:15-39.

1444

Claim 11	Disclosure
	FIG. **4** is a simplified diagram illustrating a line side transmitter module **400** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **400** is implemented as a component (e.g., section **202**) of retimer **200** illustrated in FIG. **2**. The transmitter module **400** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). A common PLL provides services all four lanes. The data path receives up to four lanes of PAM-4 data from the core and prepares them for optical transmission. In various embodiments, transmitter module **400** operates in quad-lambda mode at 28.125 GBd, transmitting data streams with PAM-4 modulation (56.25 Gb/s). Data are transmitted as optical signals. For example, the AFE block of transmitter module **400** includes an externally modulated laser for optical data transmission, and a single transmit PLL is shared among the four transmit lanes. For example, the four transmit lanes are shown in FIG. **4** as four slices. In addition to data transmission (i.e., from host Rx to line Tx), transmitter module **400** also operates in shallow line loopback mode (e.g., operating at PAM-4 to PAM-4). Riani, 8:37-60.

1445

Claim 11	Disclosure
	The PAM-4 encoding/Gray coding block follows the error injection block. For example, at this block, Gray coding, polarity inversion, PAM-4 encoding, PAM-4 MSB/LSB swapping, and/or other coding processes may be performed. Next, Tx DSP block performs various DSP processes on the encoded data. At the AFE block, data are converted to optical signal for transmission. If the transmitter module **400** is operating in a testing mode, test patterns from the test pattern memory block (instead of actual data) are transmitted. Riani, 9:38-47.

Claim 11	Disclosure
	FIG. **6** is a simplified block diagram illustrating a transmitter at the host side transmitter module **600** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, transmitter module **600** is implemented as a component (e.g., section **203**) of retimer **200** illustrated in FIG. **2**. The transmitter module **600** includes a quad lane transmitter supporting data streams of up to 28.125 GBd with PAM-4 modulation (56.25 Gb/s). Four PLLs service the 8 transmit lanes allowing for 4 independent output frequencies. The data path receives up to 8 lanes of NRZ data or 4 lanes of PAM-4 data from the core and prepares them for transmission. The cross point **601** supports the receive data streams to be transmitted on any of host transmit lanes. The received data are first processed by the FIFO buffer as shown. In various embodiments, the AM framer supports 25GE, 50GE, 100GE, 200GE and 400GE protocols. The AM framer provides protocol lock status. Data are encoded for transmission using PAM-4 communication protocol. For example, PAM-4 encoding block provides processes for polarity inversion, PAM-4 encoding, and PAM-4 MSB/LSB swap. The transmitter module **600** further provides squelch control. The pattern generator block generates test patterns for each communication. The pattern generator logic supports PRSB58/31/23/16/15/13/11/9/7, PRBS13Q, SSPRQ, and/or other formats. For example, PRBS can be generated independently on PAM-4 MSBs and LSBs. In a specific embodiment, 2×63 bit programmable-repeatable fixed patterns are used. Square wave, JP03A, JP03B, transmitter linearity test patterns, and CID jitter tolerance pattern can be used as well. The test pattern memory can be configured in a binary mode, which provides up to 32,768 bits for four transmitter lanes, and each transmitter can have a unique pattern with identical lengths. The test pattern memory can operate in linear mode: up to 16,384 UI (each 8-bits) on one transmitter lane. The host side transmitter module **600** may additionally include features such as error injection and continuous identical symbol detector. Riani, 11:4-45.

1447

Claim 11	Disclosure
	It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, re-timer form factors are suitable for QSFP56 200G modules. Certain re-timers implement PAM-4 electrical outputs and includes an integrated EML driver. Re-timers can support both 56.25 Gb/s data streams with PAM-4 modulation and 28.125 Gb/s data streams with NRZ modulation. Lane-based diagnostic link monitor can be implemented to analyze receive error statistics and estimate BER. Riani, 12:19-30.

1448

Claim 11	Disclosure
	***Abstract*—This paper presents a 40-Gb/s transmitter (TX) and receiver (RX) chipset for chip-to-chip communications in a 65-nm CMOS process. The TX implements a quarter-rate multi-multiplexer (MUX)-based four-tap feed-forward equalizer (FFE), where a charge-sharing-effect elimination technique is introduced into the 4:1 MUX to optimize its jitter performance and power efficiency. The RX employs a two-stage continuous-time linear equalizer as the analog front end and integrates a low-cost sign-based zero-forcing engine relying on edge-data correlation to automatically adjust the tap weights of the TX-FFE. By embedding low-pass filters with an adaptively adjusting bandwidth into the data-sampling path and adopting high-linearity compensating phase interpolators, the clock data recovery achieves both high jitter tolerance and low jitter generation. The fabricated TX and RX chipset delivers 40-Gb/s PRBS data at BER $< 10^{-12}$ over a channel with >16-dB loss at half-baud frequency, while consuming a total power of 370 mW.** ***Index Terms*—4:1 multiplexer (MUX), 40 Gb/s, charge-sharing effect, clock data recovery (CDR), continuous-time linear equalizer (CTLE), edge-data correlation, feed-forward equalizer (FFE), jitter suppression, jitter tolerance (JTOL), low-pass filters (LPFs), sign-based zero-forcing (S-ZF), transmitter (TX) and receiver (RX) chipset.** Zheng, 2963.

1449

Claim 11	Disclosure
	To alleviate these difficulties and provide potential solutions for ultra-high-speed transceiver design, this paper presents a 40-Gb/s quarter-rate SerDes TX and receiver (RX) chipset. The remainder of this paper is organized as follows. Section II describes the TX chip, mainly focusing on the improved 4:1 MUX. Section III illustrates the RX chip, where the CDR performance is enhanced by introducing jitter-suppression filters and adopting high-linearity compensating PIs. In Section IV, a low-cost sign-based zero-forcing (S-ZF) adaptation algorithm relying on edge-data cross correlation is designed to achieve adaptive tap-weight adjustment for the TX-feed-forward equalizer (FFE). Section V gives the experimental results and performance comparison, and Section VI concludes this paper. Zheng, 2964.

1450

Claim 11	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

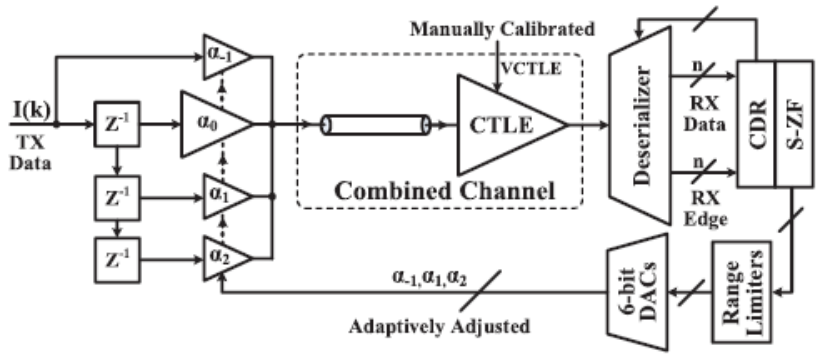
1451

Claim 11	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1452

Claim 11	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1453

Claim 11	Disclosure
	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1454

Claim 11	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 1.2-V supply at 40 Gb/s. The TX chip occupies an area of 0.6 mm^2 and consumes a total power of 145 mW with a 400-mV single-end swing. The RX chip occupies 1.92 mm^2 (including the testing circuits) and dissipates 225-mW power (excluding the testing circuits). Zheng, 2972.

1455

Claim 11	Disclosure

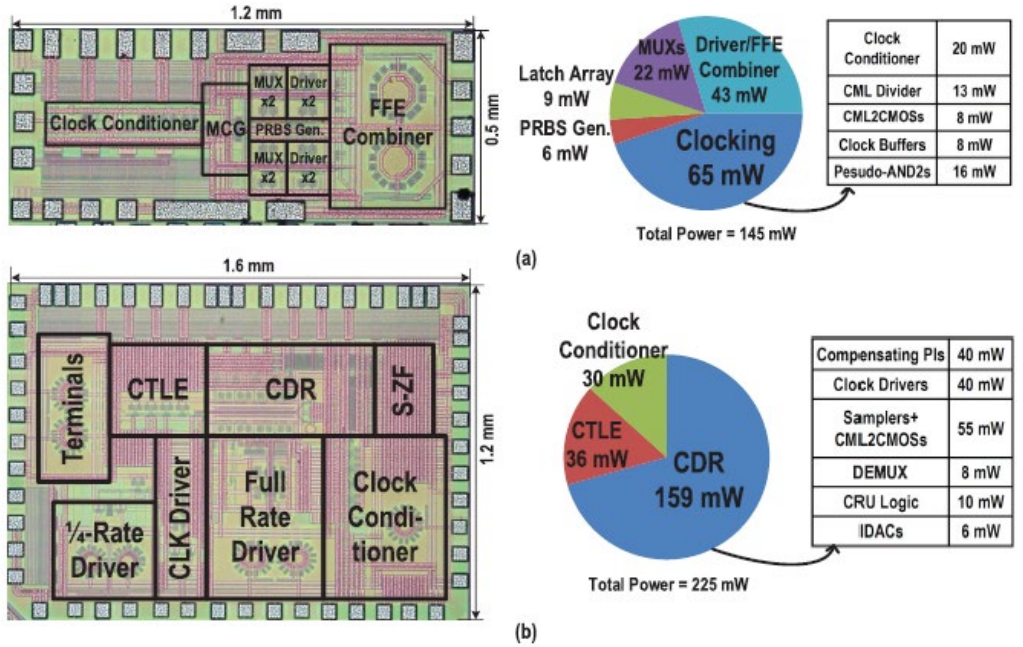


Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

Zheng, 2973.

1456

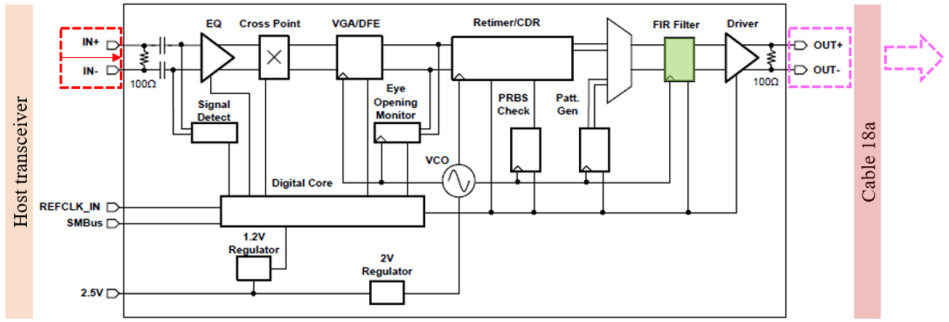
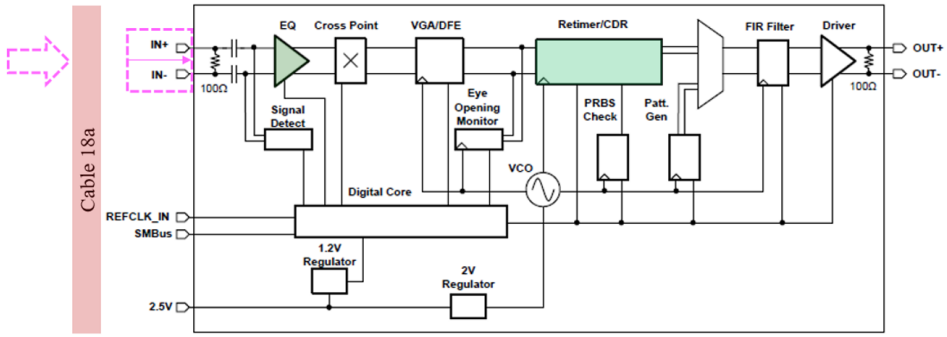
Claim 11	Disclosure
1457	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The active cable 3 receives the serial bit stream transmitted from the upstream device 2 in the upstream facing RX 30a in one retimer 30. The retimer 30 performs a waveform shaping on the serial bit stream received by the upstream facing RX 30a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 30b to the other retime 31. The retimer 31 receives the serial bit stream transmitted from the retimer 30 in the upstream facing RX 31a. The retimer 31 performs a waveform shaping on the serial bit stream received by the upstream facing RX 31a and retimes the serial bit stream after being waveform-shaped and transmits the serial bit stream from the downstream facing TX 31b to the downstream device 4." Kobayashi, 3:5-19. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a

Claim 11	Disclosure
	TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver." Kobayashi, 4:4-9. "The exchange module 202 is configured to communicate capabilities information with a remote link partner including host loss information, RS-FEC ability information, and RS-FEC request information. The exchange module 202 may transmit local host loss information to the link partner and receive remote host loss information from the link partner. In some embodiments, the exchange module 202 may be included in or otherwise established by the autonegotiation module 104 and thus the capabilities information may be exchanged during execution of an autonegotiation protocol. In some embodiments, the exchange module 202 may be included in or otherwise established by the PMD/PMA module 106 and thus the capabilities information may be exchanged during a startup protocol performed by the PMD/PMA module 106." Lusted, paragraph 23. "The autonegotiation module 104 is configured to auto-negotiate line transmission speed, mode of operation, and other communication parameters with a link partner when the communication lane 102 is brought up. Additionally, in some embodiments the autonegotiation module 104 may be configured to exchange capabilities information with the link partner including host insertion loss information, RS-FEC ability information, and RS-FEC request information. The autonegotiation module 104 may encode the capabilities information, for example, in an autonegotiation base page or in a next page. Although illustrated in FIG. 1 as coupled between the PMD/PMA 106 and the communication lanes 102, in other embodiments the autonegotiation module 104 may be included at different positions in the Ethernet port logic 100. For example, in some embodiments the autonegotiation module 104 may be coupled between the PMD/PMA 106 and the FEC 108." Lusted, paragraph 17. "In view of the above transmission characteristics of twisted-pair, current Ethernet chips are usually equipped with preemphasis technique of a transmitting end and equalization technique of a receiving end

