Exhibit A-12

Invalidity Claim Chart for U.S. Patent No. 10,877,233 (the "'233 Patent")
in view of Obviousness References

As detailed in the cover pleading, and in addition to the reasons set forth in Exs. A-1 through A-11, the claims of the '233 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 '233 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 No. 7,233,617 to Gorecki ("Gorecki") was filed on April 10, 2003, was published on April 15, 2004, and claims domestic priority to U.S. Patent Application No. 10/269,446 filed on October 11, 2002. Gorecki qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

CREDO 2015
Amphenol v. Credo
IPR2025-00835

IEEE 802.3-2015 - IEEE Standard for Ethernet was published on March 4, 2016.[1] IEEE 802.3-2015 qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

U.S. Patent 9,337,993 to Marcel Lugthart et al. ("Lugthart") was filed on December 23, 2014. Lugthart qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

U.S. Patent 8,312,302 to Paul Baker et al. ("Baker") was filed on June 30, 2011. Baker qualifies as prior art under at least 35 U.S.C. § 102(a)(1) (AIA).

U.S. Patent 8,891,670 to Brian Ishaug ("Ishaug") was filed on November 9, 2012 and published on Nov. 18, 2014. Ishaug qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

U.S. Patent 8,867,928 to David Piehler ("Piehler") was filed on May 16, 2011 and published on Oct. 21, 2014. Piehler qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

U.S. Patent 6,879,641 to J. David Unger ("Unger") was filed, June 13, 2002, published on April 12, 2005, and qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

2

U.S. Patent Application 2013/0073749 A1 to Francois Tremblay et al ("Tremblay") filed on February 22, 2011 and published on March 21, 2013. Tremblay qualifies as prior art under at least 35 U.S.C. § 102(a)(1) and (2) (AIA).

[1] *See* https://ieeexplore.ieee.org/servlet/opac?punumber=7428774.

A. INDEPENDENT CLAIM 1

Claim 1	Disclosure
1[pre] A 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.

3

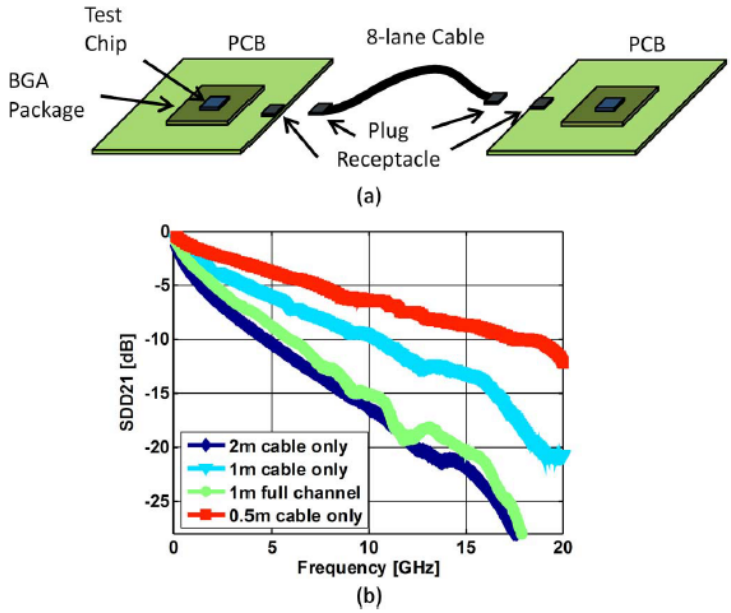
Claim 1	Disclosure
	Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and 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.

4

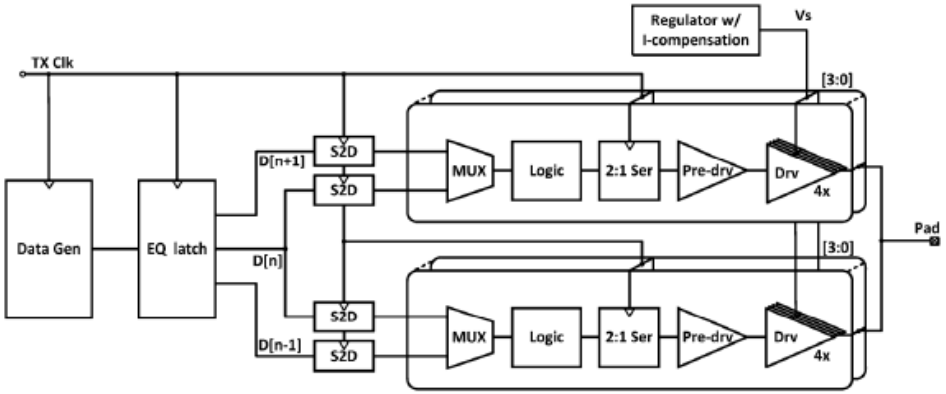
Claim 1	Disclosure
	***Abstract*—This paper details the design of an 8-lane bidirectional link for both within-the-box and external communications in 22 nm CMOS technology. A low profile connector with a high density cable assembly ensure a data rate of up to 32 Gb/s per lane while maintaining channel 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

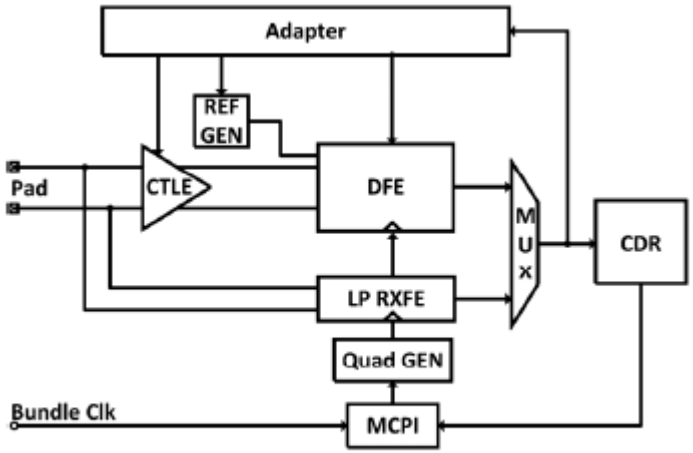
5

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
7	 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. 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.

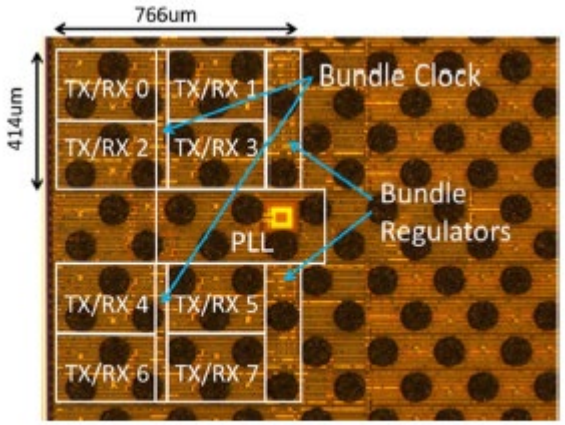
Claim 1	Disclosure
6	Musah, 3082.  Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase 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.

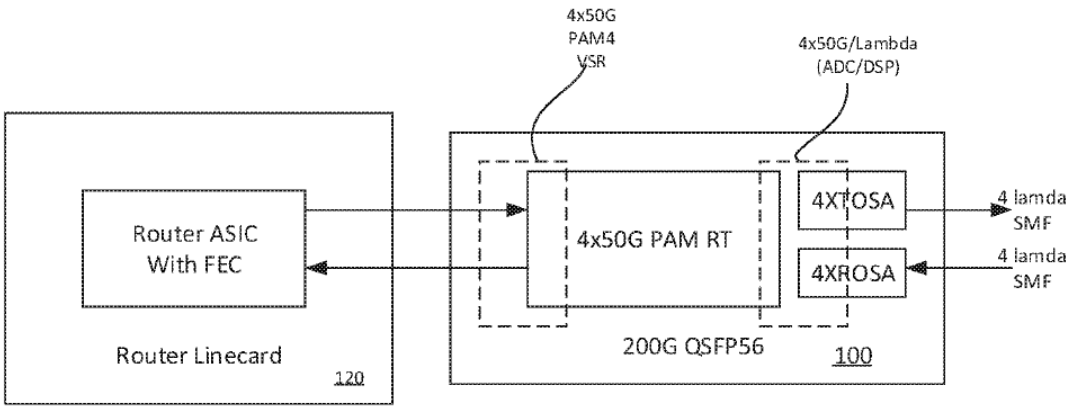
Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed 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.

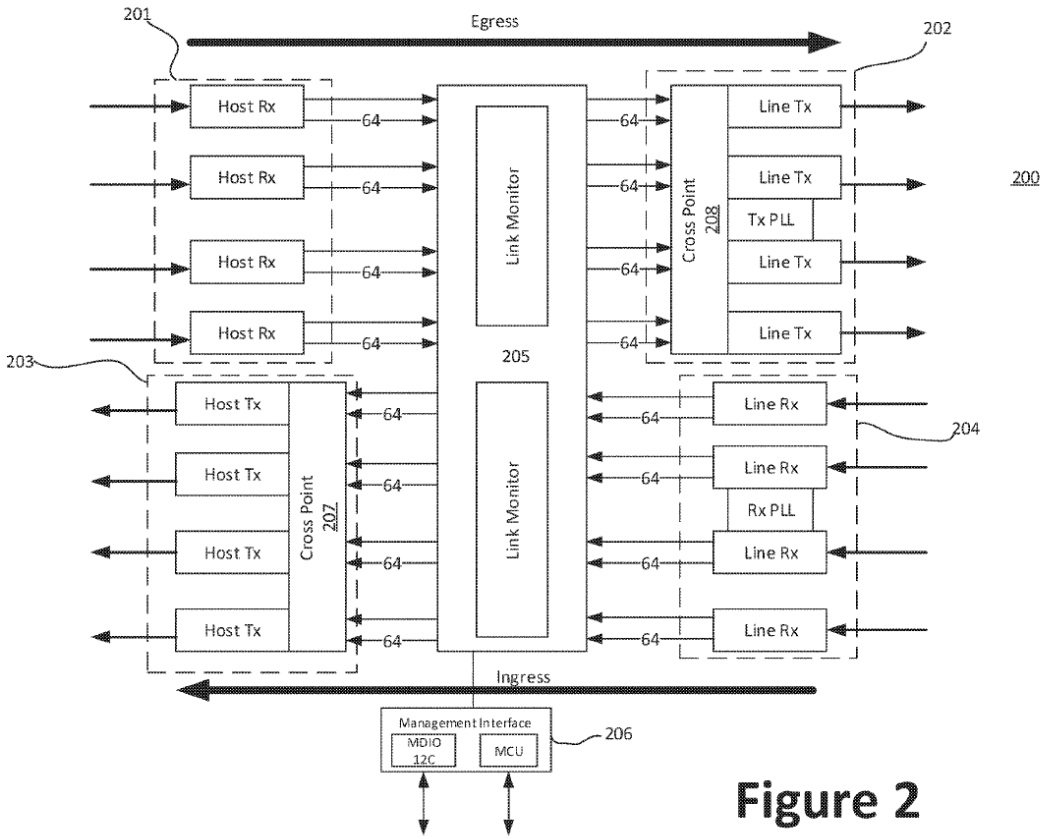
12

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.

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.

19

Claim 1	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 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.

23

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.

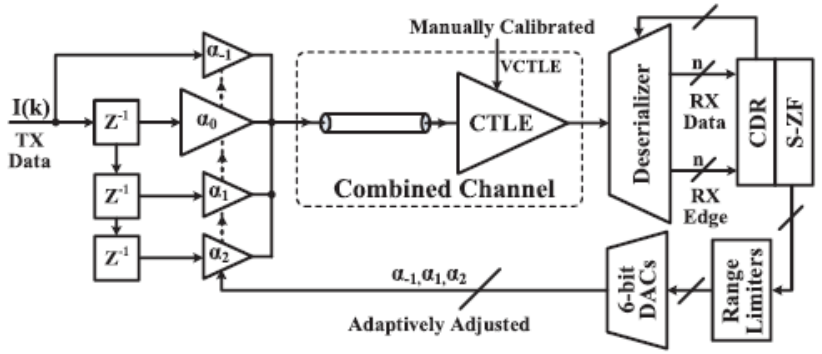
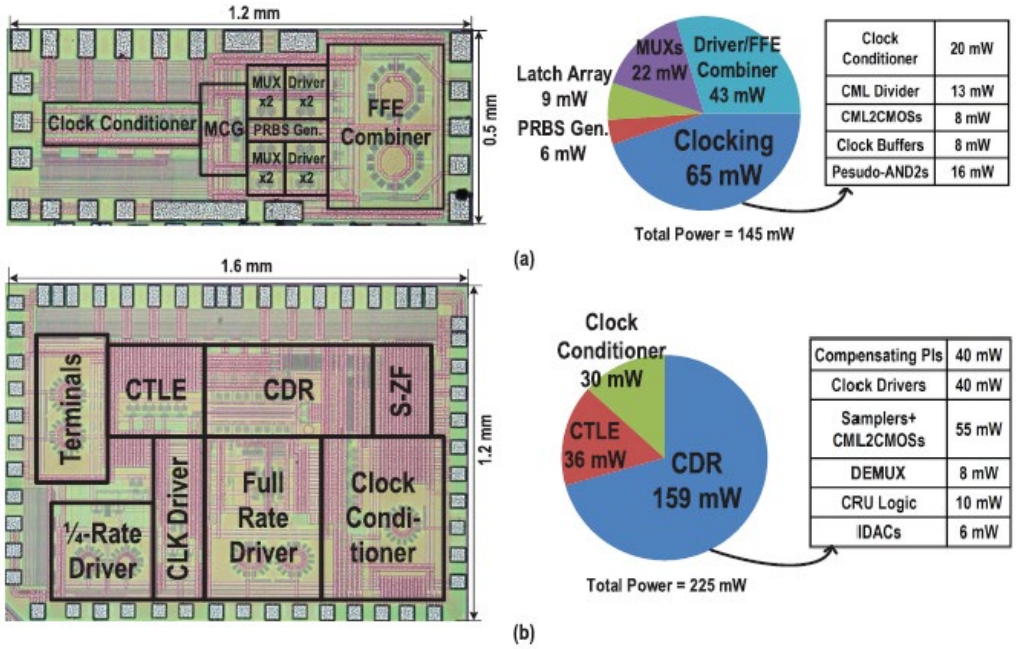


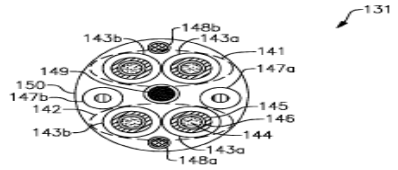
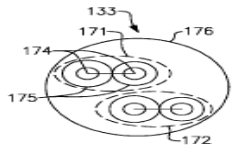
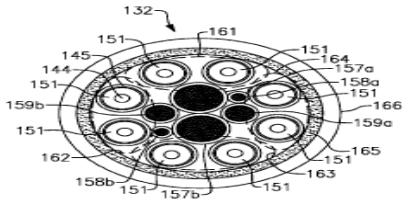
Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

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.

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.

28

Claim 1	Disclosure
	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A-3C. "The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables." Lugthart, column 11, lines 51-65.

Claim 1	Disclosure
	"Circuits, methods, and apparatus that provide for the powering of active components in connector inserts at each end of a cable may in various ways. For example, where a host is coupled to a device that is not self-powered, the host may provide power for circuitry at each end of the cable. In various embodiments of the present invention, the device may request higher voltage from the host, such that more power can be delivered. In these cases, the device may regulate the voltage received from the host to a lower voltage, and then provide the lower voltage to circuitry at one or both ends of the cable. Where the host is connected to a device that is self-powered, the host and the self-powered device may power their respective connector insert circuits." Baker, Abstract. 0027. In high speed data networks a host network element Such as a Switch or router (also referred to as a host) can be connected to a transmission medium (for example fibre cable, copper cable or other physical medium) through an interface module that processes data coming in from the transmission medium to place it into a condition suitable for the host and that processes data from the host to place it into a condition suitable for the transmission medium. Tremblay, paragraph 27. 0029. Example embodiments will now be described with reference to FIG. 1 which illustrates a host 102 that is con nected through a host channel 103 to a pluggable network transceiver interface module 100, which in turn is connected to a transmission medium 104. In the illustrated embodiment, the transmission medium 104 includes fibre optic cable 124 for transmitting high speed data from the interface module 100 and fibre optic cable 126 for transmitting high speed data to the interface module 100, however embodiments described herein can also be applied to non-fibre transmission mediums Such as copper for example. Tremblay, paragraph 29.

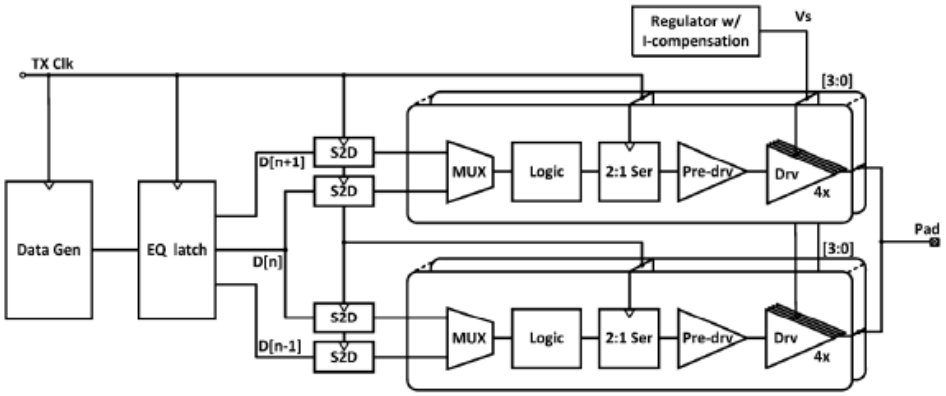
Claim 1	Disclosure
1[a] a first data recovery and remodulation (DRR) device that exchanges inbound and outbound multi-lane data streams with a first host interface port via a first end connector plug;	For example, and without limitation, at least under Credo's apparent infringement theory, the following 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

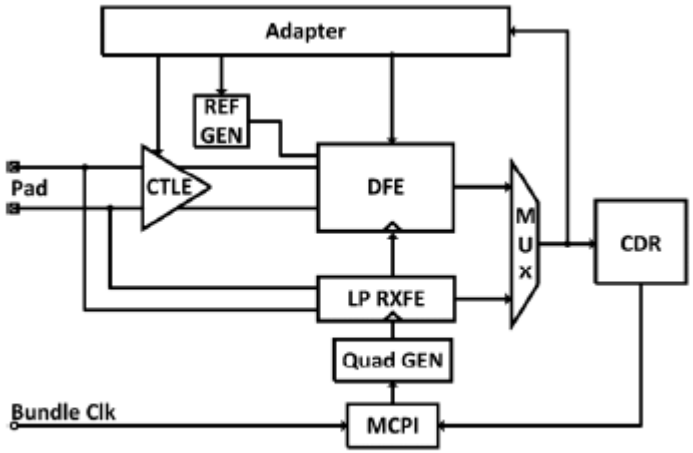
Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, 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
	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.

33

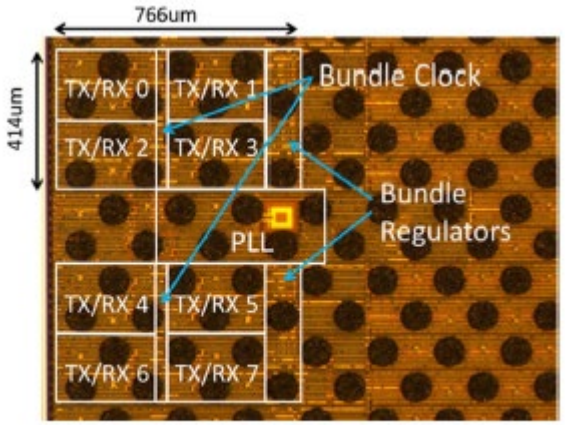
Claim 1	Disclosure
	With differential RX termination, an unequalized voltage-mode driver output stage dissipates about a quarter of the current-mode driver output stage power [12]. The differential equalization technique [13] is utilized to improve the voltage-mode driver power efficiency when equalization 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
35	 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
36	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.

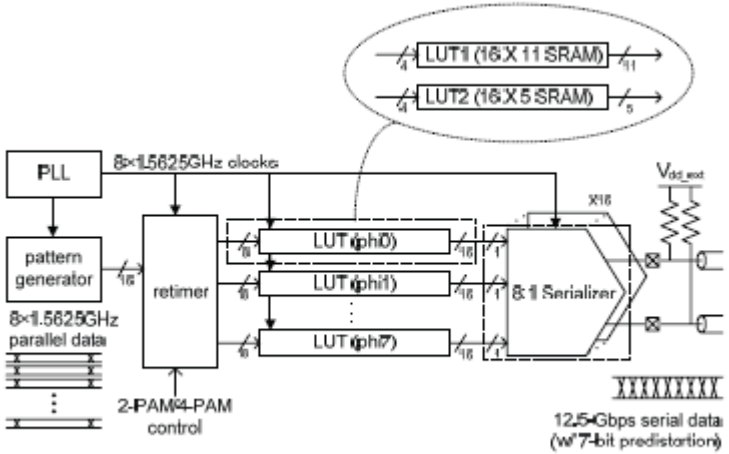
Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses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

40

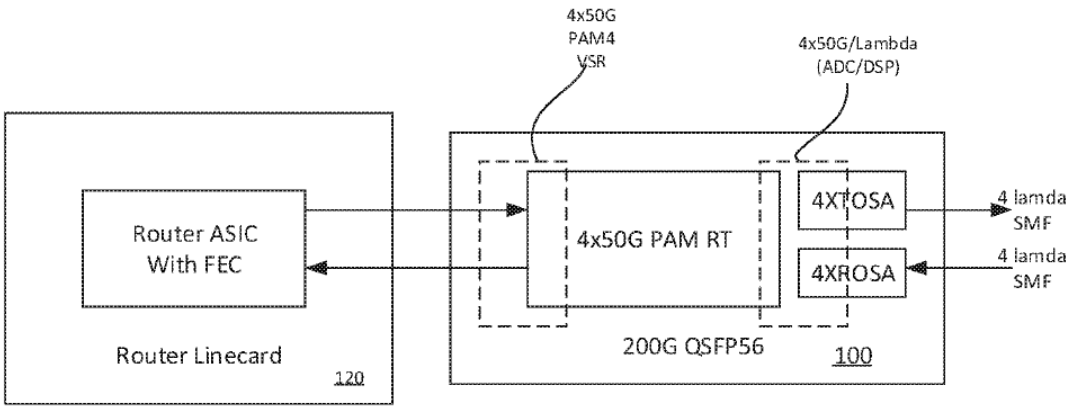
Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. 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.

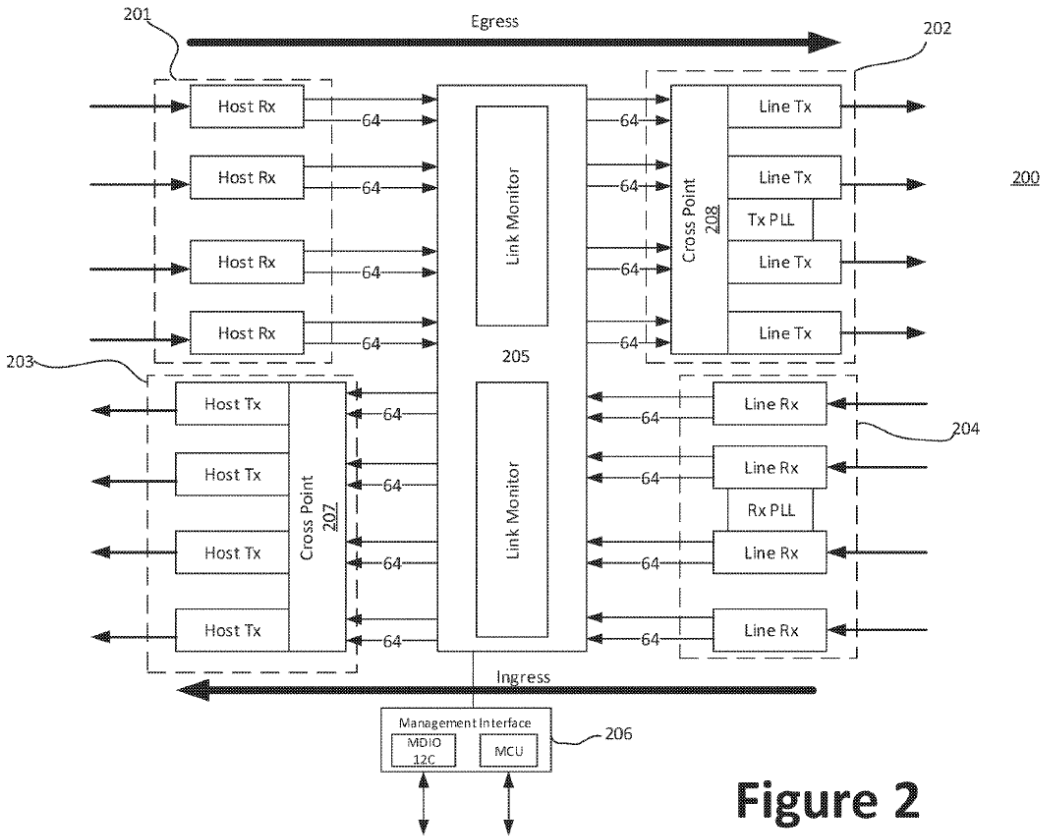
41

Claim 1	Disclosure
	 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.

Claim 1	Disclosure
	 **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.

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.

48

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.

49

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.

51

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.

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.

54

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.

55

	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.	

56

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.

57

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
59	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.

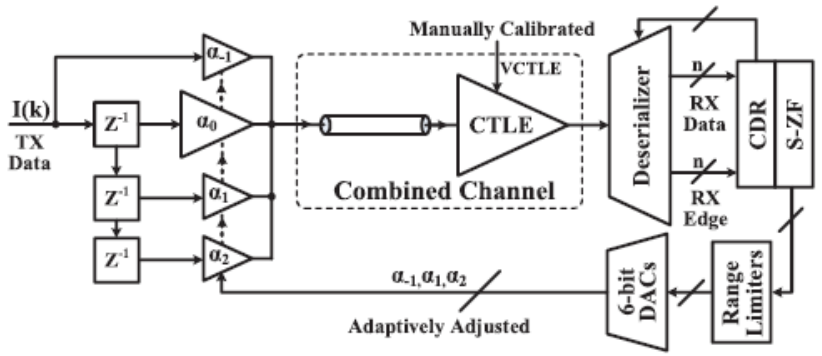
Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated 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

60

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.

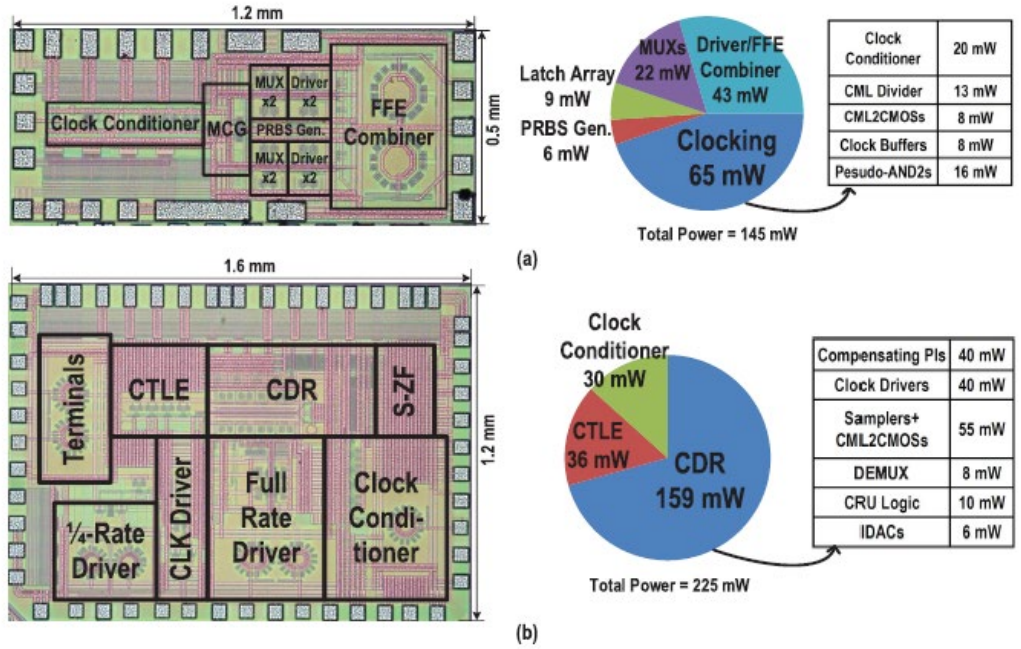
61

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.

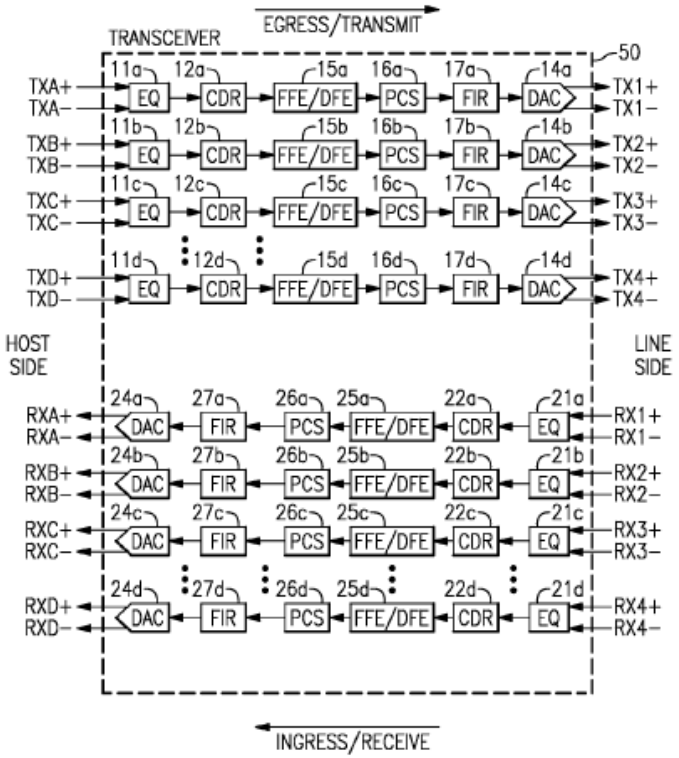
Claim 1	Disclosure
63	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based 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.