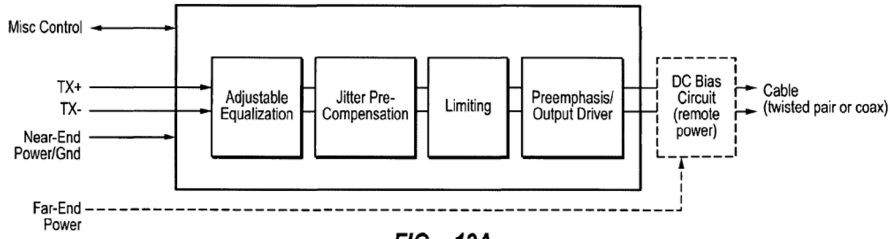
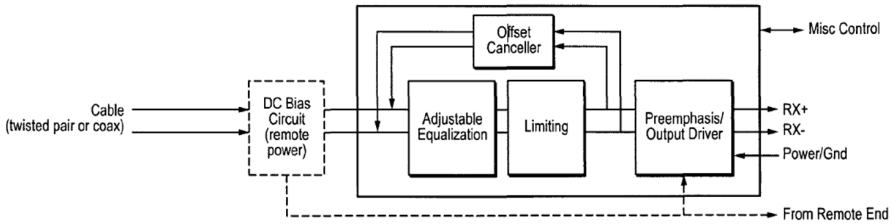
1458

Claim 11	Disclosure
	to compensate the difference in transmission attenuation of physical layer signals on different frequencies on the same medium during Ethernet baseband transmission. Main functions of the preemphasis technique of the transmitting end and the equalization technique of the receiving end include: amplifying signals on different frequencies differently at the transmitting end to compensate different transmission attenuation in the twisted pair; applying the equalization technique to different frequency components in signals received at the receiving end to compensate the attenuation." Yu, paragraph 34. "An embodiment of the present invention provides an equalizer setting apparatus, which uses a filter circuit 101, an analog-to-digital converter 102, a digital memory 103, a digital comparator 104 and a first digital accumulator 105, the filter circuit 101, the analog-to-digital converter 102, the digital memory 103, the digital comparator 104 and the first digital accumulator 105 are connected in sequence, the output end of the equalizer 201 is connected with the input end of the filter circuit 101, the output end of the first digital accumulator 105 is connected with the set value input end of the equalizer 201, the output end of the analog-to-digital converter 102 is connected with the input end of the digital comparator 104, so that the problems that the setting value of the equalizer is accurately set in the prior art, which is time-consuming and labor-consuming are solved, the method has the advantages of achieving accurate setting of the set value of the equalizer in a single time, fast, simple, convenient, accurate and low-cost manner and saving labor force to a certain extent." Yin, paragraph 26 (English translation). *DS125DF1610 Retimer Publication* The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal

1459

Claim 11	Disclosure
	received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."  Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram

1460

Claim 11	Disclosure
	U.S. Patent No. 7,762,727 to Aronson Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id.* at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id.* at 14:16-15:24.  FIG. 12A  FIG. 12B Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side]

1461

Claim 11	Disclosure
	PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23. Aronson discloses the host-side equalization and pre-equalization use fixed cable-independent parameters to address host-side losses. *See, e.g.*, Aronson at 12:10-17, 14:16-23, 15:17-24. "Fixed" parameters reduced complexity and manufacturing costs. See, e.g., U.S. Patent No. 7,239,665 at 1:7-18 (discussing slow convergence with adaptive equalization); U.S. Patent App. Pub. No. 2011/0228821 at [0009]-[0011] (discussing tradeoffs between fixed and adaptive equalization); U.S. Patent No. 6,621,862 at 1:52-55 (large gate count for adaptive equalization can cause "high chip costs"); U.S. Patent App. Pub. No. 2006/0045176 at [0034] ("The adaptive equalizer can be trapped in a failure state (i.e., incorrect state) and not able to recover from the failure state."); U.S. Patent App. Pub. No. 2017/0187463 at [0058] ("[E]qualization techniques… come with different engineering tradeoffs, including trade-offs in power consumption, performance, and/or cost."). **5. The PCIe 3.0 Equalization Adaptation Protocol** The PCIe 3.0 Base Specification [10] requires that a compliant Tx be able to produce any of the output equalization (TxEQ) levels shown in Figure 21, whenever requested to do so by the receiving port, during link equalization training. The same specification describes a protocol of TxEQ adaptation, where the host (upstream agent) or the end device (downstream agent) can each request that any of the TxEQ values in Figure 21 be presented to them by the other agent on the link. The intent of adaptation is to allow both agents to adjust the link partner's TxEQ to an optimal value for their receiver, and for the specific channel and operating conditions at hand. Samaan, p. 24.

1462

L. DEPENDENT CLAIM 12

Claim 12	Disclosure
12[a] The cable manufacturing method of claim 11, wherein the first and second transceivers are each configured to employ cable-dependent equalization parameters for generating the first and second electrical transit signals and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

1463

Claim 12	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1464

Claim 12	Disclosure
	Table 10- Control field dependency on Explicit Control bit Control field | Explicit Control value Tx Adaptive Input Eq Enable | 1 Tx Adaptive Input Eq Recall | 1 or 0 Tx Input Eq control | 1 Tx CDR control | 1 Rx CDR control | 1 Rx Output Eq control, pre-cursor | 1 Rx Output Eq control, post-cursor | 1 Rx Output Amplitude control | 1 CMIS, 38. 1.5.4.1 Control Set Usage There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

1465

Claim 12	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1466

Claim 12	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1467

Claim 12	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1468

Claim 12	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1469

Claim 12	Disclosure

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

CMIS, 71.

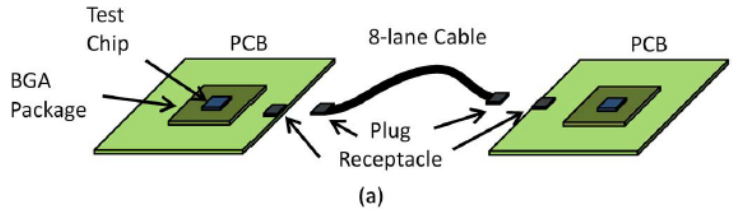
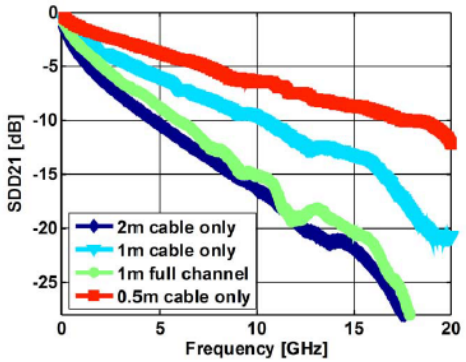
1470

Claim 12	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1471

Claim 12	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

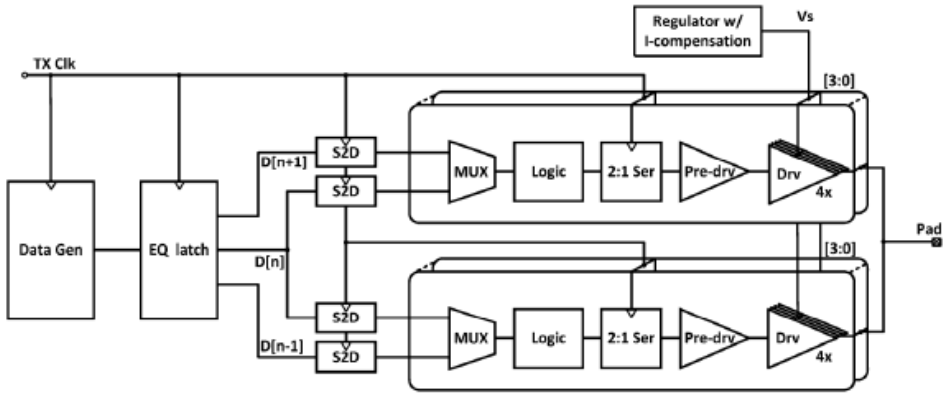
1472

Claim 12	Disclosure
1473	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

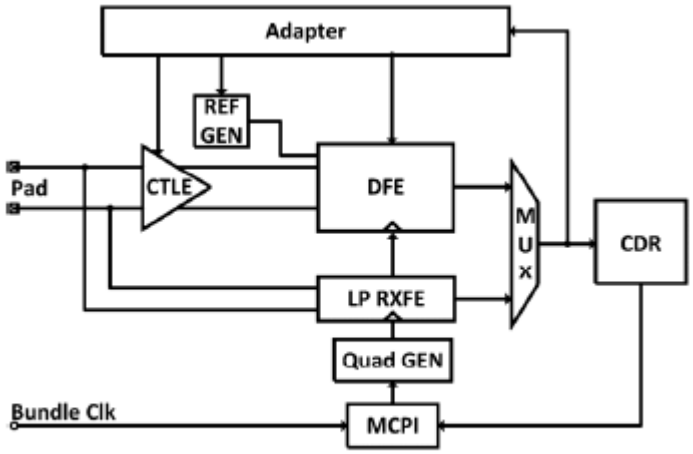
Claim 12	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 12	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1475

Claim 12	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1476

Claim 12	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

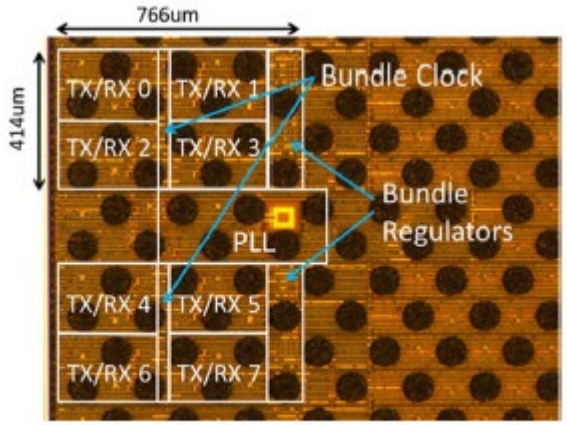
1477

Claim 12	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1478

Claim 12	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1479

Claim 12	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

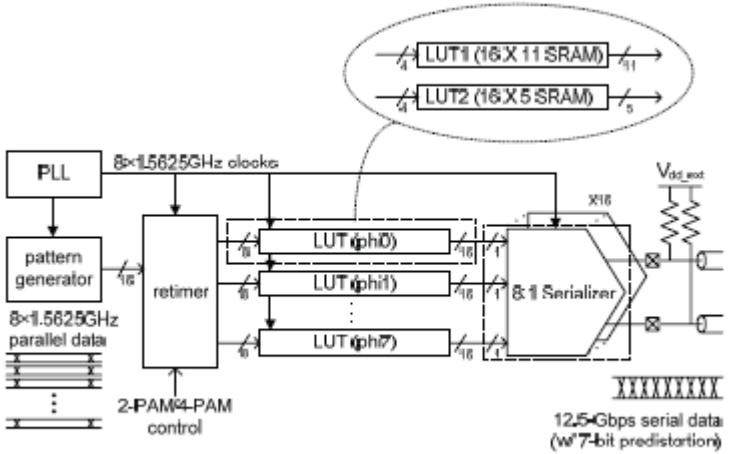
1480

Claim 12	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1481

Claim 12	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1482

Claim 12	Disclosure
1483	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 12	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1484

Claim 12	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1485

Claim 12	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1486

Claim 12	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1487

Claim 12	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 12	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1489

Claim 12	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1490

Claim 12	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1491

Claim 12	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

1492

Claim 12	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1493

Claim 12	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1494

Claim 12	Disclosure

IV. Channel Equalization

The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18.

Zheng, 2970.

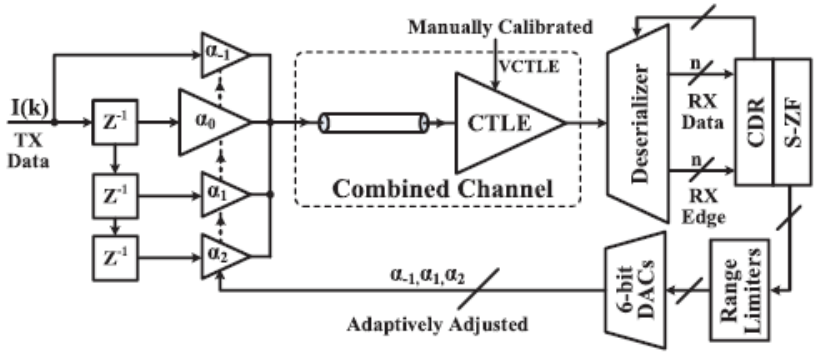


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1495

Claim 12	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

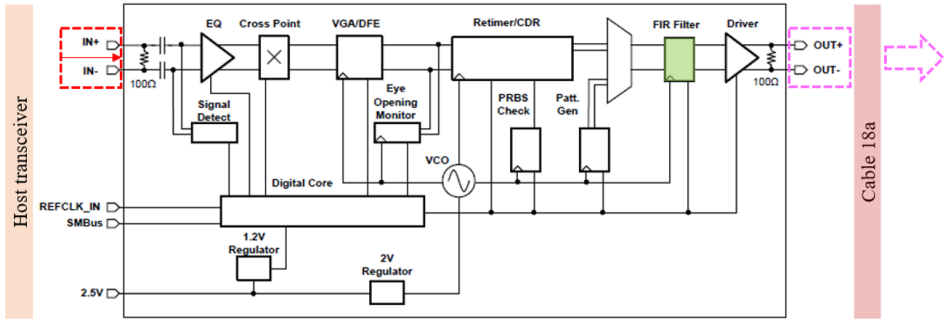
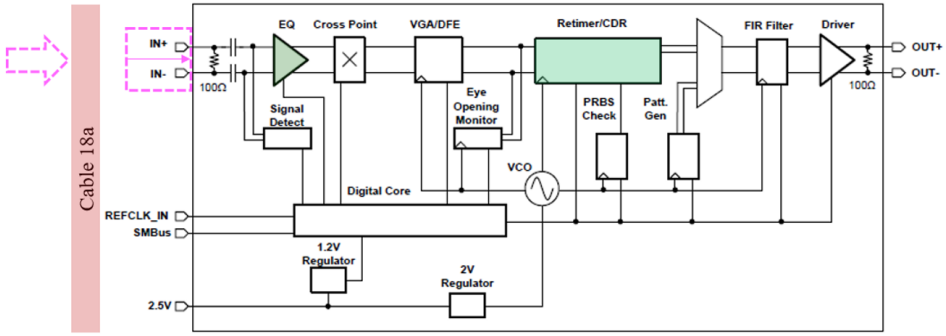
1496