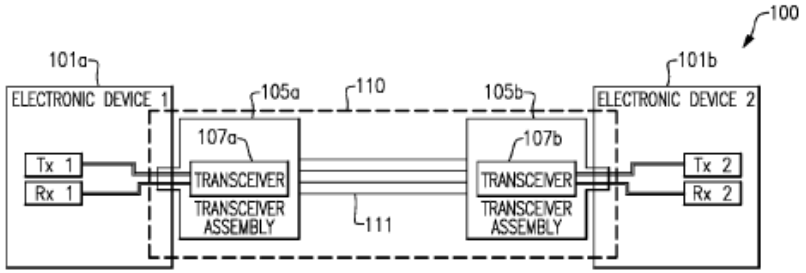
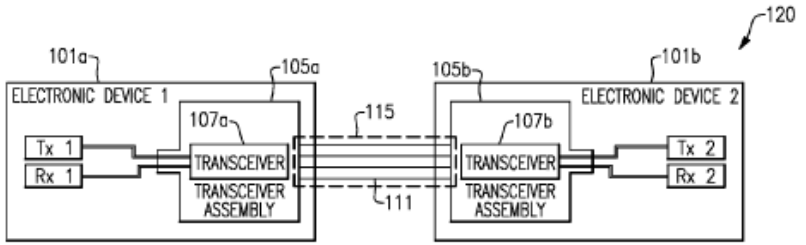
Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

64

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

Claim 1	Disclosure
66	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.

Claim 1	Disclosure
	 FIG.1D Lugthart, Figure 1D.

Claim 1	Disclosure
68	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

Claim 1	Disclosure
	According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. A first host-side signal of the first Lugthart, 3:66–4:5.

69

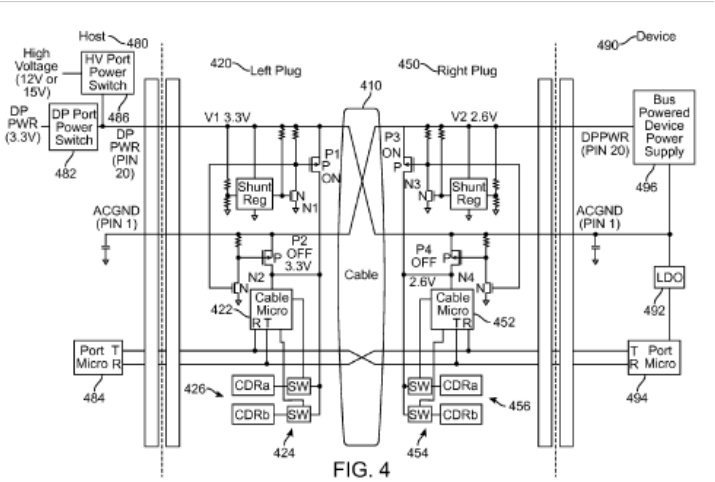
Claim 1	Disclosure
	FIG. **1**D is a schematic diagram of another embodiment of a transceiver **50**. The transceiver **50** includes a first transmit path equalizer **11***a*, a second transmit path equalizer **11***b*, a third transmit path equalizer **11***c*, a fourth transmit path equalizer **11***d*, a first transmit path CDR circuit **12***a*, a second transmit path CDR circuit **12***b*, a third transmit path CDR circuit **12***c*, a fourth transmit path CDR circuit **12***d*, a first transmit path FFE/DFE block **15***a*, a second transmit path FFE/DFE block **15***b*, a third transmit path FFE/DFE block **15***c*, a fourth transmit path FFE/DFE block **15***d*, a first transmit path PCS block **16***a*, a second transmit path PCS block **16***b*, a third transmit path PCS block **16***c*, a fourth transmit path PCS block **16***d*, a first transmit path FIR **17***a*, a second transmit path FIR **17***b*, a third transmit path FIR **17***c*, a fourth transmit path FIR **17***d*, a first transmit path DAC **14***a*, a second transmit path DAC **14***b*, a third transmit path DAC **14***c*, a fourth transmit path DAC **14***d*, a first receive path equalizer **21***a*, a second receive path equalizer **21***b*, a third receive path equalizer **21***c*, a fourth receive path equalizer **21***d*, a first receive path CDR circuit **22***a*, a second receive path CDR circuit **22***b*, a third receive path CDR circuit **22***c*, a fourth receive path CDR circuit **22***d*, a first receive path FFE/DFE block **25***a*, a second receive path FFE/DFE block **25***b*, a third receive path FFE/DFE block **25***c*, a fourth receive path FFE/DFE block **25***d*, a first receive path PCS block **26***a*, a second receive path PCS block **26***b*, a third receive path PCS block **26***c*, a fourth receive path PCS block **26***d*, a first receive path FIR **27***a*, a second receive path FIR **27***b*, a third receive path FIR **27***c*, a fourth receive path FIR **27***d*, a first receive path

70

Claim 1	Disclosure
	FIR **27***c*, a fourth receive path FIR **27***d*, a first receive path DAC **24***a*, a second receive path DAC **24***b*, a third receive path DAC **24***c*, and a fourth receive path DAC **24***d*. The illustrated transceiver **50** operates as a re-timer. For example, the transceiver **50** can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver **50** can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. In the illustrated configuration, the transceiver **50** receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver **50** receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver **50** is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1–53.

71

	FIG. **2**A illustrates an embodiment in which the cable **110** is an active cable including actively powered componentry for improving performance of the cable **110**. Thus, the cable **110** shown in FIG. **2**A includes one or more conductive lines **111** as well as first and second transceiver assemblies **105***a*, **105***b* positioned at either end of the conductive lines **111**. For example, the first transceiver assembly **105***a* includes a host side that is electrically connected to the first electronic device **101***a* and a line side that is electrically connected to a first end of the conductive lines **111**. Additionally, the second transceiver assembly **105***b* includes a host side that is electrically connected to the second electronic device **101***b* and a line side that is electrically connected to a second end of the conductive lines **111**. Although FIG. **2**A illustrates a configuration including the conductive lines **111**, in other configurations the conductive lines **111** can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies **105***a*, **105***b* can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers **107***a*, **107***b*, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. **1**A-**1**D. FIG. **2**B is a schematic diagram of a high-speed point-to-point communication system **120** according to another embodiment. In contrast to the communication system **100** shown in FIG. **2**A which includes an active cable **110**, the communication system **120** of FIG. **2**B includes a passive cable **115**. As shown in FIG. **2**B, active componentry including the transceivers **107***a*, **107***b* is situated externally to the cable **115**. For example, the first transceiver assembly **105***a* can be included in the first electronic device **101***a*, and the second transceiver assembly **105***b* can be included in the second electronic device **101***b*.

Claim 1	Disclosure
	Lugthart, 14:3–38. to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:19–23.  FIG. 4

Claim 1	Disclosure
	Baker, Figure 4. FIG. **4** illustrates an electronic system according to an embodiment of the present invention. This figure includes host **480** coupled to device **490** via cable **410**. Cable **410** includes left plug **424** connected to host **480**, and right plug **454** connected to device **490**. Host **480** may be capable of providing one or more voltages to both device **490** and cable circuitry in cable **410**. Host **480** may provide a lower voltage of 3.3 V, or a higher voltage of 12 or 15 volts. In other embodiments of the present invention, host **480** may provide various voltage levels to host **490** and cable circuitry in cable **410**. Baker, 5:25–36. In this example, power supply **496** receives only 3.3 V on line V**1** from host **480**. At this voltage, power supply **496** may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply **496** does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits **426** and **456**. These clock and data recovery circuits may receive and retime data received from host **480**, device **490**, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12.

Claim 1	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. Tremblay, paragraph 31.
1[b] a second DRR device that exchanges inbound and outbound multi-lane data streams with a second host interface port via a second end connector plug; 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:

75

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

76

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.

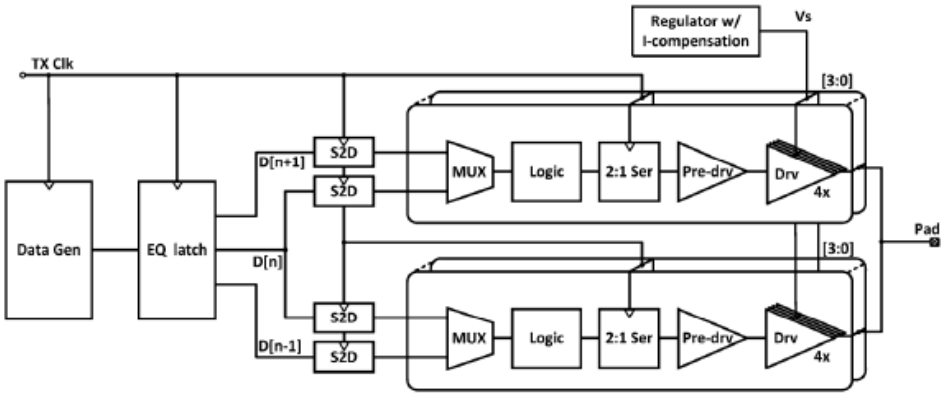
77

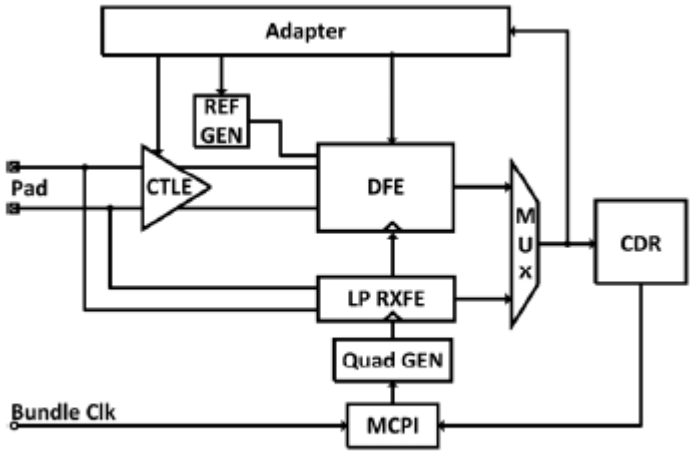
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

78

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.

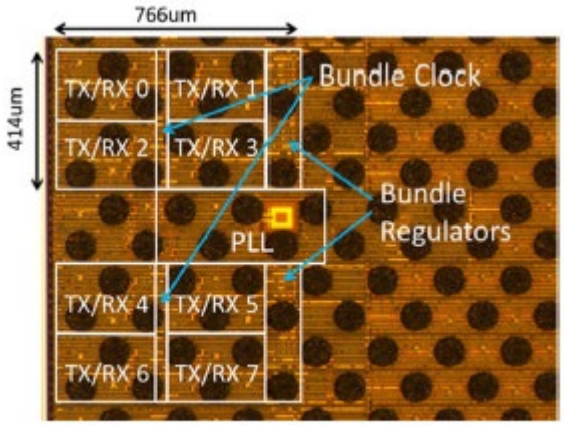
79

Claim 1	Disclosure
80	 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
	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.

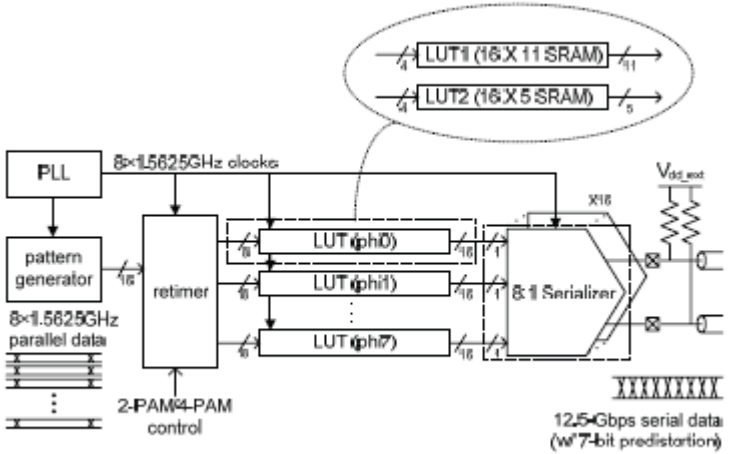
Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed 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.

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.

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.

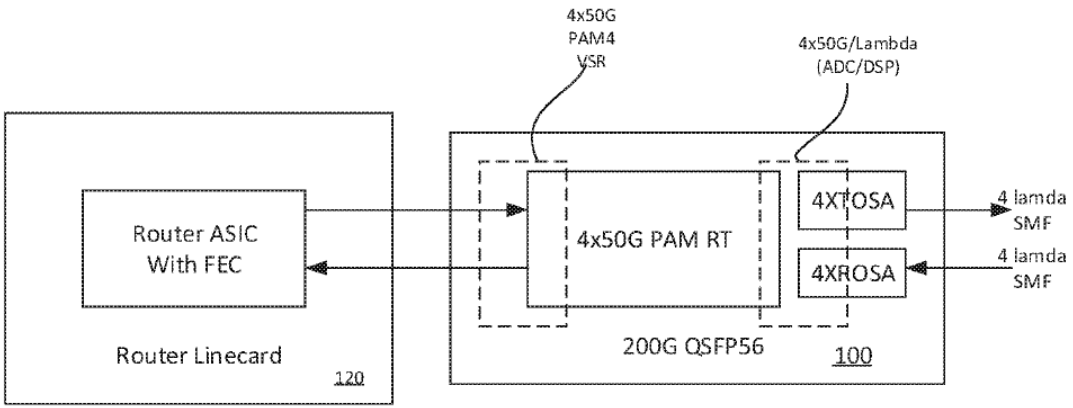
86

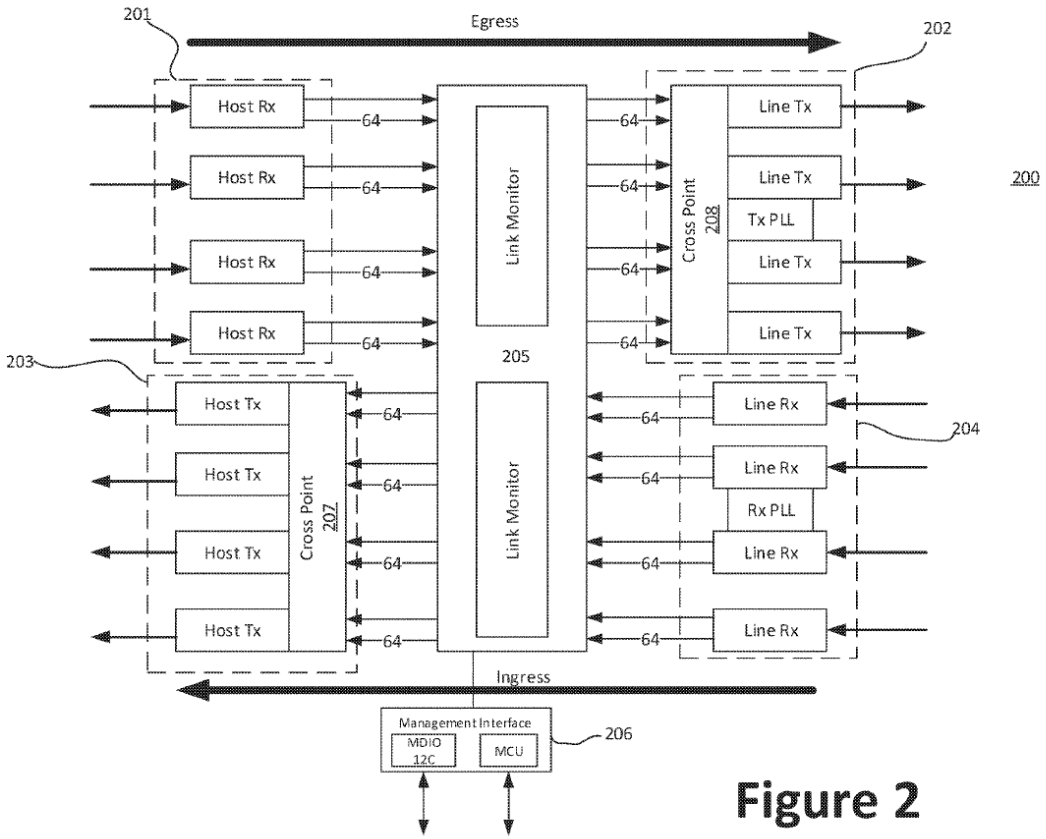
Claim 1	Disclosure
87	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

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

88

Claim 1	Disclosure
89	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 1	Disclosure
06	 **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.

92

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.

94

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.

95

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.

96

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.

97

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.

98

Claim 1	Disclosure
66	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.

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.

100

	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.	

101

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.

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
104	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

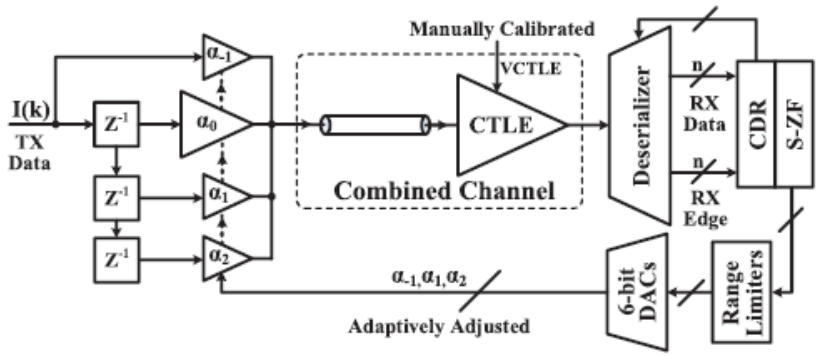
105

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.

106

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.

107

Claim 1	Disclosure
108	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.

Claim 1	Disclosure

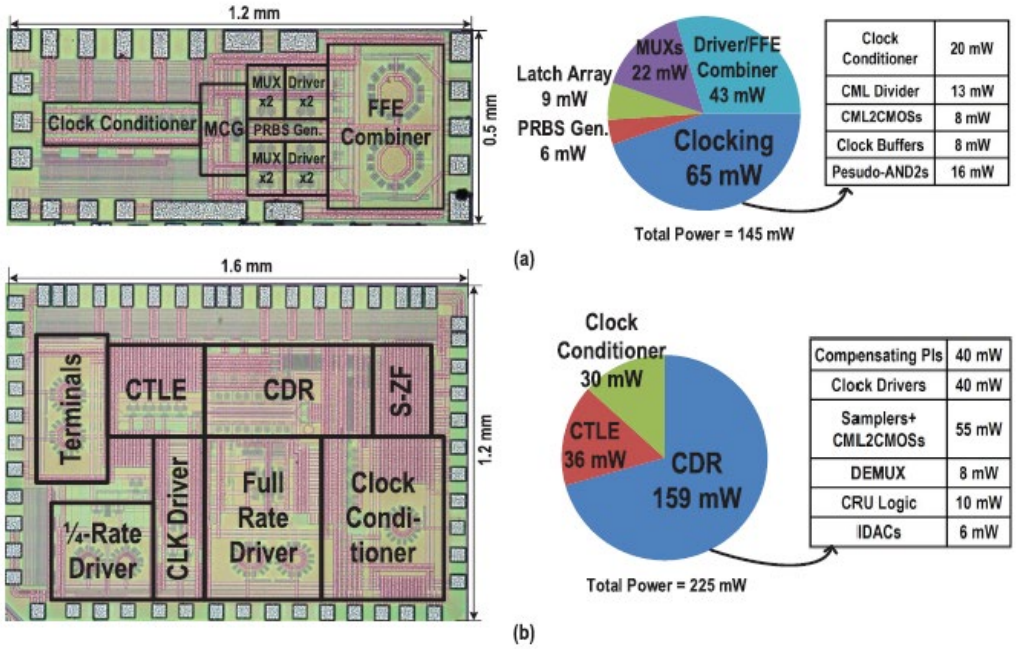


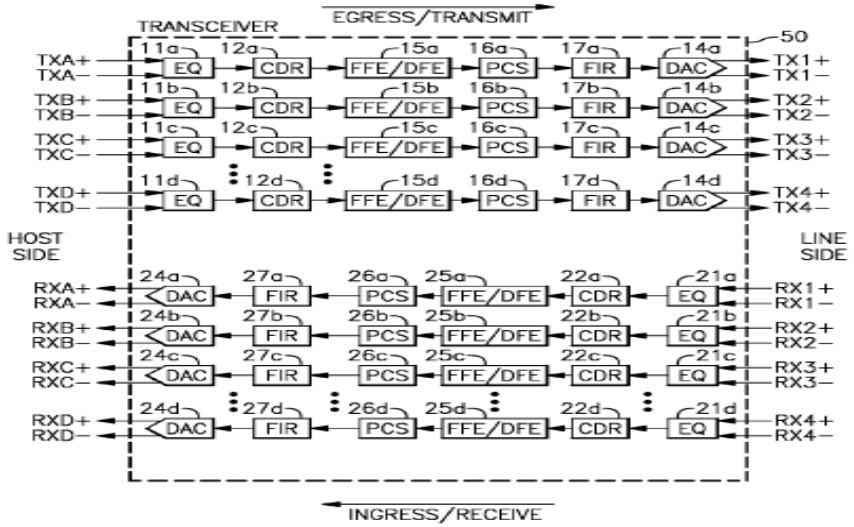
Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip.

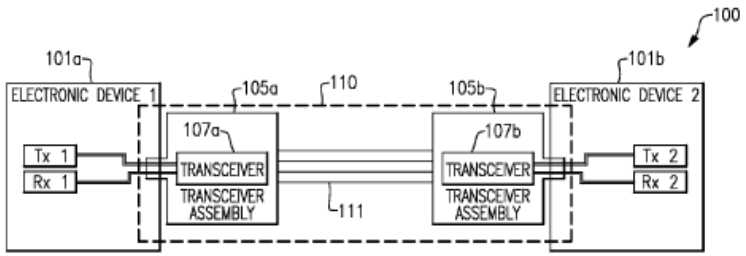
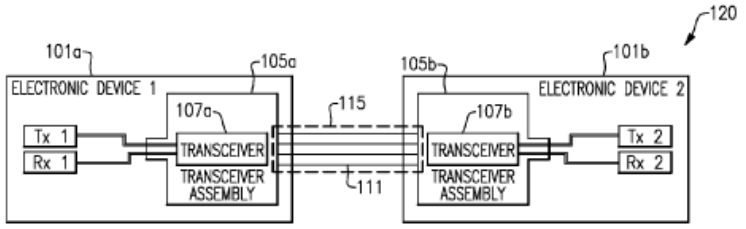
Zheng, 2973.

110

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.

111

Claim 1	Disclosure
	 FIG.1D Lugthart, Figure 1D.

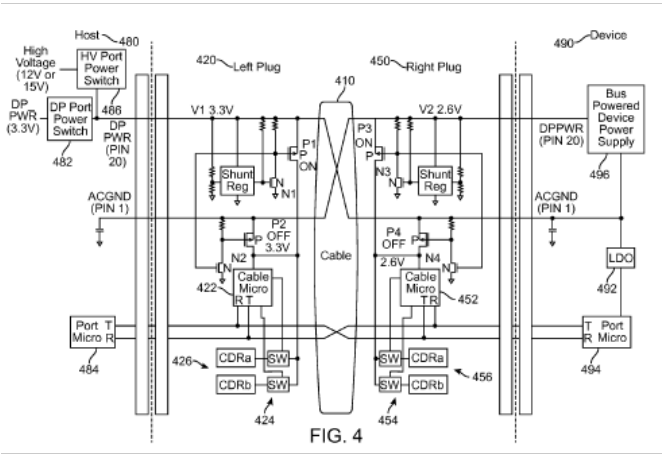
Claim 1	Disclosure
113	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B. According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver

Claim 1	Disclosure
	including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:44–4:5. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances.

114

Claim 1	Disclosure
	In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1–53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b.

115

Claim 1	Disclosure
	Lugthart, 14:3–38. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:19–23.  FIG. 4

Claim 1	Disclosure
	Baker, Figure 4. FIG. 4 illustrates an electronic system according to an embodiment of the present invention. This figure includes host 480 coupled to device 490 via cable 410. Cable 410 includes left plug 424 connected to host 480, and right plug 454 connected to device 490. Host 480 may be capable of providing one or more voltages to both device 490 and cable circuitry in cable 410. Host 480 may provide a lower voltage of 3.3 V, or a higher voltage of 12 or 15 volts. In other embodiments of the present invention, host 480 may provide various voltage levels to host 490 and cable circuitry in cable 410. Baker 5, 26–36. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12.

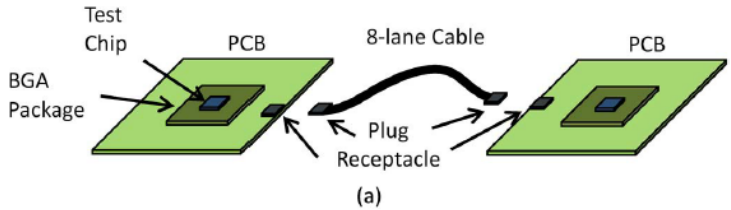
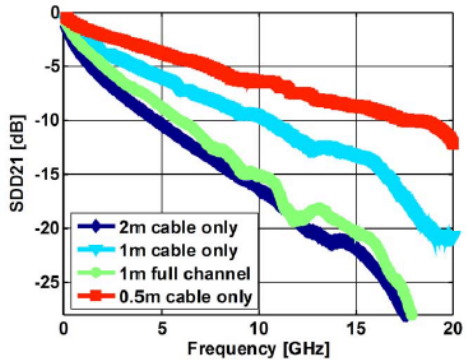
117

Claim 1	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. Tremblay, paragraph 31.
1[c] electrical conductors connecting the first and second DRR devices to convey electrical transit signals therebetween,	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
	***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

119

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.

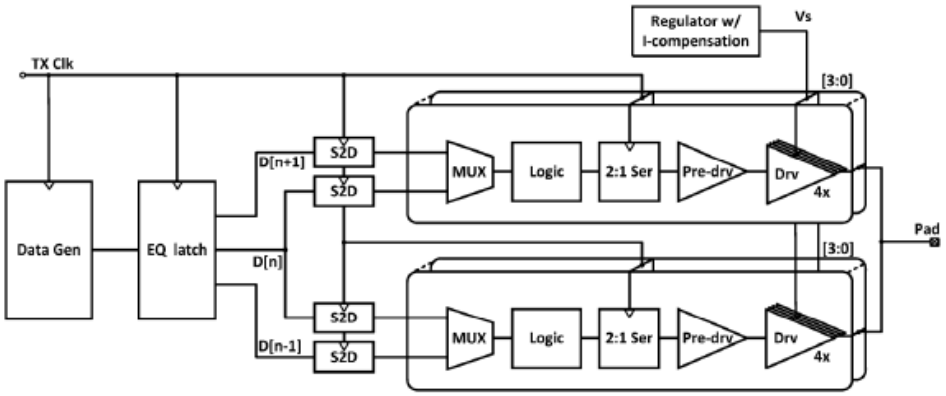
Claim 1	Disclosure
121	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

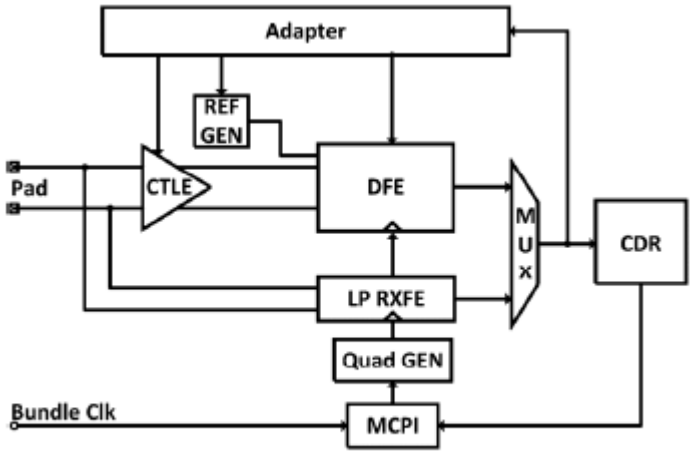
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

122

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.

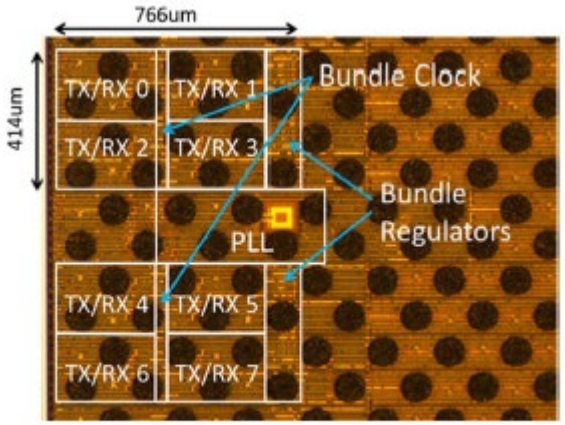
123

Claim 1	Disclosure
124	 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
125	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.

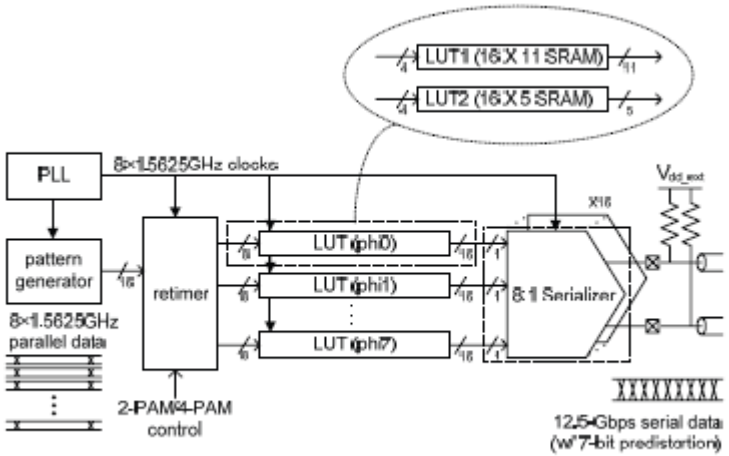
Claim 1	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed 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.

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.

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.