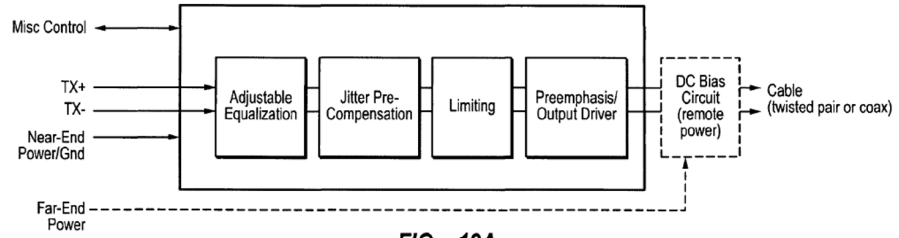
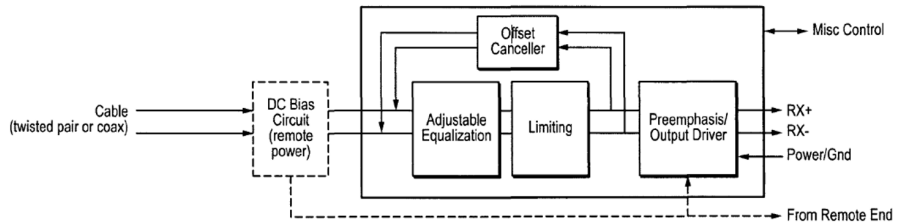
Claim 12	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50 itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK"

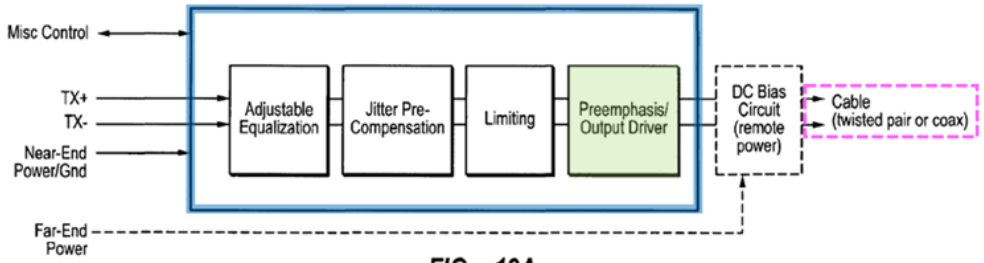
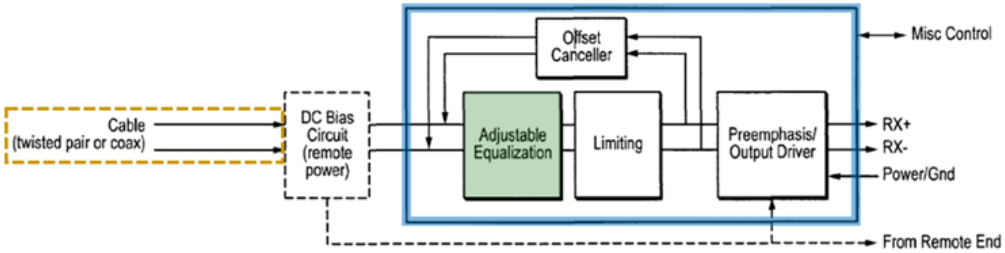
1497

Claim 12	Disclosure
	Kobayashi, 4:4-22. **"**The link activation module 204 is configured to determine a total insertion loss for the connection with the remote host based on the local host loss, the remote host loss, and cable loss. The link activation module 204 compares the total insertion loss to a FEC limit and to a specification limit. The link activation module 204 may bring up the communication link with the link partner without FEC enabled if the total insertion loss is below the specification limit and below the FEC limit. Additionally, the link activation module 204 may bring up the communication link with the link partner with FEC enabled if the total insertion loss is below the specification limit but not below the FEC limit. Further, the link activation module 204 may not bring up the communication link with the link partner if the total insertion loss is not below the specification limit (e.g., the link activation module 204 may keep the communication link down or bring the communication link down)." Lusted, paragraph 24. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. <u>*DS125DF1610 Retimer Publication*</u> The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal

1498

Claim 12	Disclosure
1499	for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."  Figure 2. DS125DF1610 Simplified Data Path Diagram  Figure 2. DS125DF1610 Simplified Data Path Diagram

Claim 12	Disclosure
1500	*U.S. Patent No. 7,762,727 to Aronson* Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id.* at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id.* at 14:16-15:24.  FIG. 12A  FIG. 12B Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it

Claim 12	Disclosure
	receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id*. at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id*. at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23. Using "cable dependent" parameters for equalization would have been the conventional and obvious way to address losses over the cable. *See, e.g.*, Aronson at 14:33-43 (pre-emphasis/output driver performs pre-equalization to overcome "loss on the cable" using cable-dependent parameters that "match the particular length [and/or] the characteristics of the copper cabling"), Fig. 12A (below).  *FIG. 12A* Using "cable dependent" parameters for equalization would have been the conventional and obvious way to address losses over the cable. Aronson at 15:1-8 (equalization "compensate[s] for the cable" losses using cable dependent parameters based on "cable length and characteristics"), Fig. 12B (below).  *FIG. 12B*

1501

Claim 12	Disclosure
12[b] for clock and data recovery on the second and first electrical transit signals.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1.

1502

Claim 12	Disclosure
	1.4.4 MgmtInit State The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

1503

Claim 12	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1504

Claim 12	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

1505

Claim 12	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1506

Claim 12	Disclosure
	1.7 Module Memory Map This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module. CMIS, 44. **1.5.5 Signal Integrity Control Field Background** The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated. The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11). Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable Control / Tx Adaptive Input Eq Enable value to use this control Tx Input Eq Adaptation Freeze — 1 Tx Input Eq Adaptation Store — 1 Tx Adaptive Input Eq Recall — 1 Tx Input Eq control — 0 The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module. CMIS, 40.

1507

Claim 12	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1508

Claim 12	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1509

Claim 12	Disclosure

1510

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

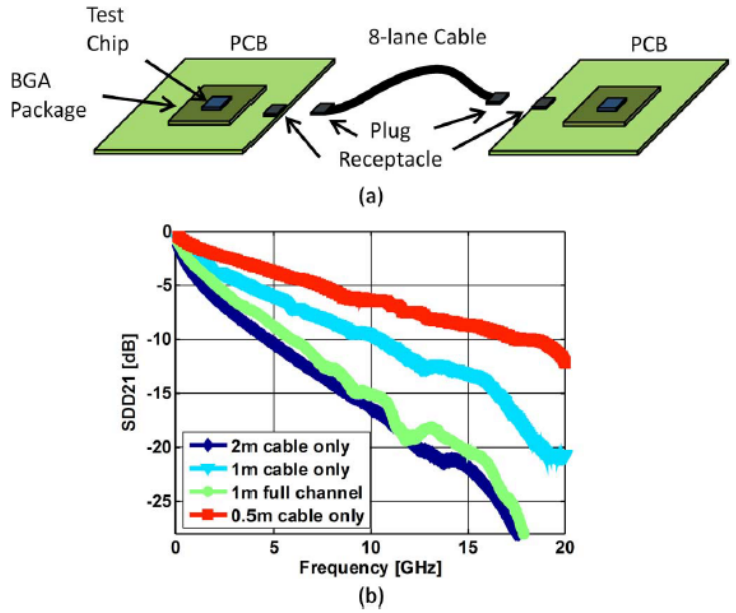
CMIS, 71.

Claim 12	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1511

Claim 12	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

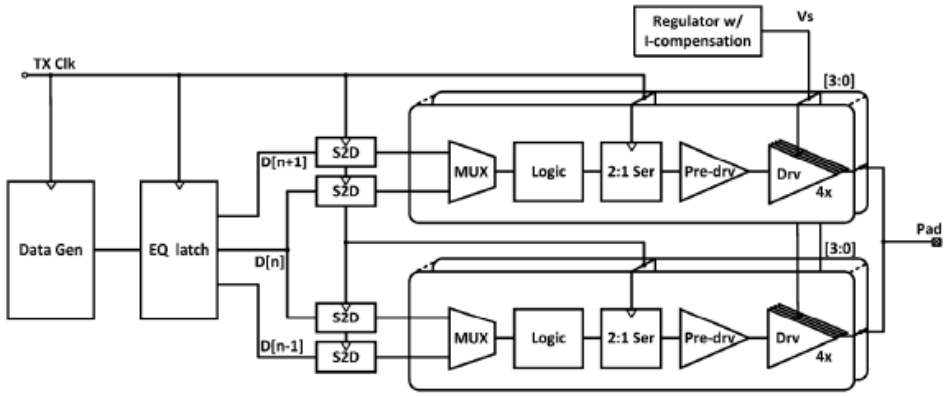
1512

Claim 12	Disclosure
1513	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

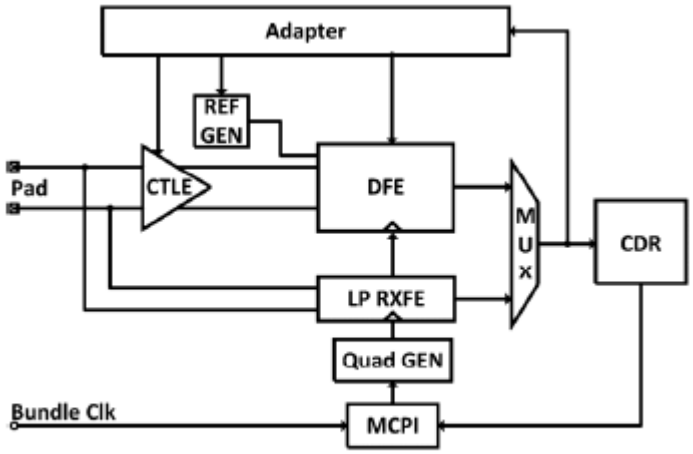
Claim 12	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 12	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1515

Claim 12	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1516

Claim 12	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

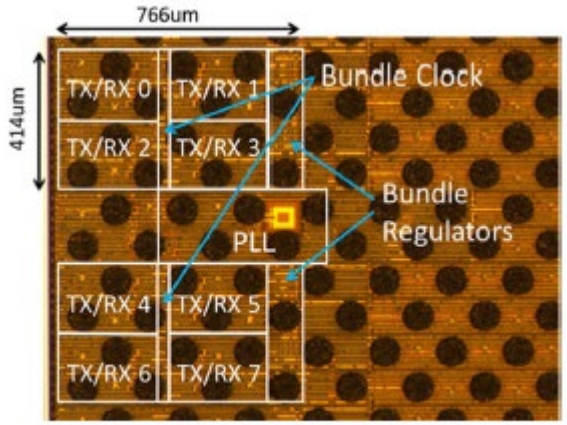
1517

Claim 12	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1518

Claim 12	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1519

Claim 12	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

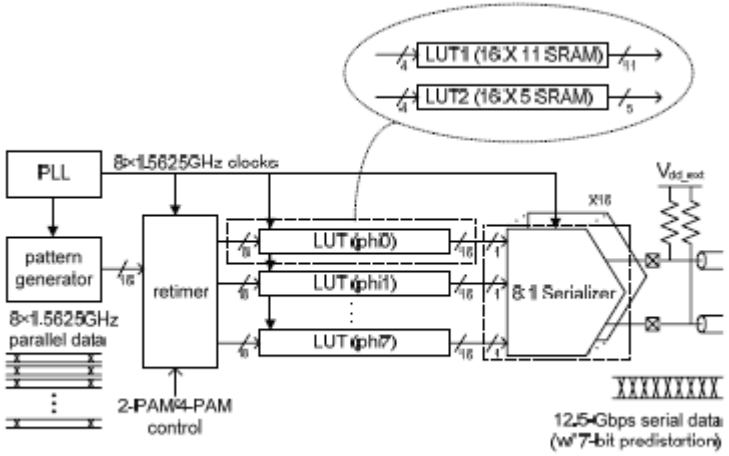
1520

Claim 12	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1521

Claim 12	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1522

Claim 12	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1523

Claim 12	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1524

Claim 12	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1525

Claim 12	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1526

Claim 12	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

Claim 12	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 12	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1529

Claim 12	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1530

Claim 12	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1531

Claim 12	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

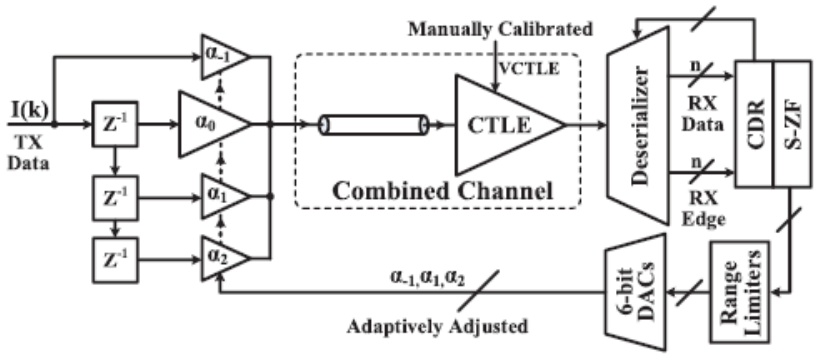
1532

Claim 12	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1533

Claim 12	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1534

Claim 12	Disclosure
1535	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

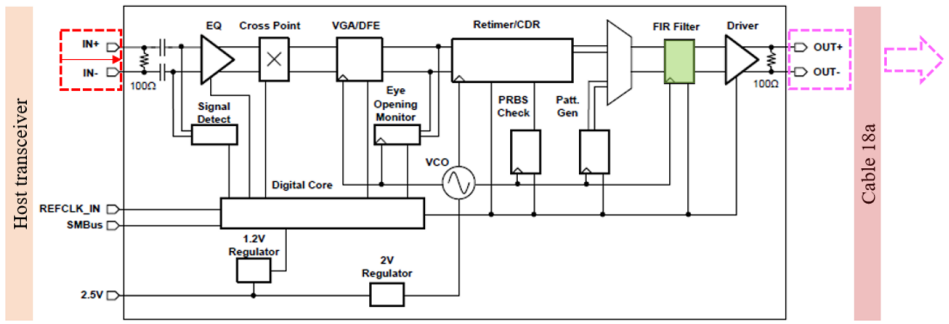
Claim 12	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1536