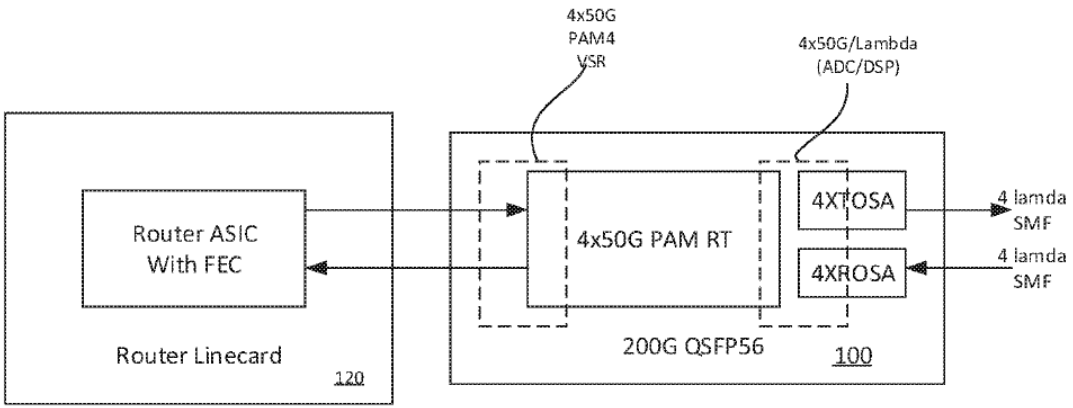
130

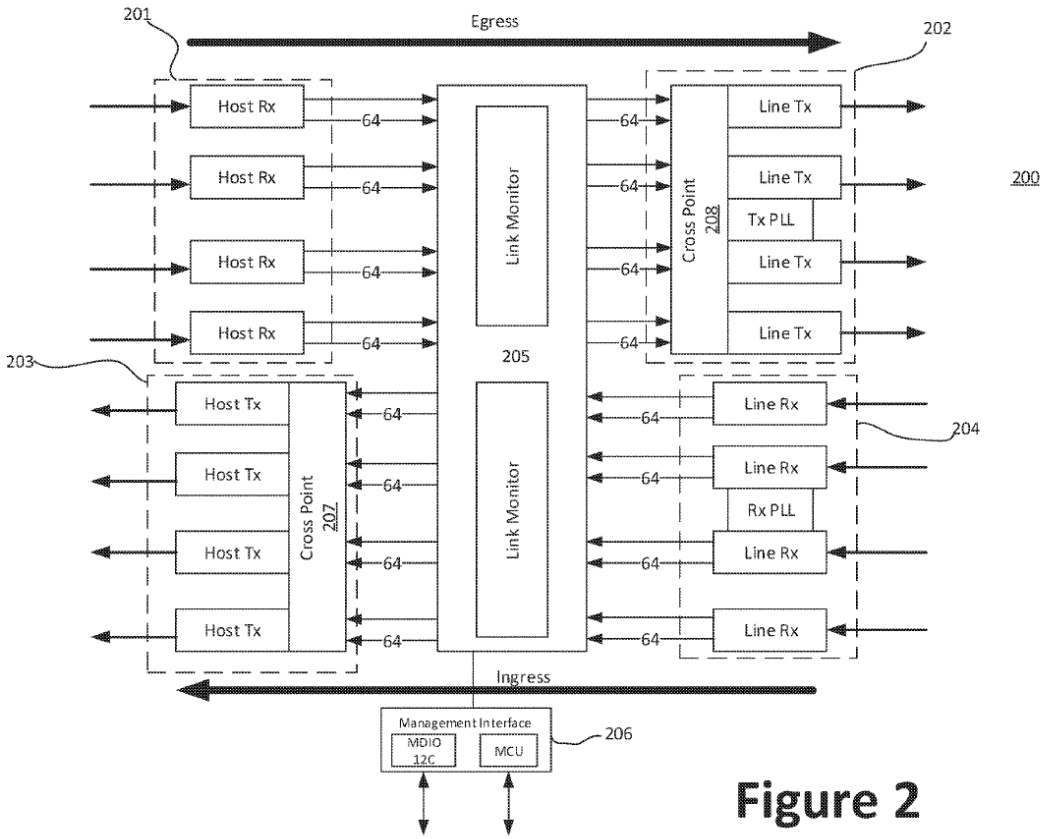
Claim 1	Disclosure
131	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

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.

133

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

Claim 1	Disclosure
135	 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.

137

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.

138

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.

139

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.

140

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.

141

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.

142

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.

143

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.

144

	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.	

145

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.

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.

147

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.

148

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

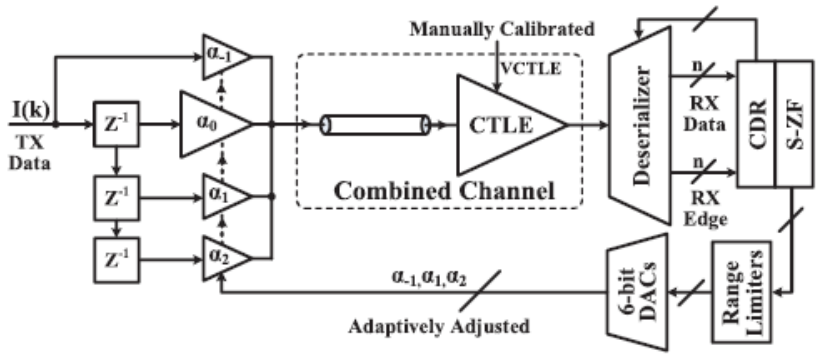
149

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.

150

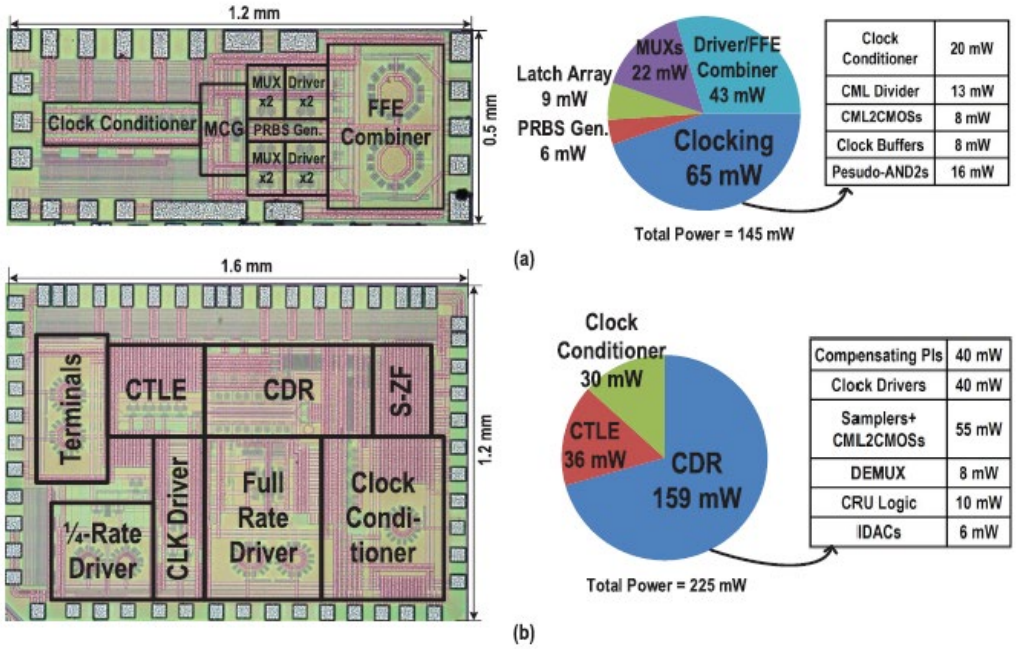
Claim 1	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear 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.

151

Claim 1	Disclosure
152	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.

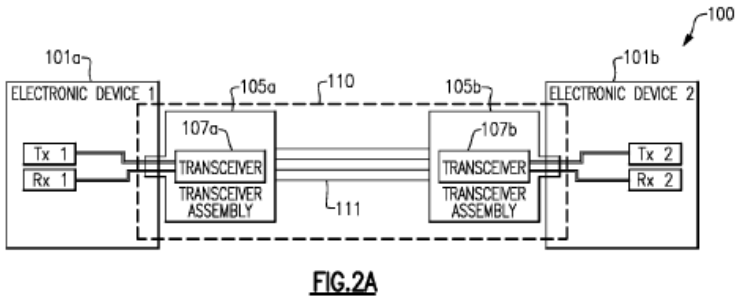
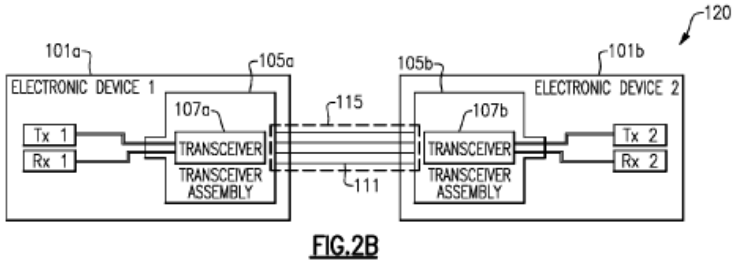
Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

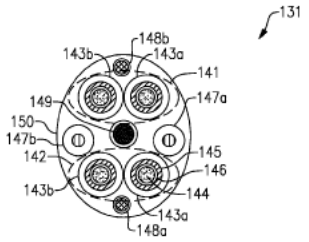
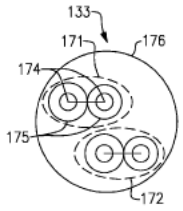
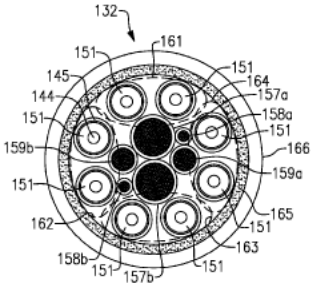
153

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

154

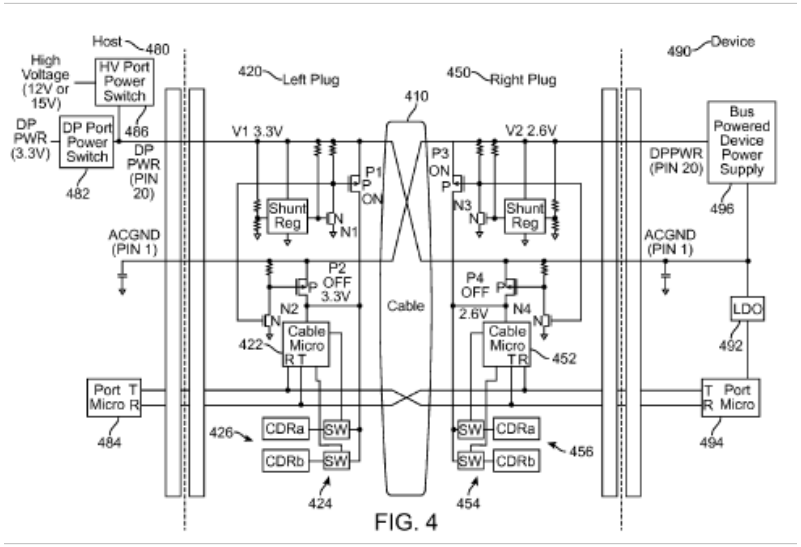
Claim 1	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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.

Claim 1	Disclosure
156	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

Claim 1	Disclosure
157	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3C. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described,

Claim 1	Disclosure
	other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. Lugthart, 14:3–20. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14: 30–38.

158

Claim 1	Disclosure
159	 FIG. 4 Baker, Figure 4. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and

Claim 1	Disclosure
	from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12. **[0029]** Example embodiments will now be described with reference to FIG. **1** which illustrates a host **102** that is connected through a host channel **103** to a pluggable network transceiver interface module **100**, which in turn is connected to a transmission medium **104**. In the illustrated embodiment, the transmission medium **104** includes fibre optic cable **124** for transmitting high speed data from the interface module **100** and fibre optic cable **126** for transmitting high speed data to the interface module **100**, however embodiments described herein can also be applied to non-fibre transmission mediums such as copper for example. In the illustrated embodiment, the host **102** includes a network interface card or board **106** that allows data to be exchanged between the host **102** and the network interface module **100**. In at least one example embodiment, the host interface board **106** includes data processing circuitry, which could for example be implemented an ASIC **108** ("Application Specific Integrated Circuit"), and a SERDES ("SERializer-DESerializer") chip or module **110**. SERDES module **110** receives parallel data from the host ASIC **108** over a parallel data channel **170** and converts it into high speed serial data for transmission over the host channel **103** to the network interface module **100**, and receives high speed serial data from the network interface module **100** through the host channel **103** which it then converts into parallel data and passes on to the ASIC **108** over the parallel data channel **170**. Tremblay, paragraph 29.
1[d] the first DRR device converting between said electrical transit signals and 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:

160

Claim 1	Disclosure
inbound and outbound multi-lane data streams for the first host interface port, 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

161

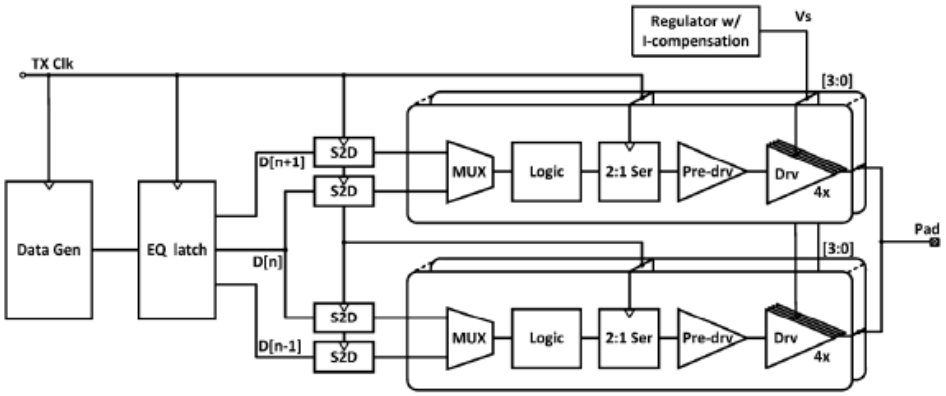
Claim 1	Disclosure
	lanes support asymmetric BW operation through the use of unequal transmit and receive rates or by configuring the lanes to transmit or receive. Link data rate scalability is achieved with minimal power overhead by operating in dual modes and by having multiple configurations for clocking, 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.

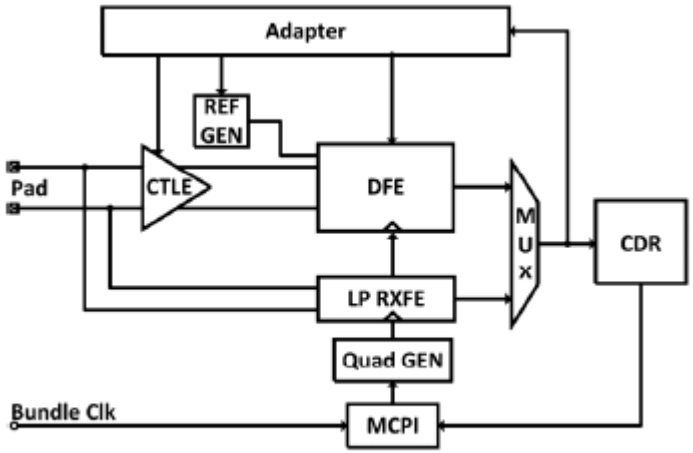
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

163

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.

164

Claim 1	Disclosure
165	 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
	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.

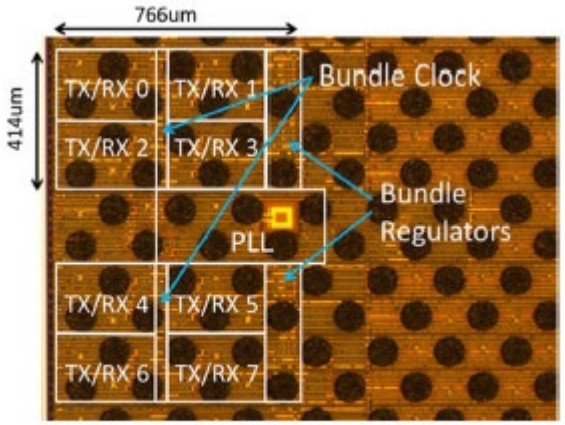
166

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.

167

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.