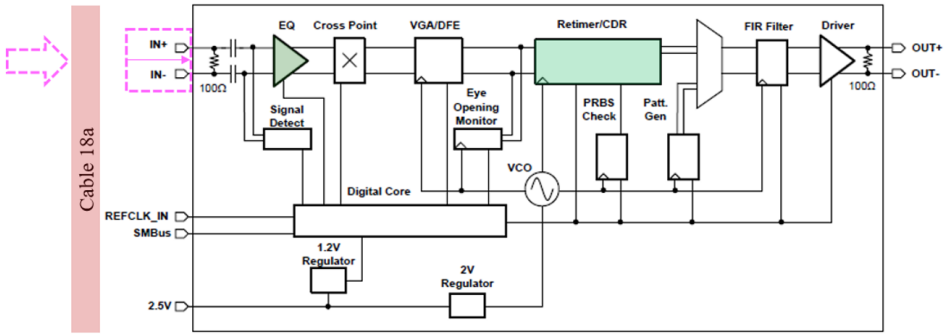
Claim 12	Disclosure
1537	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "As shown in FIG. 2, the retimer 50 includes an upstream facing RX 51, a downstream facing TX 52, a TX serial bit rate generator 53, a FIFO 54, a link training sequence pattern generator 55, and a selection circuit 56. The upstream facing RX 51 includes an equalizer, a CDR, and a de-serializer. The downstream facing TX 52 includes a serializer and a driver. The upstream facing RX 51 generates a parallel word (10 bits, for example), a recovered parallel clock, and an incoming link training sequence marker (IN_LT_SEQ_MARK) based on the serial bit stream transmitted from the upstream TX and outputs them. The TX serial bit rate generator 53 generates a serial bit clock and a parallel word clock based on a local timing reference 60 (a type of reference clock) generated in the retimer 50 and outputs them. The TX serial bit rate generator 53 also generates a training sequence marker (SELF_LT_SEQ_MARK) of the retimer 50

Claim 12	Disclosure
	itself and outputs it. The SELF_LT_SEQ_MARK is synchronized with the IN_LT_SEQ_MARK" Kobayashi, 4:4-22. "The link activation module 204 is configured to determine a total insertion loss for the connection with the remote host based on the local host loss, the remote host loss, and cable loss. The link activation module 204 compares the total insertion loss to a FEC limit and to a specification limit. The link activation module 204 may bring up the communication link with the link partner without FEC enabled if the total insertion loss is below the specification limit and below the FEC limit. Additionally, the link activation module 204 may bring up the communication link with the link partner with FEC enabled if the total insertion loss is below the specification limit but not below the FEC limit. Further, the link activation module 204 may not bring up the communication link with the link partner if the total insertion loss is not below the specification limit (e.g., the link activation module 204 may keep the communication link down or bring the communication link down)." Lusted, paragraph 24. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55. <u>*DS125DF1610 Retimer Publication*</u> The transceivers of the DS125DF1610 each employ cable-dependent equalization parameters for at least one of the remodulation of the inbound data stream for transit (by the DS125DF1610's FIR filter) and the

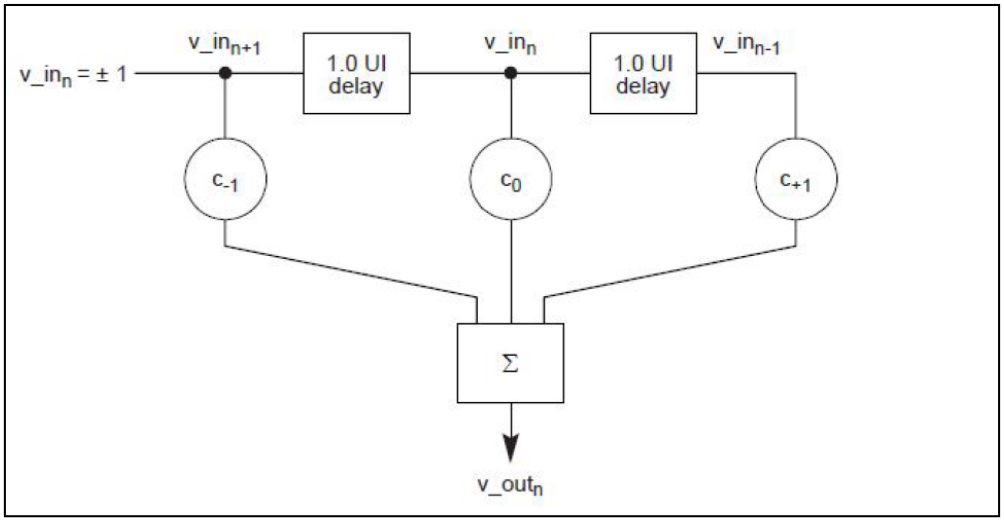
1538

Claim 12	Disclosure
	CDR performed on the electrical input signal (by the DS125DF1610's EQ). DS125DF1610's Figure 2 is annotated twice below – once to show the retimer processing the transit signal for transmission over the cable (top) and the other showing the retimer processing the transit signal received from the other side of the cable (bottom). In the top use case, FIR Filter equalizes the remodulated signal provided as the transit signal transmitted over cable 18a. *See, e.g.*, DS125DF1610 at §§3, 7.3.8. In the bottom use case, EQ equalizes the transit signal received over cable 18a for Retimer/CDR circuit. *See id.* at DS125DF1610, §7.3.3. Equalizing the transit signal before performing CDR and pre-equalizing the transit signal before transmission over the cable meets "employing … equalization parameters for … the remodulation of the inbound data stream for transit, and the [CDR] performed on the electrical transit signal."  Figure 2. DS125DF1610 Simplified Data Path Diagram

1539

Claim 12	Disclosure
	 Figure 2. DS125DF1610 Simplified Data Path Diagram

1540

Claim 12	Disclosure
1541	 **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.

M. DEPENDENT CLAIM 13

Claim 13	Disclosure
13. The cable manufacturing method of claim 12, wherein the first and second transceivers are each configured to adapt the cable-dependent equalization parameters during operation.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

1542

Claim 13	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1543

Claim 13	Disclosure

Table 10- Control field dependency on Explicit Control bit

Control field	Explicit Control value
Tx Adaptive Input Eq Enable	1
Tx Adaptive Input Eq Recall	1 or 0
Tx Input Eq control	1
Tx CDR control	1
Rx CDR control	1
Rx Output Eq control, pre-cursor	1
Rx Output Eq control, post-cursor	1
Rx Output Amplitude control	1

CMIS, 38.

1.5.4.1 Control Set Usage

There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set.

When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset.

CMIS, 38.

1544

Claim 13	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1545

Claim 13	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1546

Claim 13	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1547

Claim 13	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: **Table 14- Rx Output Amplitude Codes** Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1548

Claim 13	Disclosure

1549

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

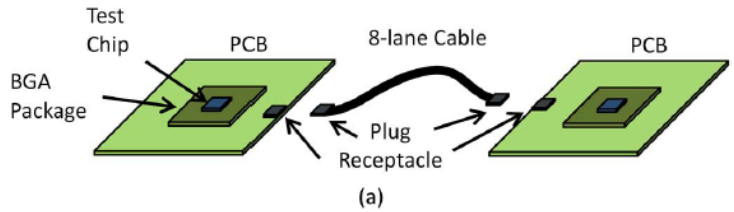
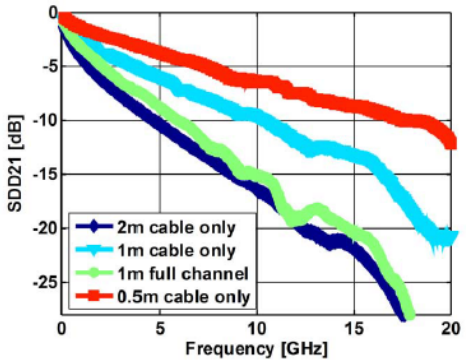
CMIS, 71.

Claim 13	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1550

Claim 13	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1551

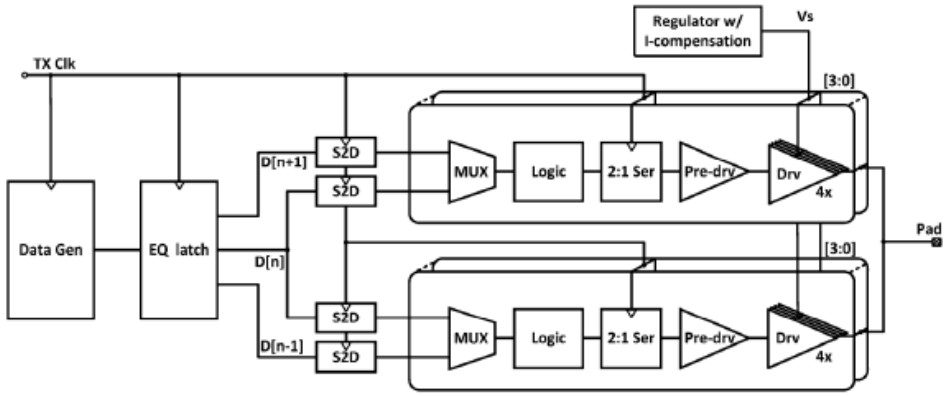
Claim 13	Disclosure
	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

1552

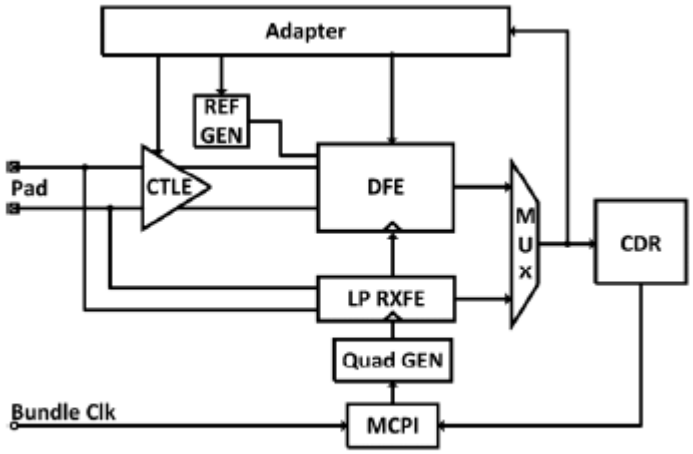
Claim 13	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

Claim 13	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1554

Claim 13	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1555

Claim 13	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

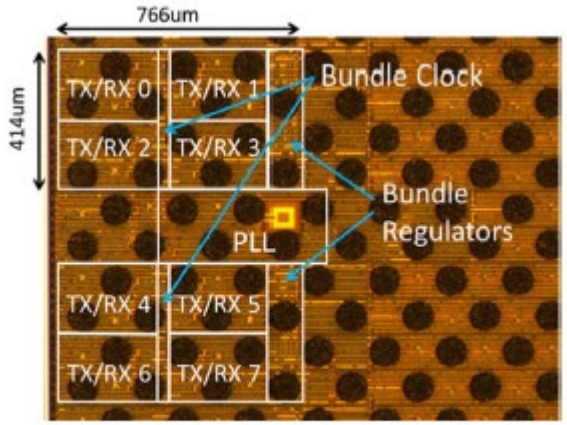
1556

Claim 13	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1557

Claim 13	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1558

Claim 13	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

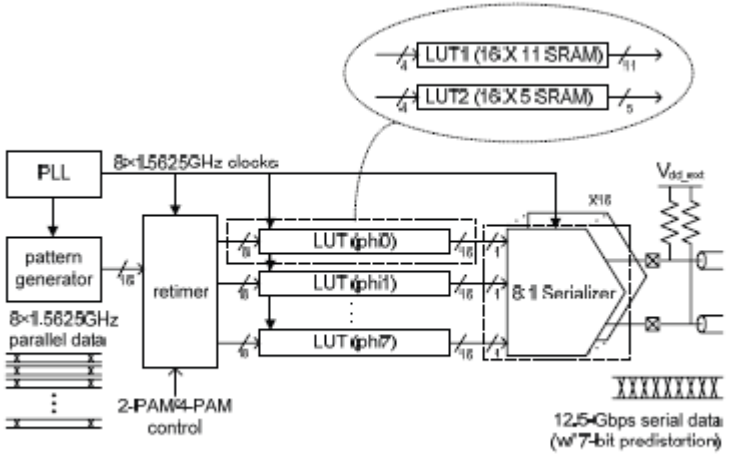
1559

Claim 13	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1560

Claim 13	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1561

Claim 13	Disclosure
1562	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 13	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1563

Claim 13	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1564

Claim 13	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1565

Claim 13	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1566

Claim 13	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

Claim 13	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1568

Claim 13	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1569

Claim 13	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1570

Claim 13	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

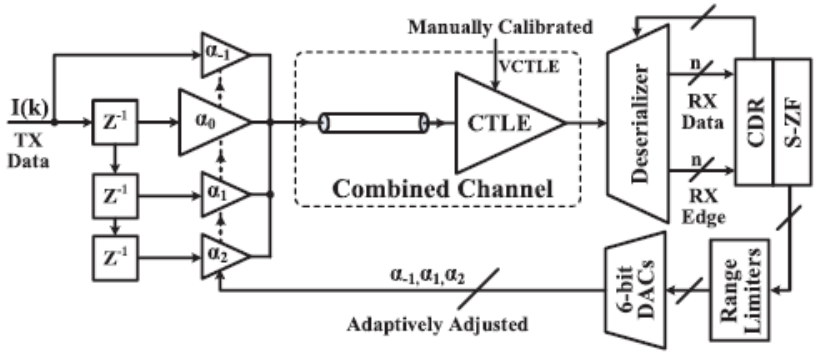
1571

Claim 13	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1572

Claim 13	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1573

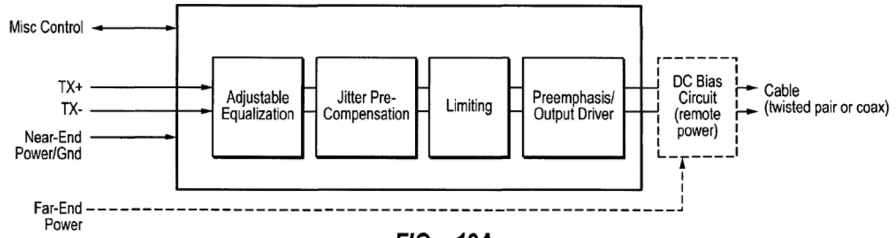
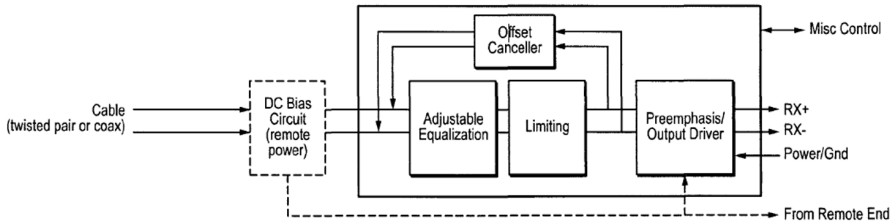
Claim 13	Disclosure
	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1574

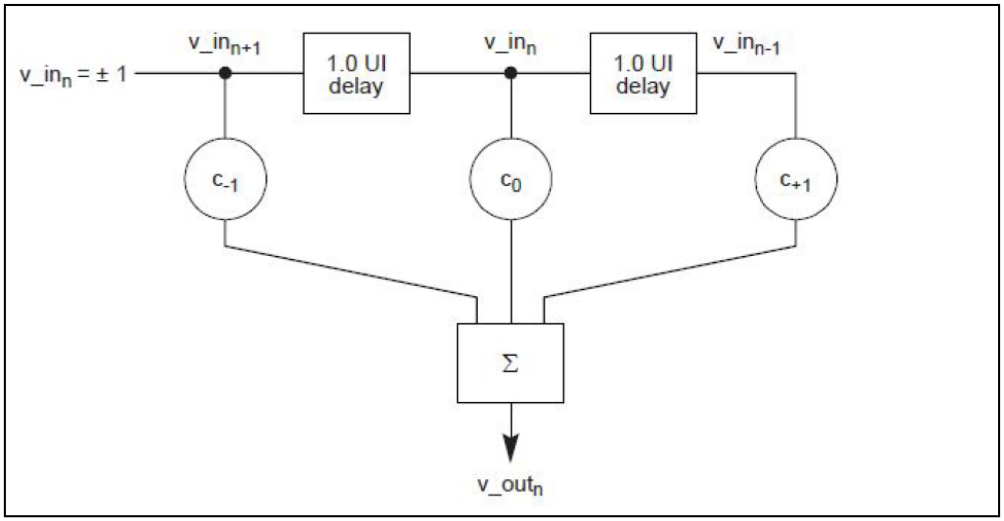
Claim 13	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1575

Claim 13	Disclosure
1576	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55.

Claim 13	Disclosure
	U.S. Patent No. 7,762,727 to Aronson Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id.* at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id.* at 14:16-15:24.  FIG. 12A  FIG. 12B Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id.* at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id.* at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23.

1577

Claim 13	Disclosure
	 **Figure 22 The PCIe 3.0 Tx Equalization FIR filter representation. (Note: The 10G-KR Filter is similar in definition).** This PCIe 3.0 TxEQ adaptation [10] protocol has 4 phases, defined as follows:- Samaan, p. 25.

1578

N. DEPENDENT CLAIM 14

Claim 14	Disclosure
14[a] The cable manufacturing method of claim 12, wherein the first and second transceivers are each configured to use preset cable-dependent equalization parameters during operation, and	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1. **1.4.4 MgmtInit State** The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp.

1579

Claim 14	Disclosure
	CMIS, 22. **1.5.4 Control Set Overview** A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1580

Claim 14	Disclosure
1581	Table 10- Control field dependency on Explicit Control bit Control field | Explicit Control value Tx Adaptive Input Eq Enable | 1 Tx Adaptive Input Eq Recall | 1 or 0 Tx Input Eq control | 1 Tx CDR control | 1 Rx CDR control | 1 Rx Output Eq control, pre-cursor | 1 Rx Output Eq control, post-cursor | 1 Rx Output Amplitude control | 1 CMIS, 38. 1.5.4.1 Control Set Usage There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

Claim 14	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1582

Claim 14	Disclosure

1.7 Module Memory Map

This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module.

CMIS, 44.

1.5.5 Signal Integrity Control Field Background

The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated.

The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11).

Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable

Control	Tx Adaptive Input Eq Enable value to use this control
Tx Input Eq Adaptation Freeze	1
Tx Input Eq Adaptation Store	1
Tx Adaptive Input Eq Recall	1
Tx Input Eq control	0

The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module.

CMIS, 40.

1583

Claim 14	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1584

Claim 14	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1585

Claim 14	Disclosure

1586

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

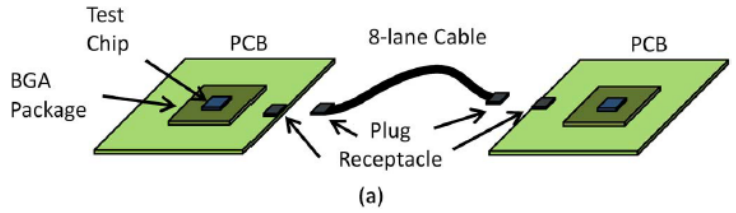
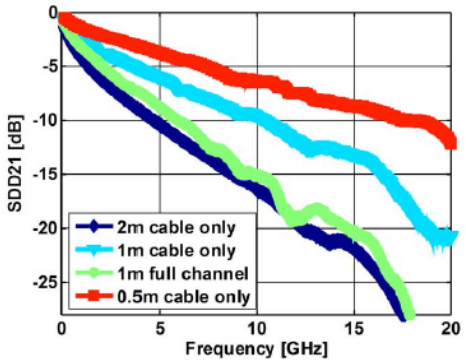
CMIS, 71.

Claim 14	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1587

Claim 14	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1588

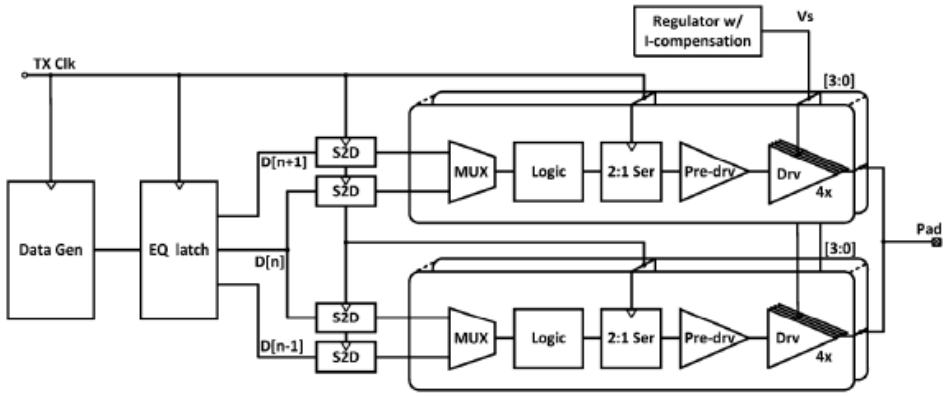
Claim 14	Disclosure
1589	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 14	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

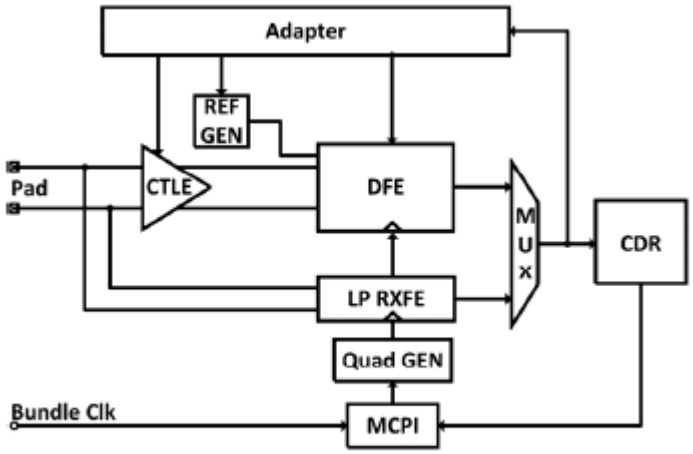
1590

Claim 14	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1591

Claim 14	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

1592

Claim 14	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

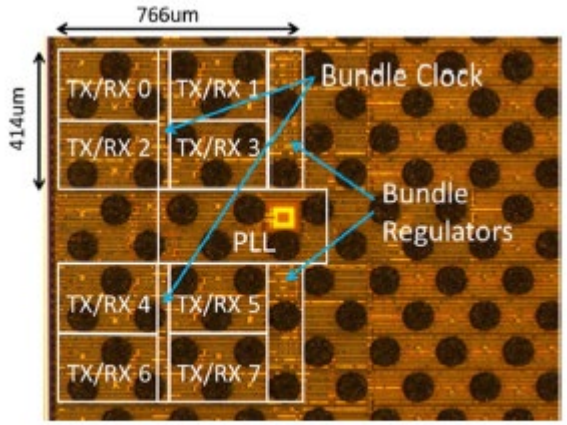
1593

Claim 14	Disclosure
	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

1594

Claim 14	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1595

Claim 14	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

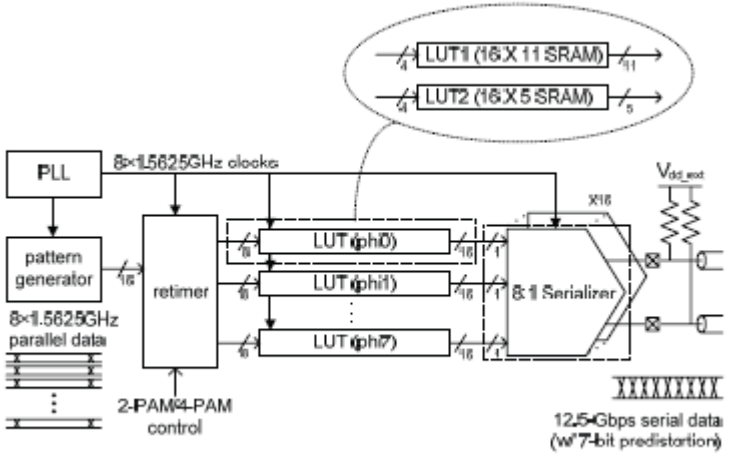
1596

Claim 14	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13μm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

1597

Claim 14	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1598

Claim 14	Disclosure
1599	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 14	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1600

Claim 14	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1601

Claim 14	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1602

Claim 14	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1603

Claim 14	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

1604

Claim 14	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1605

Claim 14	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1606

Claim 14	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1607

Claim 14	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

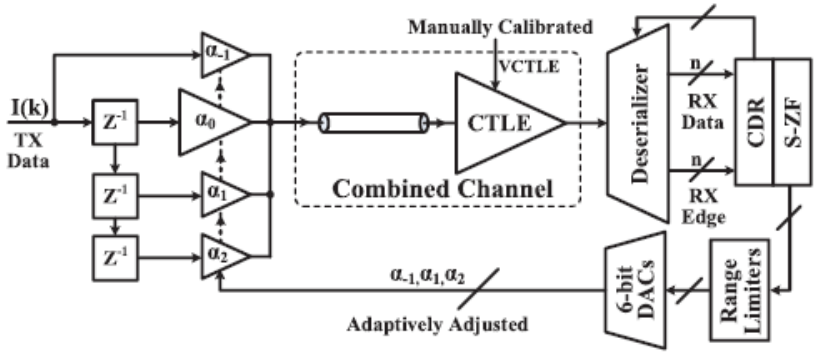
1608

Claim 14	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1609

Claim 14	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

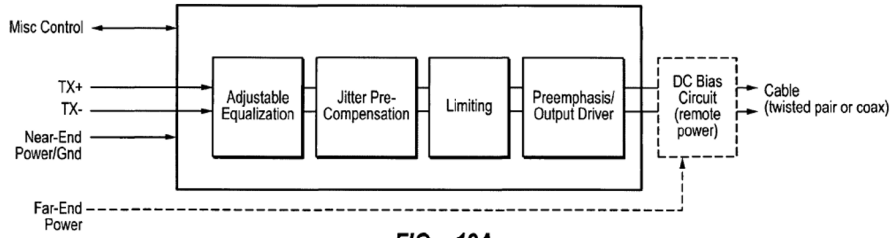
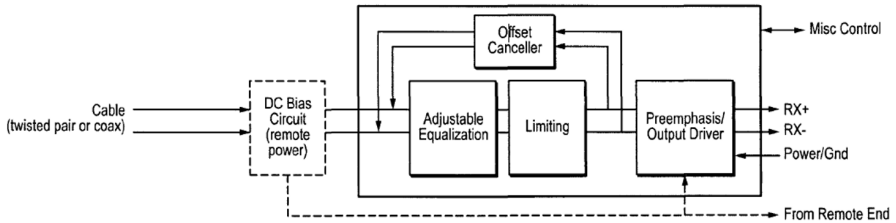
1610

Claim 14	Disclosure
1611	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 14	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1612

Claim 14	Disclosure
1613	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55.

Claim 14	Disclosure
1614	*U.S. Patent No. 7,762,727 to Aronson* Aronson discloses an active cable for high-speed transmissions between network nodes. See, e.g., Aronson at Abstract, 1:12-21. While Aronson's title recites an optical cable, Aronson also discloses copper twisted-pair or coaxial cables like Lugthart '414. *See, e.g., id.* at 1:29-2:23, 7:17-20. Figures 12A-12B (below) illustrate features "incorporated into the ICs in a copper active cable design." *See, e.g., id.* at 14:16-15:24.  FIG. 12A  FIG. 12B Figure 12A's circuit processes TX signals from the host and transmits signals over the cable, where equalization "compensat[es] for high frequency loss in the host board traces," and pre-equalization overcomes loss on the cable and "match[es]" loss from "the particular length [and] characteristics of the copper cabling." *See* Aronson at 14:16-43. Figure 12B's circuit handles dataflow in the other direction: it receives RX signals from the cable, performs equalization "to compensate for" cable losses, and—before transmitting to the host—performs pre-emphasis to "overcome high frequency losses on long [host-side] PCB traces." *See, e.g., id.* at 15:1-6, 15:17-24. Signal pre-emphasis and/or equalization can be "fixed, adjustable, or adaptable." *See, e.g., id.* at Abstract, 14:19-23, 14:33-34, 15:1-8, 15:17-23.