168

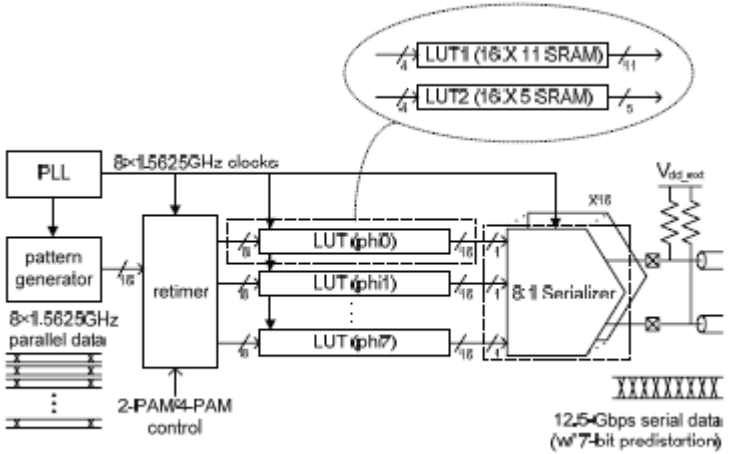
Claim 1	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses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

170

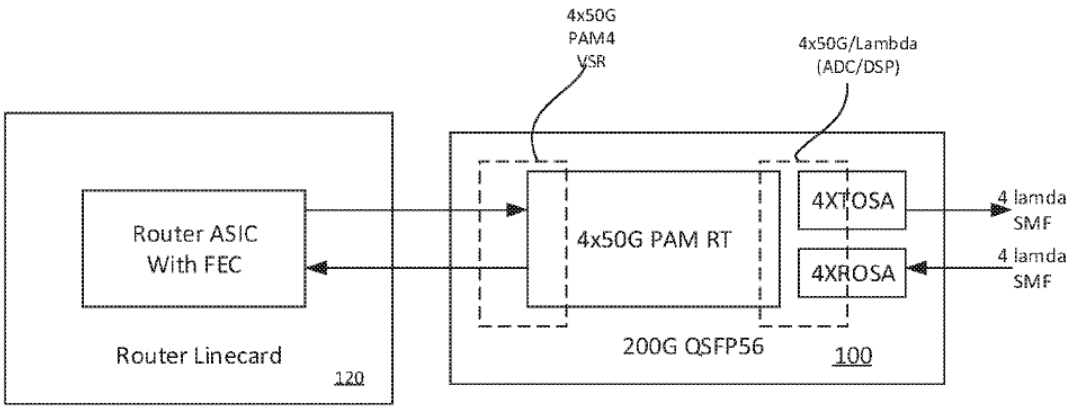
Claim 1	Disclosure
	This paper demonstrates a 12.5-Gbps transmit DAC with 7-bit resolution and 4-tap equalization using look-up tables (LUT). The LUT-based equalizer supports both linear and nonlinear forms of equalization and also offers flexibility to easily reprogram equalization tap settings [4,5]. 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.

171

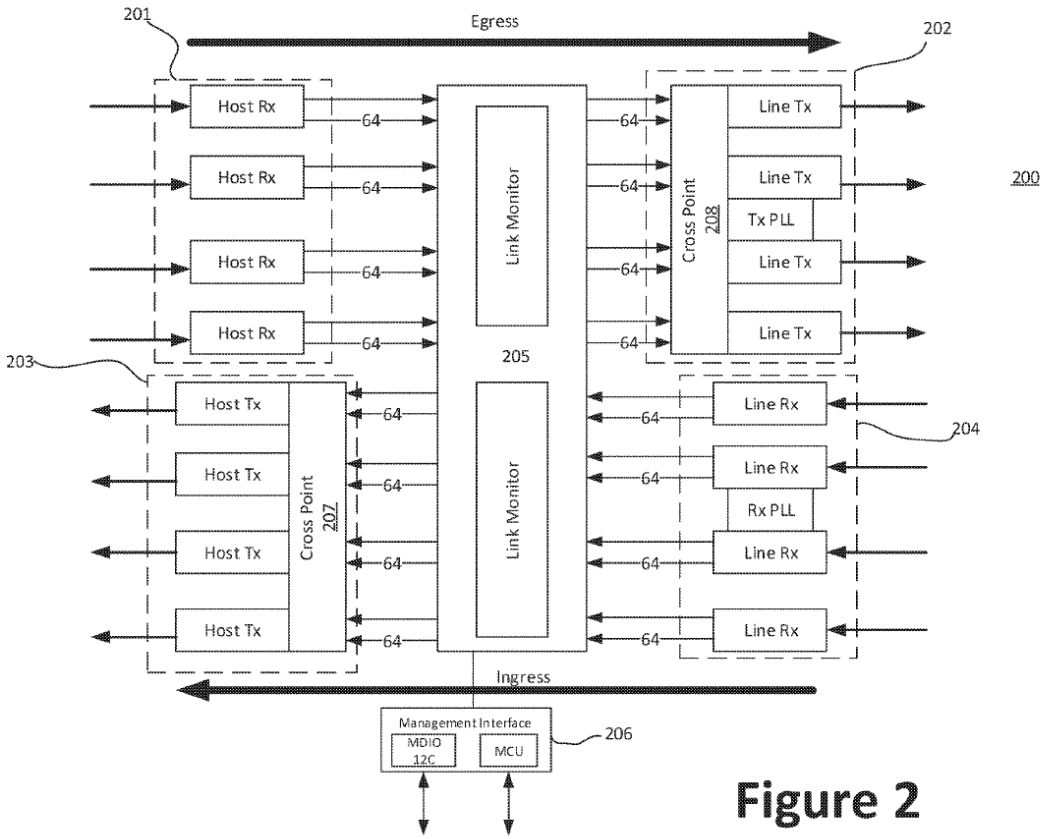
Claim 1	Disclosure
172	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

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.

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

175

Claim 1	Disclosure
176	 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.

177

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.

178

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.

179

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.

180

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.

181

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.

182

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.

183

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.

184

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.

185

	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.	

186

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.

187

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
	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.

189

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

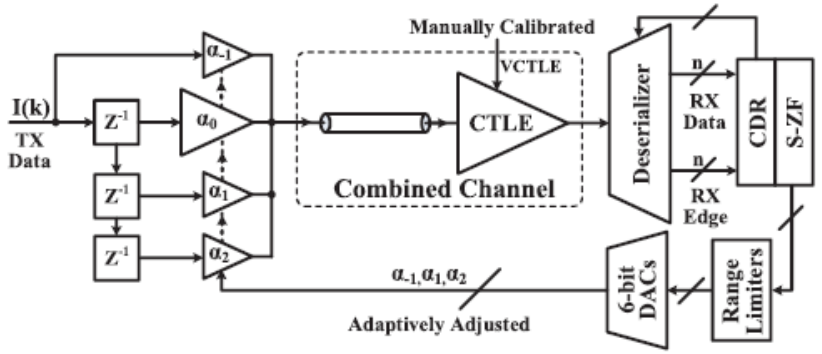
190

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.

191

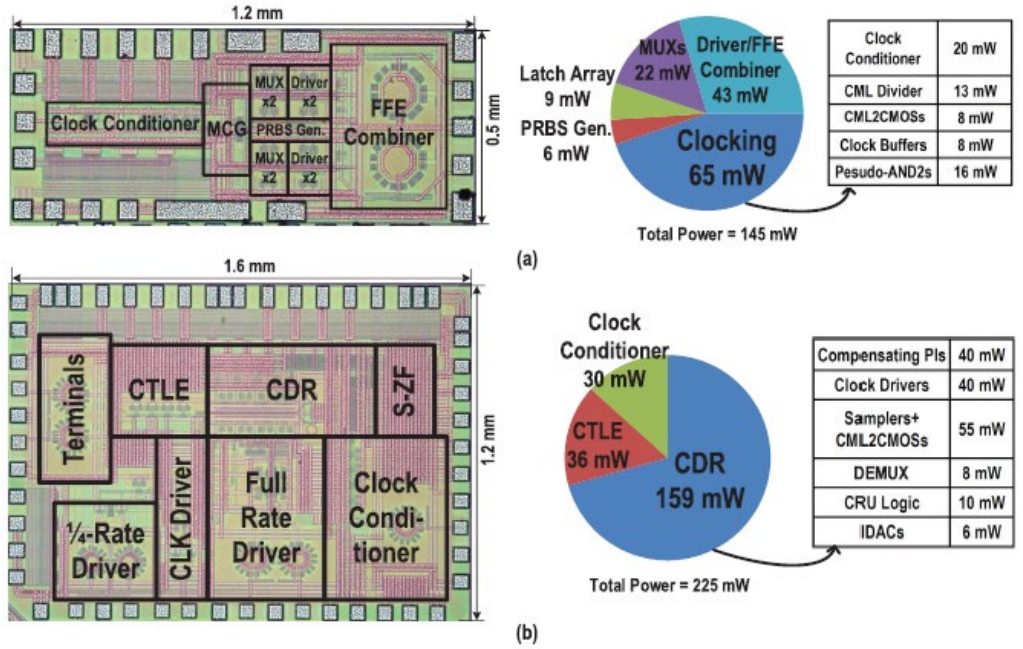
Claim 1	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear 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.

192

Claim 1	Disclosure
193	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.

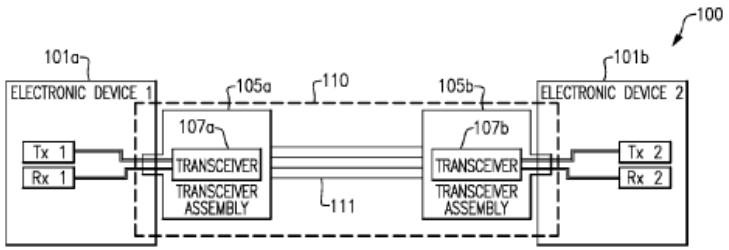
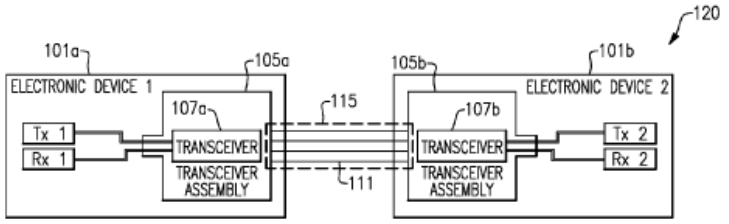
Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

194

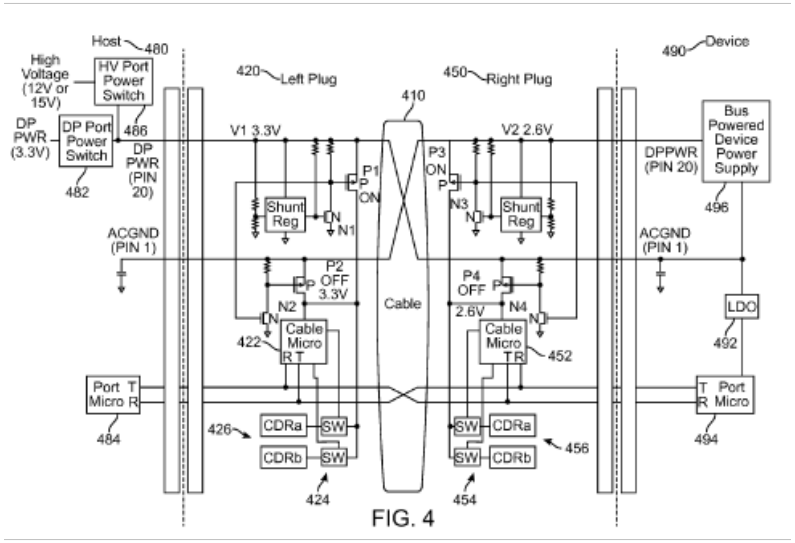
Claim 1	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

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.

196

Claim 1	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a,

197

Claim 1	Disclosure
	101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.  FIG. 4 Baker, Figure 4.

Claim 1	Disclosure
	Voltage V1 is received by low-drop-out regulator 492 in device 490, which provides power to port microcontroller 494. The host port microcontroller 484 may then communicate with cable microcontrollers s 422 and 452 and port microcontroller 494 to determine proper configuration for the cable. In a specific embodiment of the present invention, the host port microcontroller 484 may check with device port microcontroller 494 to determine whether device 490 requires a higher level of power. If it does, host port microcontroller 484 may check with cable microcontroller 422 to determine whether the cable can support the delivery of this higher level of power. If device 490 requires higher power, and the cable can deliver it, then host 480 may provide the higher level of power. In another embodiment of the present invention, host port microcontroller 484 may determine how much power will be needed by the cable and device 490. In some circumstances, one link including one pair of clock and data recovery circuits, or other circuitry, may need to be powered off. Baker, 5:49–67.

199

Claim 1	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. **[0032]** The receive section **122** includes a photo diode module **148**, a TIA ("TransImpedance Amplifier") module **146** and a LA/CDR ("Limiting Amplifier/Clock and Data Recovery") module **144**. The photo diode module **148** receives high speed optical data from fibre optic cable **126** and converts the optical data into high speed electrical data that is further converted from current signals into voltage signals by TIA module **146**. The high speed data from the TIA module **146** is processed by LA/CDR module **144** and then sent over host channel **103** as high speed serial data to a receiver portion of SERDES **110**. In embodiments having multiple channels in the form of multiple transmitting conductors **112**, the LA/CDR module **144** performs these operations for each channel of the conductors **112**. Tremblay, paragraphs 31-32.
1[e] the second DRR device converting between said electrical transit signals and said inbound and outbound multi-lane data streams	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
for the second host interface port,	***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

201

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.

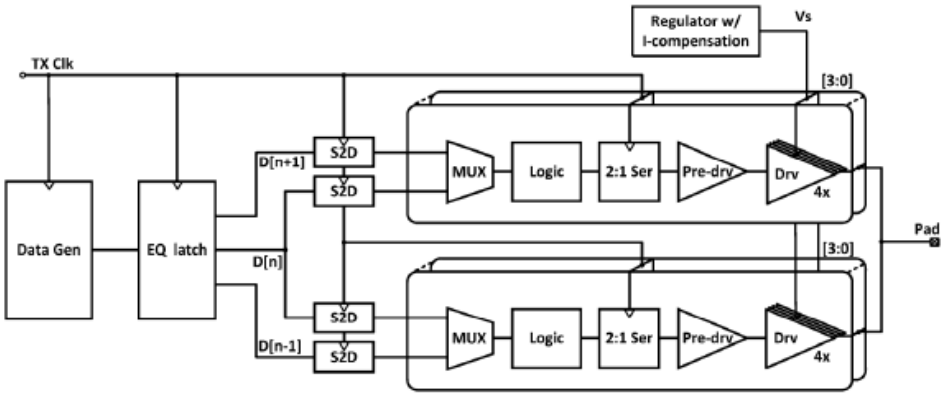
202

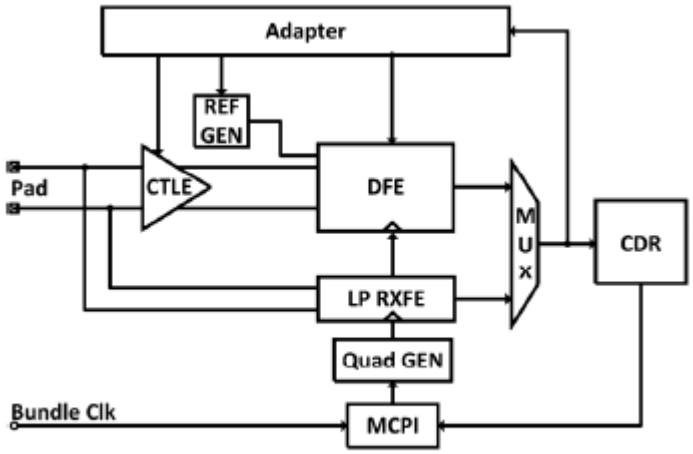
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

203

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.