Claim 14	Disclosure
14[b] wherein the method further comprises: testing an assembled cable to determine the cable-dependent equalization parameters.	For example, and without limitation, at least under Credo's apparent infringement theory, the following obviousness references disclose this limitation based on at least the following exemplary citations: Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and transceivers. CMIS, 1.

1615

Claim 14	Disclosure
	1.4.4 MgmtInit State The MgmtInit state is a module control state that is entered any time the module is brought out of the Reset state, (either hardware or software reset). The MgmtInit state is applicable to both active modules and passive copper cable assemblies. During this state, the module shall configure the memory map and initialize the management interface for access by the host. The module may perform limited power-up of the high-speed data path circuitry, however the module shall remain in Low Power Mode throughout this state. For active modules, all Data Path States shall remain in DataPathDeactivated throughout MgmtInit. The module may ignore all TWI transactions while in the MgmtInit state. Interrupt flag conformance in the MgmtInit state is defined in Section 1.6.1. All disallowed interrupt flags shall not be set during MgmtInit. Before the module exits the MgmtInit state, all memory map register locations shall be set to their power-on defaults. In Hardware Init mode, power-up defaults may include DataPathPwrUp bits. The module shall have completed MgmtInit within the module form-factor dependent management interface initialization time, defined as the time from power on (defined as the instant when supply voltages reach and remain at or above the minimum level specified in the form factor-dependent specification), hot plug, or the rising edge of the Reset signal until the module has configured the memory map to default conditions and activated the management interface. Upon completion of MgmtInit, the next state depends on the Init mode of the module. If the module is booting in Software Init mode, the next state is ModuleLowPwr. If the module is booting in Hardware Init mode, the next state is ModulePwrUp. CMIS, 22.

1616

Claim 14	Disclosure
	1.5.4 Control Set Overview A Control Set is a group of registers that are used to provide initialization settings for data paths. Each control set includes fields to select the Application, fields for signal integrity control, and fields to apply these Application and signal integrity settings. Within a Control Set, the host uses the Application Select Control registers defined in Table 56 or Table 60 to select the desired Application(s) and configure the data path(s) within the module. The Application Select Control registers include an ApSel code, a Data Path Code and an Explicit Control bit. Each supported Application advertised in Table 26 and Table 48 is identified by an ApSel code. The host uses the ApSel code to assign that Application to one or more specific lanes using the Application Select Control registers. Where an Application requires multiple lanes, the host shall write the same ApSel code into each lane of that Application. The Data Path Code in the Application Select Control registers is used to specify the lanes used for an instance of an Application. Each data path is identified by the lowest numbered lane in that data path. The host shall assign lane combinations that are in accordance with the lane assignment advertised by the module for that application. For a multi-lane data path the host shall write the same data path code in all lanes of that data path. The Explicit Control bit allows the host to provide either its own signal integrity values or use the Application-defined signal integrity values associated with the ApSel code. These Application-defined settings shall provide host electrical and media interfaces that are compliant with the standard associated with the selected Application. To use Explicit Control for a given lane (by setting the control bit to '1'), the host shall provide values for all signal integrity controls for that lane in the Control Set. The value of the Explcit Control bit determines which signal integrity control fields within the Control Set are used in an Apply operation. When the Explicit Control bit for a given lane is set to the value listed in Table 10, the associated control field value for that lane is used in an Apply operation. The control field value for that lane is ignored for all other values of the Explicit Control bit. CMIS, 37-38.

1617

Claim 14	Disclosure
1618	Table 10- Control field dependency on Explicit Control bit Control field | Explicit Control value Tx Adaptive Input Eq Enable | 1 Tx Adaptive Input Eq Recall | 1 or 0 Tx Input Eq control | 1 Tx CDR control | 1 Rx CDR control | 1 Rx Output Eq control, pre-cursor | 1 Rx Output Eq control, post-cursor | 1 Rx Output Amplitude control | 1 CMIS, 38. **1.5.4.1 Control Set Usage** There are two types of control sets. The Active Control Set reports the current settings that are used by the module to control its hardware. The Staged Control Sets are used by the host to identify new settings for future use. Apply-fields are used to copy those settings from the Staged Control Set registers into the Active Control Set registers. This apply mechanism decouples the timing and sequence of host writes to the Staged Control Set from the module actions used to configure the module hardware. Each module shall implement the Active Control Set and at least one host-configurable Staged Control Set. When the host requests power up of a data path, the module shall use the settings in the Active Control Set for data path initialization. Those settings may have been copied from the Staged Control Set by a previous host Apply action, or those settings may have been programmed into the Active Set as power-on defaults. The module shall populate both Staged Control Set 0 and the Active Control Set registers with the module-defined default Application and signal integrity settings at power on or exit from reset. CMIS, 38.

Claim 14	Disclosure
	1.5.5.1 Tx Input Eq Adaptation Store/Recall Methodology Transmit equalizer adaptation can be a time-consuming activity. In some implementations, the available time for a speed change does not include equalizer adaptation time. This specification provides an optional Tx Input Eq Adaptation Store and Recall mechanism, to facilitate storing of adapted equalizer values for recall and use at a later time. Module support of the Tx input Eq Adaptation Store and Recall mechanism is optional and declared in Table 46. This advertisement field identifies the number of Store/Recall buffers implemented in the module. These buffers are independent of Staged Set 0 and 1 and may be numbered independently. The module shall provide sufficient storage in each Store/Recall buffer to store the adapted equalizer value for each lane in the module. This storage is implementation specific and not defined in this specification. The Tx Input Eq Adaptation Store control field is located in Table 54. This field provides two bits per lane and is write only. A read of this register shall return 0. When the host would like to store the most recent adapted Tx input equalizer value, the host shall write the target store location encoding into the Tx Input Eq Adaptation Control, into each lane whose adapted value should be stored. Adaptation shall continue to occur after the store event completes unless the Tx Input Eq Adaptation Freeze bit is set. Host requests to store the Tx Input Eq Adaptation when the Tx Input Eq Adaptation Enable bit for that lane is clear shall be ignored by the module. Tx Input Eq Adaptation Store may occur at any time while a data path is powered and occurs when requested. The Tx Input Eq Adaptation Recall control field is located in Table 57 for Staged Set 0 and Table 61 for Staged Set 1. This field provides two bits per lane and is read-write. The Active Set provides a read-only indication of the current Tx Input Eq Adaptation Recall status for each lane. The host may recall any stored Tx Input Eq Adaptation by programming the applicable lane controls with the store location to be recalled. These values are not recalled until the Apply_DataPathInit or Apply_Immediate bits are 1 for that Staged Set. The Tx Input Eq Adaptation Recall field is used independent of the Explicit Control field settings for that lane. If the Tx input freeze bit is clear, the recalled Tx input eq adaptation shall be used as the starting point for continuous adaptation for the applicable lanes. If the Tx Input Eq Freeze bit is set, the recalled Tx Input Eq Adaptation shall be used as the frozen Tx Input Eq value for the applicable lanes. CMIS, 41-42.

1619

Claim 14	Disclosure
1620	**1.7 Module Memory Map** This subsection defines the Memory Map for a CMIS Module used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all CMIS devices. The interface has been designed largely after the QSFP memory map. The memory map has been changed in order to accommodate 8 electrical lanes and limit the required memory space. The single address approach is used as found in QSFP. Paging is used in order to enable time critical interactions between host and module. CMIS, 44. **1.5.5 Signal Integrity Control Field Background** The signal integrity control fields provide an opportunity for the host to specify signal integrity settings to the module. For each control, the host sets the code for the desired behavior and the device makes a best effort to provide the function indicated. The Tx Input Equalizer has multiple controls associated with it. These controls can be divided into two groups: those that are active when Tx Adaptive Input Eq is enabled and those that are active when Tx Adaptive Input Eq is disabled. Table 11 summarizes which controls are associated with each group. The module shall ignore the value in the applicable control field when the Tx Adaptive Input Eq Enable bit is not set to use that control (see Table 11). Table 11- Tx Input Eq control relationship to Tx Adaptive Input Eq Enable Control / Tx Adaptive Input Eq Enable value to use this control Tx Input Eq Adaptation Freeze — 1 Tx Input Eq Adaptation Store — 1 Tx Adaptive Input Eq Recall — 1 Tx Input Eq control — 0 The Tx Input Equalization Control is a four-bit field per lane as shown in Table 12. This field allows the host to specify fixed Tx input equalization and is ignored by the module if Tx Adaptive Input Eq Enable is set for that lane. Refer to Table 46 to determine if the module supports Fixed Tx Input Equalization Control. Refer to Table 42 for the Maximum Tx equalization supported by the module. The code values and the corresponding input equalization are based on a reference CTLE and may not directly apply to the equalizer implemented in the module. CMIS, 40.

Claim 14	Disclosure

Table 12- Fixed Tx Input Equalization Codes

Code Value	Bit pattern	Input Equalization
0	0000	No Equalization
1	0001	1 dB
2	0010	2 dB
3	0011	3 dB
4	0100	4 dB
5	0101	5 dB
6	0110	6 dB
7	0111	7 dB
8	1000	8 dB
9	1001	9 dB
10	1010	10 dB
11	1011	11 dB
12	1100	12 dB
13-15		Custom

The Rx Output Emphasis Control is a four-bit field per lane. Refer to Table 46 to determine if the module supports Rx Output Emphasis Control. Refer to Table 42 for the maximum Rx output emphasis supported by the module. Rx output emphasis is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output equalization are defined as follows:

Table 13- Rx Output Emphasis Codes

Code Value	Bit pattern	Post-Cursor Equalization	Pre-Cursor Equalization
0	0000	No Equalization	No Equalization
1	0001	1 dB	0.5 dB
2	0010	2 dB	1.0 dB
3	0011	3 dB	1.5 dB
4	0100	4 dB	2.0 dB
5	0101	5 dB	2.5 dB
6	0110	6 dB	3.0 dB
7	0111	7 dB	3.5 dB
8-10	1000-1010	Reserved	Reserved
11-15	1011-1111	Custom	Custom

Note: The pre-cursor equalizer settings in dB approximates to
Pre EQ (dB)=-20*log10((1-(C(-1)/(C(-1)+C(0)+C(1))
The post-cursor equalizer settings in dB approximates to
Post EQ (dB)=-20*log10(1-(C(1)/(C(-1)+C(0)+C(1))

CMIS, 40-41.

1621

Claim 14	Disclosure
	The Rx Output Amplitude Control is a four-bit field per lane. The output amplitude is measured with no equalization enabled. Refer to Table 46 to determine if the module supports Rx Output Amplitude Control and Table 42 to determine which codes are supported. Output amplitude is defined at the appropriate test point defined by the relevant standard. The code values and the corresponding output amplitude are defined as follows: Table 14- Rx Output Amplitude Codes Code Value / Bit pattern / Output Amplitude 0 / 0000 / 100-400 mV (P-P) 1 / 0001 / 300-600 mV (P-P) 2 / 0010 / 400-800 mV (P-P) 3 / 0011 / 600-1200 mV (P-P) 4-14 / 0100-1110 / Reserved 15 / 1111 / Custom CMIS, 41.

1622

Claim 14	Disclosure

1623

1.7.4.10 Implemented Signal Integrity Controls

Table 46 describes the implemented signal integrity controls.

Table 46- Implemented Signal Integrity Controls (Page 01h)

Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze not implemented 1b=Tx Input Eq Freeze implemented	
	3	Adaptive Tx Input Eq implemented	0b=Adaptive Tx Input Eq not implemented 1b=Adaptive Tx Input Eq implemented	
	2	Tx Input Eq fixed manual control implemented	0b=Tx Input Eq Fixed Manual control not implemented 1b=Tx Input Eq Fixed Manual control implemented	
	1	Tx CDR Bypass control implemented	0b=Tx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Tx CDR Bypass control implemented	
	0	Tx CDR implemented	0b=Tx CDR not implemented 1b=Tx CDR implemented	
162	7-6	Reserved		RO RQD
	5	Staged Set 1 implemented	Staged Control Set 1 implemented on Pg 10h	RO RQD
	4-3	Rx Output Eq control implemented	00b=Rx Output Eq control not implemented 01b=Rx Output Eq Pre-cursor control implemented 10b=Rx Output Eq Post-cursor control implemented 11b=Rx Output Eq Pre- and Post-cursor control implemented	
	2	Rx Output Amplitude control implemented	0b=Rx Output Amplitude control not implemented 1b=Rx Output Amplitude control implemented	
	1	Rx CDR Bypass control implemented	0b=Rx CDR Bypass control not implemented (if CDR is implemented, it will be enabled) 1b=Rx CDR Bypass control implemented	
	0	Rx CDR implemented	0b=Rx CDR not implemented 1b=Rx CDR implemented	

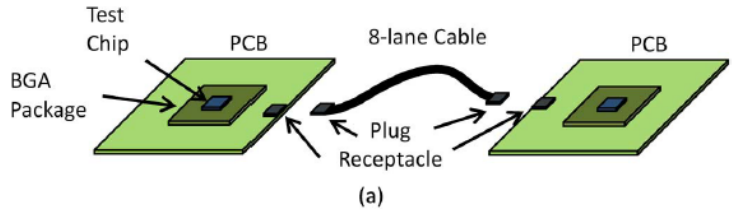
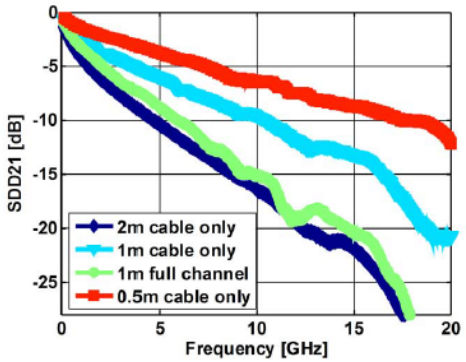
CMIS, 71.

Claim 14	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel loss below 25 dB. Channel equalization is performed by a combination of a 3-tap feed-forward equalizer (FFE), single-stage continuous-time linear equalizer (CTLE) and a 6-tap decision-feedback equalizer (DFE). Collaborative timing recovery is used to enable lane characterization without degrading jitter performance. Phase error decimation, with a conditional phase detection scheme, is used to reduce the DFE complexity by 50%. Power consumption over a wide range of data rates from 4 to 32 Gb/s is reduced by using regulated CMOS clocking with lane bundling, low swing transmitter with a source-series terminated (SST) driver and a highly reconfigurable receiver with an active inductor CTLE. At a lane data rate of 32 Gb/s, over a 0.5 m cable with 16 dB of loss, a transceiver lane consumes 205 mW from a 1.07 V supply. The power scales down to 26 mW from a 0.72 V supply at 8 Gb/s, when transmitting over a channel with 8 dB loss. The active silicon area per lane is 0.079 mm^2.** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

1624

Claim 14	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, transmitter (TX) and receiver (RX) circuits depending on data rate. Each lane consists of a TX with a low-swing voltage-mode driver and low power pre-emphasis, a data-rate scalable RX front-end (RXFE) with a continuous-time linear equalizer (CTLE) and decision-feedback equalizer (DFE), embedded clock and data recovery (CDR), lane clocking and adaptation and calibration logic. The link is targeted to applications whose data rate requirements may span from 4 to 32 Gb/s per lane. Supply voltage scaling, assuming a switching regulator, is employed to ensure that the link power decreases at least linearly with data rate [6]. However, even with supply voltage scaling, a link design solely optimized for performance may exhibit sublinear power scaling with data rate, as circuit blocks with static power consumption degrade power efficiency at lower rates. To overcome the challenge of power optimization across different data rates, a multi-mode design approach was adopted, transforming the lane from high performance mode (HP) to low power mode (LP). The HP and LP modes span 8–32 Gb/s and 4–12 Gb/s, respectively. Supply voltage scaling, power gating, bias control and transmit swing control are used to ensure optimal power-performance scaling. Musah, 3079-3080.

1625

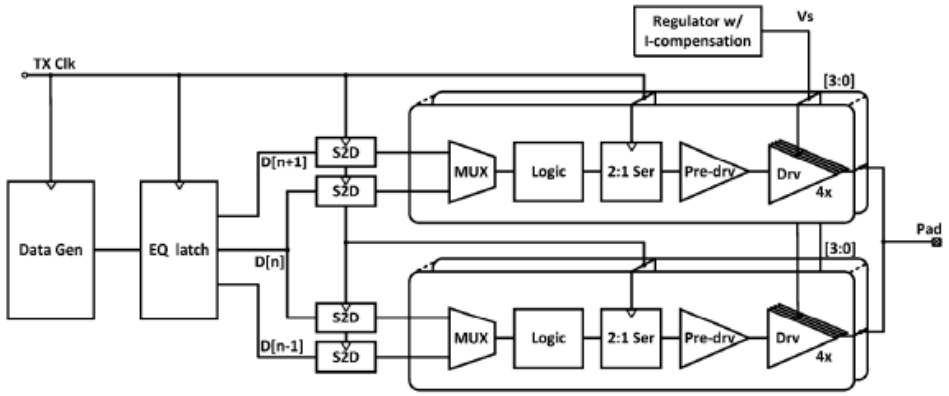
Claim 14	Disclosure
1626	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 14	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information from each lane to pursue wider jitter tracking BW, the per-lane phase contribution is decreased to one-fourth to reduce the receiver complexity. Moreover, unlike [8], this CDR scheme enables full eye margining characterization of one lane within the bundle while still transmitting data. When the collaborative CDR is enabled, the lane CDR is configured to track low-frequency inter-lane skew. Musah, 3081.

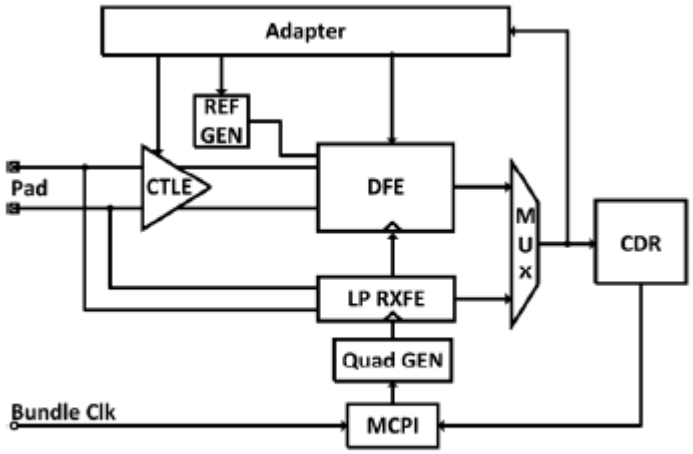
1627

Claim 14	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization is enabled. To support the bidirectional operation of the link, the voltage-mode driver performs two functions: (a) a data driver and termination in the TX mode; (b) a common-mode voltage generator and RX termination in the RX mode. The block diagram for the TX is shown in Fig. 6. An on-chip PRBS data generator (Data Gen) drives the equalization delay chain to produce the delayed versions of the signals needed for equalization. The reconfigurable 1 to 3-tap differential equalization is realized with 2 banks of 16-leg drivers, bringing the total to 32 legs. Based on the equalization MUX (EQ MUX) selection, the driver legs in the first bank are split between the primary-cursor and the post-cursor signal, while those in the other bank are split between the primary-cursor and the pre-cursor signal. Musah, 3082.

1628

Claim 14	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used to recover the clock and data in the LP mode. This results in a 4-way interleaved receiver, with each interleaved slice realized as a

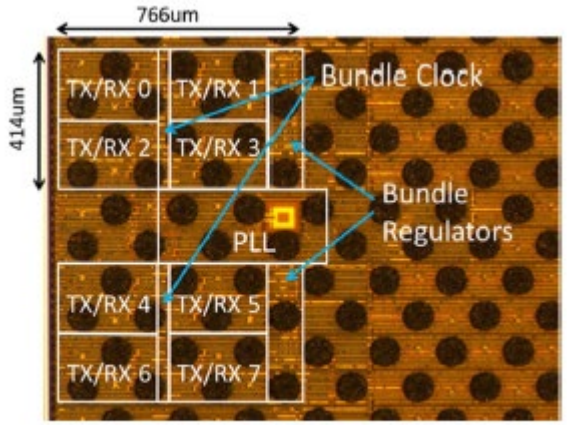
1629

Claim 14	Disclosure
1630	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (ISI). The HP RXFE is designed to equalize approximately 18 dB of the 24 dB channel loss target and includes a first-order CTLE with auto-gain control (AGC) and a 6-tap quarter-rate DFE. The DFE uses Mueller-Muller baud-rate phase detection [15] to avoid having an auxiliary equalized path for error detection as in [1]. A 6b resistor string DAC generates the reference voltage levels needed for the phase detection. A background adaptation unit monitors and controls the DFE coefficients and gain parameters. Musah, 3082.  Fig. 7. Multi-mode RX.

Claim 14	Disclosure
1631	Musah, 3082. Each comparator has a 9b segmented DAC (3b coarse, 6b fine) with programmable range and resolution to cancel the comparator offset and to provide the speculative tap value. The remaining five DFE taps are directly fed back into the summer. All DFE coefficients have a 6b control, with varying ranges, to ensure accurate ISI cancellation over various channel lengths and supply voltages. On-die logic controls foreground and background adaptation of CTLE, CML summer and sense amplifier offset voltages. Foreground calibration requires interruption of data transmission/reception while background calibration happens with the link fully operational. The logic also performs continuous AGC, DFE tap weight adaptation and reference voltage adaptation in the background using a sign-sign least mean squares (SSLMS) algorithm. To reduce design and area overhead of the RX adaptation, a unified adaptation macro unit (AMU) is configured to adapt or calibrate all RX parameters. The AMU block diagram is shown in Fig. 10. The RXFE output signals are first pre-processed to generate data, error and error valid signals. Correlation of data and error signal are used to adapt the DFE tap coefficients. If data mask is enabled, only the error signals are sent to the up/down counter enabling offset cancellation, AGC and reference voltage adaptation. De-asserting the enable signal freezes the AMU outputs allowing normal operation for foreground calibrated parameters like summer offsets. Musah, 3084.

Claim 14	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed to achieve the target BW. A CML summer following the track and hold sums the DFE currents from the non-speculative taps. To facilitate data-rate scalability of the DFE, the CML summer has to operate across a wide frequency range. Thus, the summer NMOS load is biased to be in triode and is coupled to the tail current to minimize the sensitivity of the output common-mode voltage to the summer bias. This biasing technique reduces the summer current at lower data rates to save power, while not sacrificing the sense amplifier performance. The voltage offset of the CML summer is cancelled with a digitally controlled offset current sourced into its output. The CML summer is followed by the speculative DFE tap latches and MUX. The three parallel latches can be configured for baud-rate phase detection and to receive data with or without tap speculation. To support the various latch configurations, there are three input differential pairs with inputs from the CML summer, the voltage reference generator and the coarse speculative tap coefficient. The speculative DFE tap is applied in a coarse-fine approach using a 3b device DAC (coarse) and a 6b current DAC (fine). This configuration enables wide voltage ranges that can realize equalization and comparator offset cancellation concurrently. The MUX provides additional amplification and speculative tap selection. When the speculative tap is disabled, the MUX is bypassed. Musah, 3085.

1632

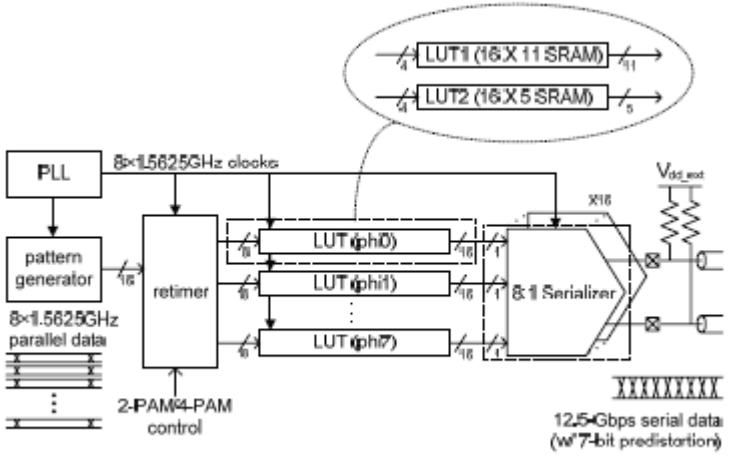
Claim 14	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

1633

Claim 14	Disclosure
1634	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane power efficiencies between 1 and 5.7 pJ/b. The area of the 8-lane transceiver is 0.634 mm^2. With a TX voltage swing of 630 $mV_{pp-diff}$ and lane CDR enabled, the 4-lane link demonstrated error-free ($BER < 10^{-12}$) transmission over a 0.5 m cable with PRBS10 data at 32 Gb/s per lane. Musah, 3088. ***Abstract*—This paper presents a 12.5-Gbps transmitter that uses a lookup table (LUT)-based equalizer to compensate for within-die imperfections. An equalization technique with 2x sampling is proposed to accommodate timing offsets in the multiphase clocks used for 8:1 serialization. LUT code remapping is also demonstrated to compensate for mismatch effects that introduce nonlinearity in the transmit DAC. Experimental results of a 7-bit resolution transmitter with 4-tap equalization, implemented in 0.13µm CMOS, show the LUT-based equalizer can significantly improve the signal integrity of an otherwise closed eye for data transmitted at 12.5-Gbps.** Chung, Abstract.

Claim 14	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. In addition to compensating for channel losses, the proposed equalizer can be used to mitigate various mismatch effects. The transmitter employs an 8-way interleaved structure using 8 multiphase clocks to alleviate on-die speed limitations. Unfortunately, these clock phases are susceptible to variations that lead to inter-symbol timing mismatch and cause eye closure. While phase adjusters can be employed to reduce phase mismatch [6,7], we can leverage the flexibility of the existing transmit equalizer to improve signal integrity by treating the mismatch effects as inter-symbol interference (ISI). Other within-die mismatch effects, such as DAC current mismatch and nonlinearity, and deviations in pulse peak positions, are also addressed by remapping LUT codes. While this paper relies on an external sampling scope and offline analysis, an on-chip sampling scope can be used to measure transmitter characteristics and drive an internal self-calibration engine [8]. Chung, 563.

1635

Claim 14	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1636

Claim 14	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random pattern generator (or input data memory) to generate the addresses for 8 parallel sets of LUTs, each running off of a different clock phase. The LUT uses a set of data (current, next, and two previous bits) as addresses that correspond to 4-tap mappings of transmit equalizer outputs. The parallel implementation was originally chosen to accommodate speed limitations of the automatically-generated memories, but the interleaved structure confers the benefit of supporting individual equalizer settings for each serializer path to compensate for phase mismatch effects. As shown in the inset of Fig.1, one LUT set comprises two 16-word SRAMs, instead of one 256-word SRAM, to further overcome speed limitations and to reduce power and area overheads. LUT1 uses a set of current and previous bits as an address to generate a portion of the transmitter output that consists of the main-cursor and first post-cursor components for two adjacent bits in each symbol output. The 11-bit code out of each LUT1 ultimately maps to one of 128 possible output levels (7-bit resolution) with 4-bit binary coding and thermometer coding for the remaining

1637

Claim 14	Disclosure
	higher-order bits. LUT2 uses a set of next and second previous bits as an address that maps to the pre-cursor and second-post cursor compensation components of the transmitter output. Since the pre-cursor and second post-cursor components typically exhibit smaller amplitudes in transmit-side equalization, the LUT2 output requires fewer bits (5-bit binary). Each LUT generates data for 1 out of 8 interleaved symbols combined together via an 8:1 serializer. The 11- and 5-bit codes out of each LUT are connected to 16 parallel serializers. Chung, 563-564. While the LUT-based equalizer can also compensate for channel loss, the focus of this section is on the compensation of within-die variations—improving the worst-case eye by compensating for phase mismatch and DAC current mismatch. Uneven spacing of the multiphase clocks is the dominant cause of signal integrity degradation, creating the worst-case eye. While this issue can also be addressed with phase adjusters [6,7], we investigate how the proposed LUT-based equalizer can also alone compensate for the clock phase mismatch effects. Chung, 565.