204

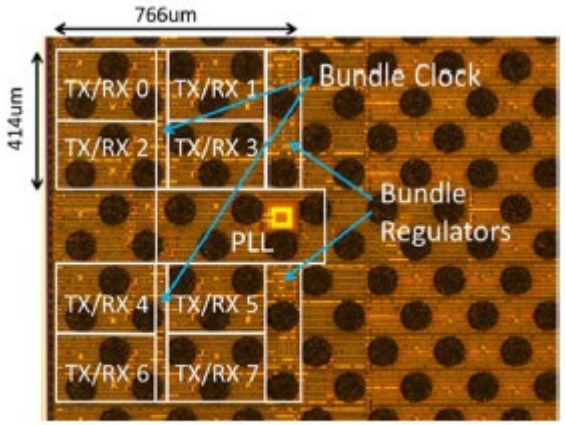
Claim 1	Disclosure
205	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is 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
206	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.

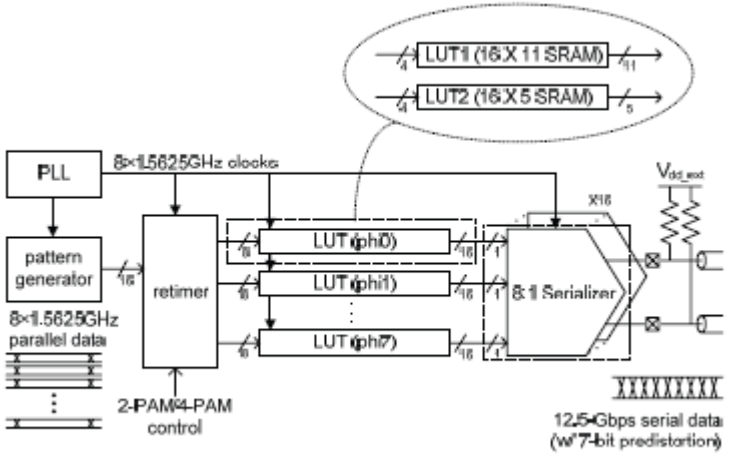
208

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.

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.

211

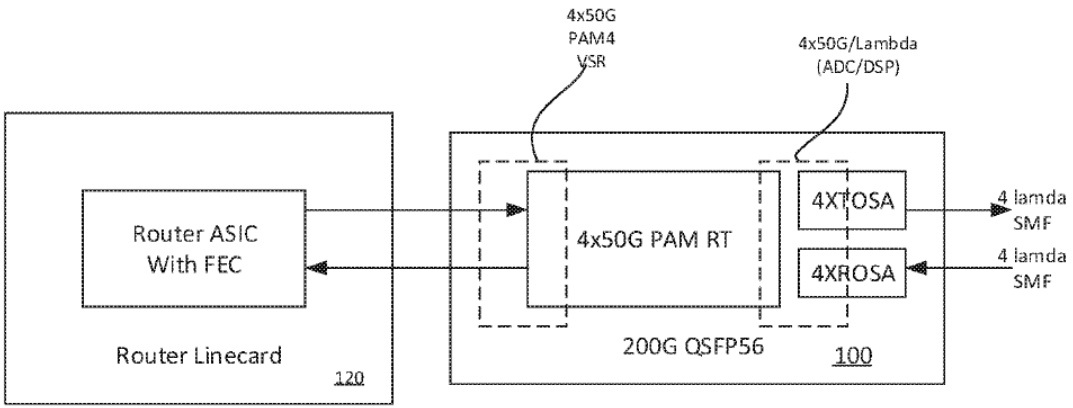
Claim 1	Disclosure
212	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

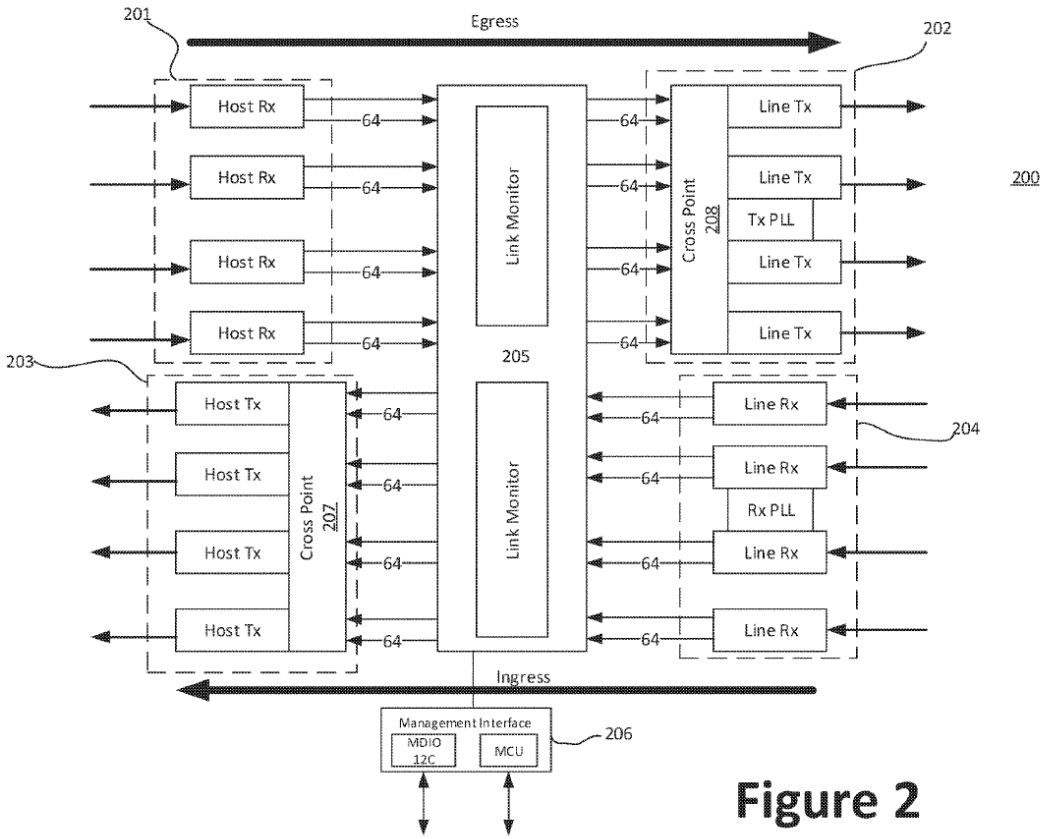
Claim 1	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

213

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.

214

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

Claim 1	Disclosure
216	 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.

217

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.

219

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.

220

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.

221

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.

222

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.

223

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.

224

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.

225

	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.	

226

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.

227

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.

228

Claim 1	Disclosure
229	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

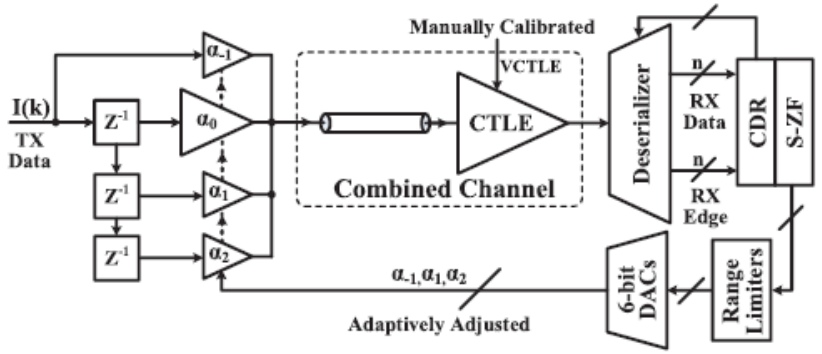
230

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.

231

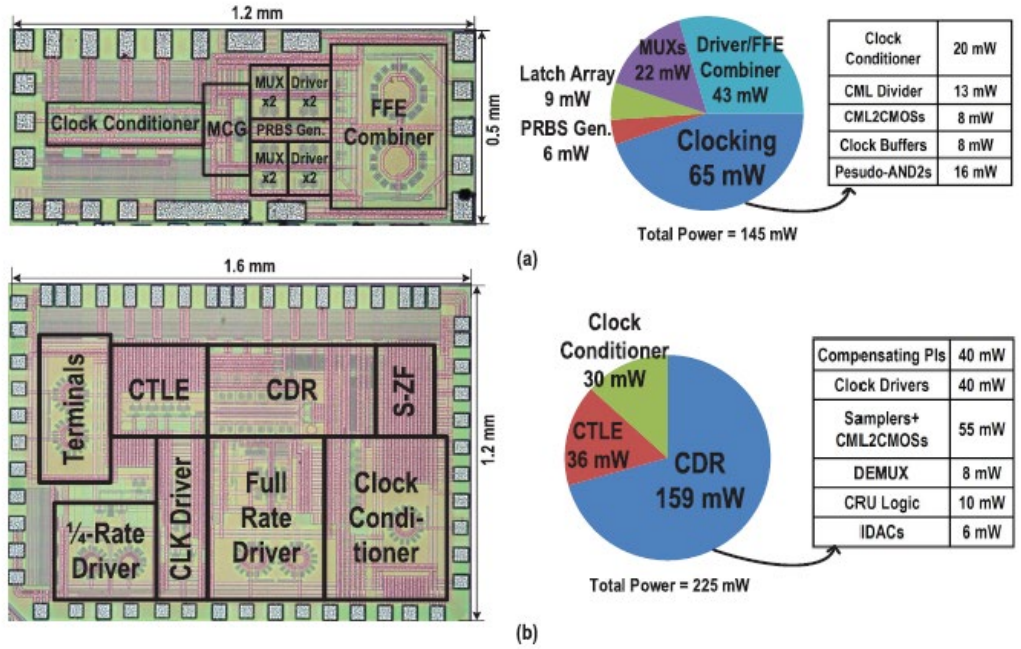
Claim 1	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear 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.

232

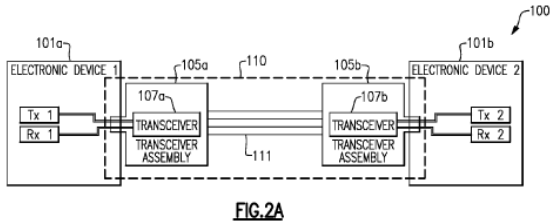
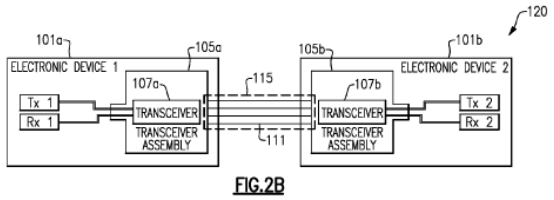
Claim 1	Disclosure
233	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based 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.

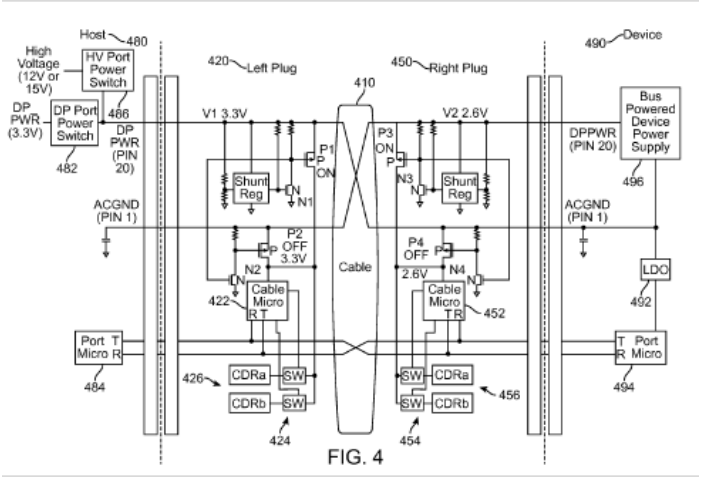
Claim 1	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

234

Claim 1	Disclosure
235	 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.

Claim 1	Disclosure
237	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP).

Claim 1	Disclosure
238	Lugthart, 15:8–23.  FIG. 4 Baker, Figure 4. Voltage V1 is received by low-drop-out regulator 492 in device 490, which provides power to port microcontroller 494. The host port microcontroller 484 may then communicate with cable microcontrollers s 422 and 452 and port microcontroller 494 to determine proper configuration for the cable. In a specific embodiment of the present invention, the host port microcontroller 484 may check with device port microcontroller 494 to determine whether device 490 requires a higher level of power. If it does, host port microcontroller 484 may check with cable microcontroller 422 to determine whether the

Claim 1	Disclosure
	cable can support the delivery of this higher level of power. If device 490 requires higher power, and the cable can deliver it, then host 480 may provide the higher level of power. In another embodiment of the present invention, host port microcontroller 484 may determine how much power will be needed by the cable and device 490. In some circumstances, one link including one pair of clock and data recovery circuits, or other circuitry, may need to be powered off. Baker, 5:49–67. **[0031]** In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. **[0032]** The receive section **122** includes a photo diode module **148**, a TIA ("TransImpedance Amplifier") module **146** and a LA/CDR ("Limiting Amplifier/Clock and Data Recovery") module **144**. The photo diode module **148** receives high speed optical data from fibre optic cable **126** and converts the optical data into high speed electrical data that is further converted from current signals into voltage signals by TIA module **146**. The high speed data from the TIA module **146** is processed by LA/CDR module **144** and then sent over host channel **103** as high speed serial data to a receiver portion of SERDES **110**. In embodiments having multiple channels in the form of multiple transmitting conductors **112**, the LA/CDR module **144** performs these operations for each channel of the conductors **112**.

Claim 1	Disclosure
	Tremblay, paragraphs 31-32.
1[f] the first and second DRR devices providing pre-equalization of the electrical transit signals using transmit filter coefficient values stored in nonvolatile memories.	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.

240

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.

241

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.

Claim 1	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.

243

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.

Claim 1	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.

245

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.

246

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.

247

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) 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 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

249

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.

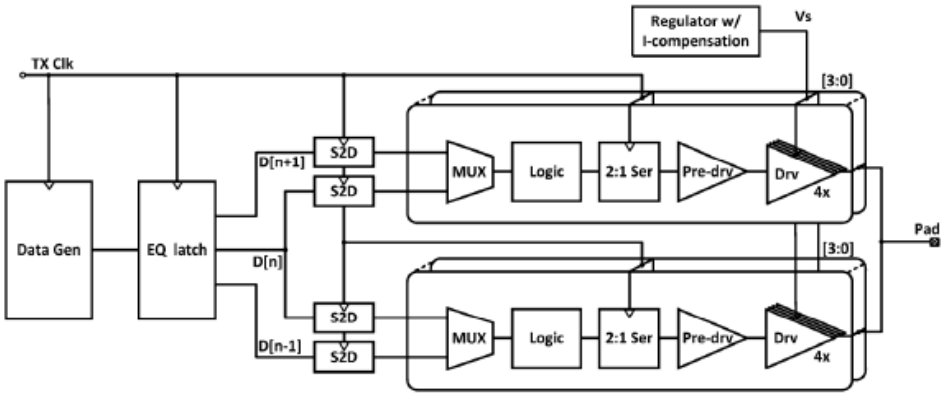
250

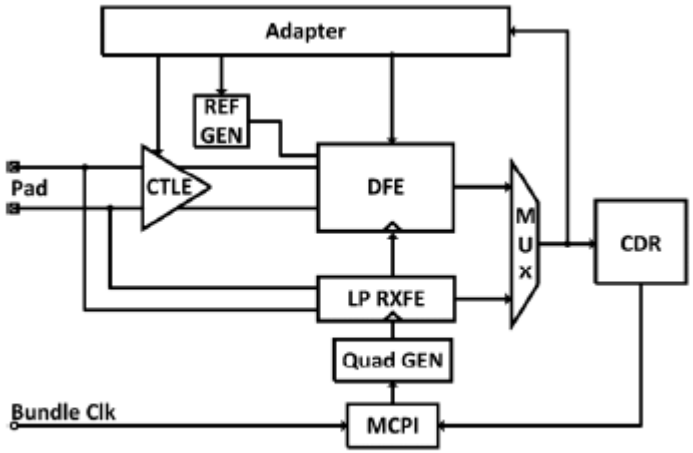
Claim 1	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

251

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.