1638

Claim 14	Disclosure
	The digital signal processing (DSP) module of receiver module **300** includes a line R×D macro, which is a part of the data path. The DSP module provides many functionalities. Depending on the implementation, the DSP module may include an FFE block, a PAM-4 slicer, a PAM4-DFE, a PAM-4 reflection canceler, a PAM-4 level dependent equalizer, a PAM-4 1+D de-precoder, and/or a soft FEC decoder. Riani, 7:47-54. As an example, the FEC block may be implemented with an 6/10-tap FFE. The assumption is that the interface is 32-unit intervals (UI) interleaved. The 6/10-tap FFE, with gain and DC offset compensation, provides the equalization needed. The 10-taps are fully adaptive. The PAM-4 slicer block compares the FFE output data with thresholds (e.g., {Tn, Tm, Tp}), and produces the PAM-4 decisions. For example, thresholds {Tn, Tm, Tp} are adaptive and can be used to compensate for the level-dependent noise nature of optical channels. The PAM-4 decision feedback equalizer (DFE) is configured to provide nonlinear adaptive feedback on the inner levels in the data path. The line receiver module **300** may also include a PAM-4 reflection canceller to reduce

1639

Claim 14	Disclosure
	group-delay effects. The line receiver module **300** may additionally include a PAM-4 level dependent equalizer, which provides equalization for the two inner levels of the PAM-4 eyes. The PAM-4 1+D de-precoder block is configured to operate in DFE mode. In the DFE mode, to overcome DFE Error propagation, a 1/(1+D) pre-coder is employed on the transmit driver and on the receiver end. The final DFE decisions pass through the (1+D) filter to cancel the effect of pre-coder. The line receiver module **300** may also include a soft FEC that work with parity check code, which inserts parity information every 20 FEC symbols. It is to be appreciated that the soft FEC can provide a 1.5 dB gain with 5% overhead. In various implementations, SFEC supports in 200GE mode where the host data is KP4 FEC encoded. For example, the use of SFEC is described in patent application Ser. No. 15/691,023, filed Aug. 30, 2017, which is incorporated by reference herein. Riani, 7:55-8:17.

1640

Claim 14	Disclosure
	The host receivers are implemented with various features. For example, an OIF 56G-VSR compliant interfaces is implemented with adaptive CTLE/AGC that is able to equalize 10 dB channels loss plus up to 4 dB total package loss. Decision feedback equalizer (DFE) may be included for trace reflection cancellation. Independent CDRs can be configured in each lane with full-rate LC VCO. Similarly, independent power down control is provided for each lane. Host receivers may also include support for reference-less operation. Additional features include SNR monitor, adaptive equalizer convergence monitor, and loss of signal (LOS) detector with programmable thresholds, PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and PAM-4 MSB/LSB swap. Host receivers may also have diagnostic features, such as PRBS13Q, PRBS58/31/23/16/15/13/11/9/7 checkers, PRBS auto-polarity detection, and continuous identical symbol detection. Riani, 9:48-64.

1641

Claim 14	Disclosure
	FIG. **5** is a simplified block diagram illustrating a host side receiver module **500** according to embodiments of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, host side receiver module **500** implements section **201** in FIG. **2**. Receiver module **500** may operate in different modes with an aggregate bandwidth of 200 Gbps. For example, receiver module **500** supports eight lanes of NRZ data or four lanes of PAM-4 data. As shown in FIG. **5**, the receiver lanes are capable of receiving data streams at 28.125 Gbps in NRZ mode or 28.125 GBd in PAM-4 mode. For example, data received from the eight data lanes are transmitted through eight parallel data transmission lanes. The AFE **501** receives optical signal through optical communication channels. The AFE **501** is configured to handle various types of data format, such as 56 Gbps of PAM-4 data. In various embodiments, AFE **501** supports adaptive CTLE/AGC for equalizing 10 dB loss channels and up to 4 dB total package loss.

1642

Claim 14	Disclosure
	The AFE block **501** may also include a decision feedback equalizer covering h3 to h12 used to equalize reflections. The AFE block **501** may also include independent a clock data recovery (CDR) module with full rate inductive/capacitive (LC) voltage controlled oscillator (VCO). Each of the eight data lanes may have its own independent power down control. The AFE block **501** may additionally support reference-less operation (receiver is capable of coarse tuning and frequency locking its VCO to incoming data without an external pilot tone). The AFE block **501** may also include a signal-to-noise ratio (SNR) monitor and/or an adaptive equalizer convergence monitor. Additionally, the AFE block **501** may also include a loss of signal detector with programmable threshold (e.g., 50 mV to 100 mV). It is to be appreciated that the functionalities of AFE block **501** may be provided by the analog support block **502**. Riani, 10:6-41.

1643

Claim 14	Disclosure
	Data streams, once processed by the AFE block **501**, are transmitted through data path **504**. In various embodiments, data path block **504** provides PAM-4 inverse gray coding, polarity inversion, PAM-4 decoding, and/or PAM-4 MSB/LSB swapping. DSP block **503**, as shown, is coupled to AFE block **501**, data path block **504**, and CID detector block **505**. For example, DSP block **503** is capable of processing 32 UI of data and error samples per clock cycle. In various implementations, DSP block **503** provides six shared plus ten DFE LMS engines and two shared search engines. The DSP block **503** may additionally provide calibration and adaption (e.g., by determining operating parameters) for LFEQ, CTLE, AGC gain, sample offsets, thresholds, phases, and DFE. The DSP block **503** may also include eye monitor with 256 steps over 2 unit-intervals and 512 voltage steps. To perform signal processing, the DSP block **503** may be implemented with capture buffer for 32 unit-intervals of aligned data, error, and eye monitor samples and qualifiers. Riani, 10:42-60.

1644

Claim 14	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated by the on-chip PRBS generator, which are then interleavedly latched by the compact latch array to produce the 16-path quarter-rate data for the following four 4:1 MUXs. The desired timing relationship (see the signal positions in the latch array), which enables each MUX to share the same timing margin, is satisfied by 90°-spaced quarter-rate clock relatching. The full-rate UI-spaced outputs of the 4:1 MUXs are first buffered by the pre-drivers and then sent to the four-tap FFE combiner. In the clock bundle, a clock conditioner is employed

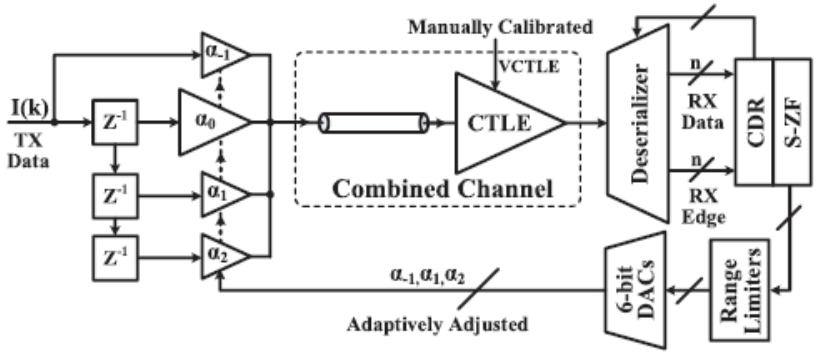
1645

Claim 14	Disclosure
	to convert the incoming single-end half-rate clock into differential outputs, which are then fed into a divider (DIV2) to generate the quart-rate I, Q clocks. After being transformed into full swing by the CML2CMOS converters, these clocks are further applied to four driving buffers and four pseudo-AND2s to produce 50% and 25% duty cycle clocks for the latch array and the 4:1 MUXs, respectively. The main feature of the TX chip is the compact implementation of the multiple 4:1 MUX-based four-tap FFE, which not only relaxes the stringent timing requirement of the final serialization stage, but also provides a robust approach to support a wide operation range. On the other hand, the doubled self-drain capacitance in the 4:1 MUX significantly reduces the bandwidth of the MUX, which is the key factor that constrains the maximum operation speed. Additionally, the output performance highly relies on the quality of the multi-phase gating clocks. The remainder of this section will focus on the enhancement of the 4:1 MUX, including topology consideration, unit cell improvement, and clocking techniques. Zheng, 2964.

1646

Claim 14	Disclosure
	III. Receiver Chip *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear equalizer (CTLE), a quarter-rate CDR, an FFE adaptation unit, and some testing circuits for the recovered data and clock measurements. The received signal is first equalized by the CTLE and then sliced by eight data/edge samplers, where the sampling clocks are generated by two quarter-rate compensating PIs and the sampling positions are adjusted by a CDR logic using bang–bang phase detectors (BBPDs). In addition, a newly developed S-ZF algorithm along with three 6-bit DACs is adopted to produce the bias voltages for the TX-FFE. The rest of this section focuses on the optimization techniques for the CDR, and the S-ZF algorithm will be elaborated in Section IV. Zheng, 2968.

1647

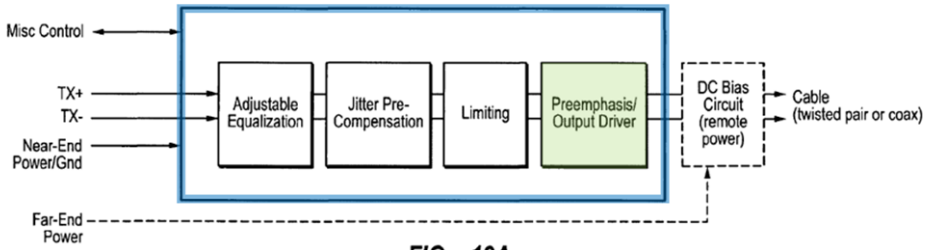
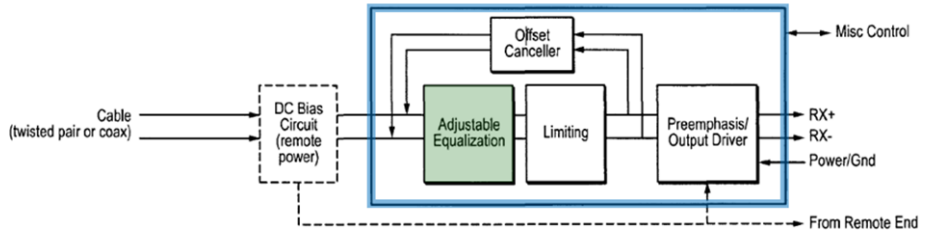
Claim 14	Disclosure
	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based S-ZF algorithm in the RX side. The digital tap weights generated by the S-ZF engine are first constrained by three range limiters and then applied to three 6-bit DACs to produce the bias voltages for the TX-FFE taps. These bias voltages are transferred to the TX through PCB traces. The TX-FFE is performed by a CML-based four-tap FFE combiner, where the tap weights are adjusted by changing the bias voltages of the current sources (see Fig. 1). The RX-CTLE schematic details and its frequency responses are described in Fig. 18. Zheng, 2970.  Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1648

Claim 14	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, the main tap weight is pre-fixed to accelerate the convergence speed. In each path, the edge and data streams with proper time shift are applied to a correlation detector (CD) to generate the residual correlation $\mathrm{ResCor}_l(n)$, which denotes $\mathrm{sign}[e(n)] \cdot D(n-l)$ in (1). These parallel correlation coefficients are first summed and then fed into a 16-bit integrator to execute the iteration of (1), where λ is determined by the subsequent truncation operation. In this design, a set of consecutive 4-bit data/edge of the 1/16-rate demultiplexed data/edge are employed, which ensures that the data/edge information used for equalization adaptation comes from different samplers. This decentralized error collection method reduces the possibility of non-optimal adaptation caused by imperfections, such as fabrication mismatch, duty cycle distortion, and I, Q quadrature error. Fig. 20 further details the operation principle and function table of the CD. Clearly, if there is no transition $[D(n) \oplus D(n+1) = 0]$, $\mathrm{ResCor}_l(n)$ is assigned 0. In case of a data transition $[D(n) \oplus D(n+1) = 1]$, $\mathrm{ResCor}_l(n)$ is assigned $+1$ or -1 when the polarities of $D(n-l)$ and $E(n)$ are identical $[D(n-l) \oplus E(n) = 0]$ or opposite $[D(n-l) \oplus E(n) = 1]$. Zheng, 2972.

1649

Claim 14	Disclosure
1650	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing bandwidth-adaptively adjusting LPFs into the clock path for data sampling, the CDR achieves high performance on both low-frequency jitter tracking and high-frequency jitter suppression. To further improve the CDR performance, a TA-based compensating PI is designed to optimize the phase-step uniformity and reduce the phase-spacing shift between edge-sampling and data-sampling clocks. A combined TX-FFE and RX-CTLE is employed to compensate for the channel loss, where a low-cost edge-data correlation-based S-ZF adaptation algorithm is proposed to automatically adjust the TX-FFE's tap weights. Zheng, 2976. "The filter circuit 206 may be designed to have various characteristics as required. For example, assuming attenuation of the high frequency region in the coaxial cable transmission line beforehand, the filter circuit 206 can have a characteristic that strengthens high frequency range components (transmission signal pre-equalization treatment). On the other hand, the filter circuit 206 can also have another characteristic that suppresses excessive high frequency components to suppress unintentional emission. By providing a filter circuit with these types of characteristics in the connector base 20, a desired output signal P31 spectrum can be obtained according to the usage." Ikeda, 7:44-55.

Claim 14	Disclosure
	U.S. Patent No. 7,762,727 to Aronson One obvious way to implement the equalization coefficients used by a FIR filter and an equalizer is during manufacturing-testing. *See, e.g.*, Aronson at 14:35-39 (Fig. 12A's pre-equalization parameters for transit over the line-side are "adjusted at factory setup"), 15:1-8 (Fig. 12B's equalization may be "fixed" or "adjustable but factory set" "at the time of the cable manufacture").  FIG. 12A  FIG. 12B

1651