252

Claim 1	Disclosure
253	 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

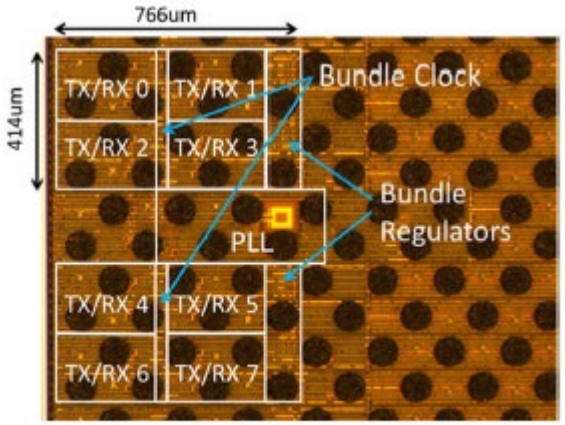
Claim 1	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

255

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.

256

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.

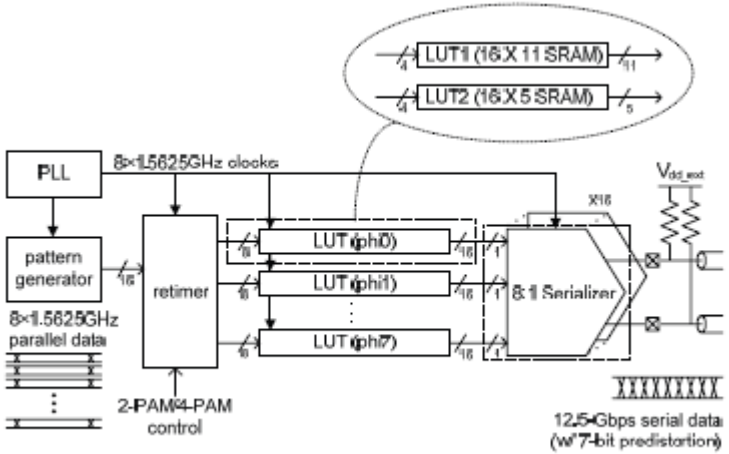
257

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.

258

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.

259

Claim 1	Disclosure
260	 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.

262

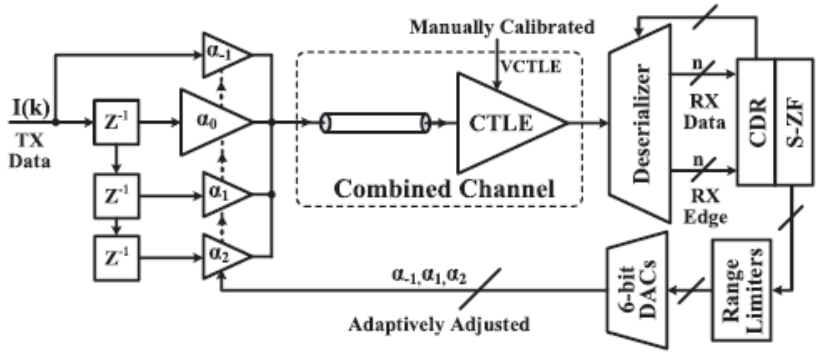
Claim 1	Disclosure
	II. TRANSMITTER CHIP *A. Overall Architecture* Fig. 1 shows the block diagram of the TX chip. It contains a multi-MUX-based four-tap FFE combiner, a latch array, an on-chip PRBS generator, and a clock bundle. The parallel quarter-rate data D0⟨n⟩, D1⟨n⟩, D2⟨n⟩, and D3⟨n⟩ are generated 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

263

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.

264

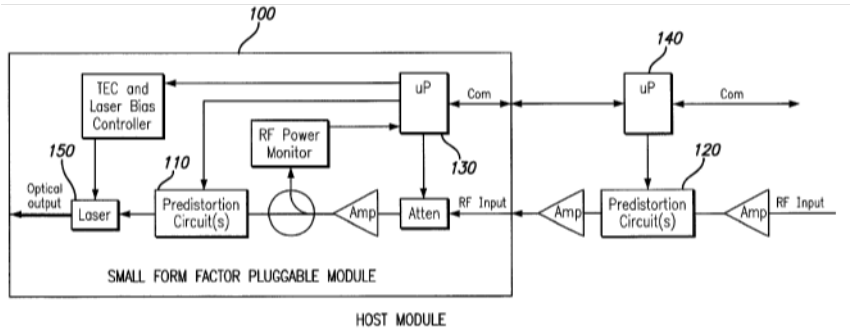
Claim 1	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time linear 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
266	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.

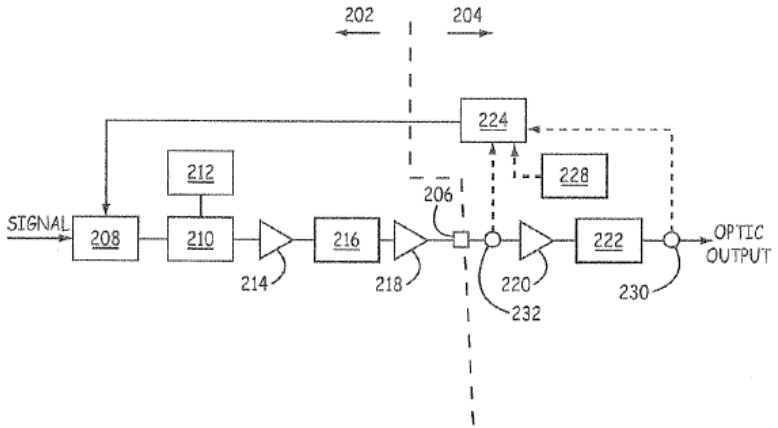
Claim 1	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

267

Claim 1	Disclosure
268	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. 7,233,617 to Gorecki* Gorecki discloses that "tap weights or coefficients (values) may be determined or controlled[.]" *See* Gorecki at 6:56-57. Gorecki discloses the storing of filter coefficient values in EEPROM, and use those values to set filter coefficients to provide pre-equalization. *See* Gorecki at 7:18-33 ("[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM. In this way, the transmitter may access the memory to retrieve the necessary information during start-up/power-up, initialization or re-initialization."). An EEPROM [electrically erasable programmable read-only memory] is non-volatile memory. *See, e.g.*, U.S. Patent App. Pub. No. 2015/0106536 at [0081] ("nonvolatile memory can include… electrically erasable PROM (EEPROM)"); U.S. Patent No. 8,000,176 at 2:53-54 ("EEPROM is a non-volatile memory"); U.S. Patent App. Pub. No. 2005/0013317 at [0012] ("non-

Claim 1	Disclosure
	volatile memory may include… an EEPROM"). The L-G combination teaches programming the coefficients of the FIR filter 17 using the values stored in non-volatile EEPROM to pre-equalize the TX signal in transceiver 107a and 107b.  FIG. 1 Ishaug, Figure 1.

Claim 1	Disclosure
	A method includes reading operation parameters from a non-volatile memory located in a pluggable module that is coupled to a host module; processing the operational parameters with a processor located in the host module to control operation of a predistortion circuit located in the host module; adding predistortion to a signal with the predistortion circuit located in the host module and then sending the predistorted signal to the pluggable module. An apparatus includes a host module including a predistortion circuit and a processor coupled to the predistortion circuit; and a pluggable module coupled to the host module, wherein the pluggable module includes a non-volatile memory containing operational parameters for the predistortion circuit of the host module, wherein the operational parameters are processed by the processor of the host module to control the predistortion circuit of the host module. Ishaug, Abstract. As shown in FIG. 1, both the pluggable module and the host module have predistortion circuits 110, 120. However, the pluggable module does not need to contain a predistortion circuit. The pluggable module includes a processor 130 having non-volatile memory. In addition, the host module is shown as a standalone unit including a processor 140 having non-volatile memory. However, the host may be incorporated into a larger system that provides the analog RF input. There is also gain, attenuation and RF power monitoring shown in the pluggable and host modules that are not required, but may be needed to set the drive level to the laser modulation at an appropriate level. Finally, the optical transmitter is shown as a laser 150. The laser may be a directly modulated laser or an externally modulated laser in which the output is modulated using techniques such as electro-absorption modulation, Mach-Zender modulation or any appropriate technique for producing analog modulated optical signal. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other “smart” equipment to

Claim 1	Disclosure
271	determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:34–5:4.  FIG. 3A Piehler, Figure 3A.

Claim 1	Disclosure
	Linearized optical transmitter units are described for a hybrid optical fiber coaxial cable network. The linearized optical transmitter unit can comprise a directly-modulated or externally-modulated laser optically coupled to an optical conduit directed to an optical fiber communications link and electrically coupled to an electrical RF source line that provides an RF source to drive the laser or an external modulator for a light beam from the laser. A linearization information electrical component comprising memory and/or a processor, and a data output configured to transmit linearization enabling data for input into a direct digital synthesis engine that enables the direct digital synthesis engine to generate an RF signal wherein nonlinear responses of the transmitter and/or the optical fiber communications link are pre-compensated, in which the data is specific for the optical transmitter and/or the optical fiber communications link. An electronic communication channel can connected to the data output of the electrical component to communicate the linearization enabling data external to the linearized optical transmitter unit suitable for communication to the direct digital synthesis engine. The linearized optical transmitter can be assembled in a plug-in module. Corresponding HFC systems and methods are also described. Piehler, Abstract. FIG. 3A shows edge QAM 202 and optical transmitter 204 connected by interface 206, e.g., a 25 ohm RF interface. While the division between edge-QAM 202 and optical transmitter 204 can be considered arbitrary, the division provides a convenient design framework since the edge-QAM 202 provides electrical RF signals based on user specific content and optical transmitter 204 provides laser module specific performance. The interface involves transfer of appropriate signals to provide the respective functionalities including, for example, pre-compensation for optical transmitter and fiber nonlinearities for efficient production of optical signals with low distortion. As described below, the division of the optical transmitter provides for the formation of separate modules, which can be plug-in components. A digital switch or other device passes a signal to edge QAM 202 to provide user specific digital content.

Claim 1	Disclosure
	Piehler, 9:38–53. In additional or alternative embodiments, electrical component 224 may comprise a microprocessor and further receive signal information that is used in creation of the pre-distortion parameters for dynamic correction of nonlinearities. As shown in FIG. 3A, electrical component 224 can receive optionally signals from sensor 228 and/or from an optical coupler or tap 230 and RF coupler or tap 232. Sensor 228 can measure, for example, laser temperature, laser bias voltage and/or other local conditions that can influence optical transmitter nonlinearities. Optical coupler 230 and RF coupler may be integrated into the system at a position downstream and upstream, respectively, of laser driver/laser assembly 222. One or both of the couplers/tap may be used to provide signal information, as indicated by dashed lines connecting the couplers and microprocessor. If both couplers 230, 232 are used, the microprocessor may compare the upstream and downstream signal information to isolate distortion created by laser driver/laser assembly 222 and RF amplifier 220, and use the comparison data to create data related to the pre-distortion correction that is provided to processing unit 208. The correction may include predetermined parameters specific to laser modulator 222 or be based only on the comparison data. Electrical component 224 may be an integrated circuit or other microprocessor programmed for active processing and/or supply of data or a passive device, e.g., random access memory, non-volatile memory, or other memory. Information in electrical component 224 may comprise parameters for the specific laser module in the device. For example, the laser may be tested at the point of manufacturing or at other times to measure the distortion that it creates. The distortion has consistent components that are specific to the laser; these components may be analyzed and appropriate compensatory corrections created that are placed in, or otherwise introduced by, the electrical component. The pre-distortion correction provided to the DDS may comprise initial values for pre-distortion based on the laser's intrinsic nonlinearity at normal operating conditions as well as other nonlinearities which may exist in the system. The pre-distortion corrections may thus compensate for the net nonlinearity of the optical transmitter which includes, the laser modulator, other optical components, and other RF components including RF amplifiers. Examples of corrections are compensations for one or more of second, third, and higher nth order distortions in the optical transmitter.

Claim 1	Disclosure
	Piehler, 10:35–11:12. "The filter coefficients are stored in memory (**115**) for future use. In one embodiment, the coefficients are stored in non-volatile memory such as some type of programmable read only memory. However, the present invention is not limited to any one type of memory. The memory may be semiconductor, magnetic, or optical." Unger at 2:50-55. "In one embodiment, these coefficients are read from memory (**206**) where they were stored by the method of FIG. **1**. The memory may be incorporated into the filter, part of the same circuit card as the modulator and upconverter, or it can be located in another part of the system and coupled to the predistortion filter through a data bus." Unger at 3:35-41, *see also id.* at Figs. 1-2, 4. "In another embodiment, the media access controller is responsible for loading the filter coefficients that are stored in memory (**403**) into the predistortion filter (**203**). The embodiment of FIG. 4 shows the memory (**403**) as being a separate block. One embodiment incorporates the memory into the media access controller (**401**). Another embodiment incorporates the memory into the predistortion filter (**403**)." Unger at 4:47-53. **[0038]** With reference to the drawings, FIG. **2** is a flow chart of the operation of an example embodiment, wherein the LA/CDR unit **144** in conjunction with the microcontroller

274

Claim 1	Disclosure
	150 performs pre-equalization and optimizes power consumption of the transmission over host channel **103** via physical conductors **114** based on the feedback received from the SERDES module **110** via the low-speed digital interface **152**. In FIG. **2**, the operational flow of the LA/CDR unit **144** and the microcontroller **150** begins at step **202** when the microcontroller receives feedback via the low speed digital interface **152**. At step **204**, the microcontroller **150** instructs the LA/CDR unit **144** to equalize its transmitted data signal in response to the feedback, applying sufficient transmit pre-emphasis and other pre-equalization effects to the transmitted data signal to achieve predetermined target eye characteristics at the SERDES module **110**. Once the feedback from the SERDES module **110** indicates that the predetermined target eye characteristics have been achieved, the microcontroller **150** at step **206** instructs the LA/CDR unit **144** to begin reducing the launch amplitude of the transmitted data signal as much as it can without letting the received signal at the SERDES module **110** fall short of the predetermined target eye characteristics. The process is iterative, with feedback being monitored at step **202** and equalization being adjusted at step **204** while the launch amplitude is being reduced at step **206**. In alternative embodiments, the process may operate according to a different sequence of steps while achieving the same result, namely the optimization of signal quality and power consumption over the host channel **103**. Tremblay, paragraph 38.

275

B. DEPENDENT CLAIM 2

Claim 2	Disclosure
2[a] The cable of claim 1, further comprising a first controller that configures the first DRR device in response to a power-on event,	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.

276

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.

277

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.

278

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.

279

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.

281

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.

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

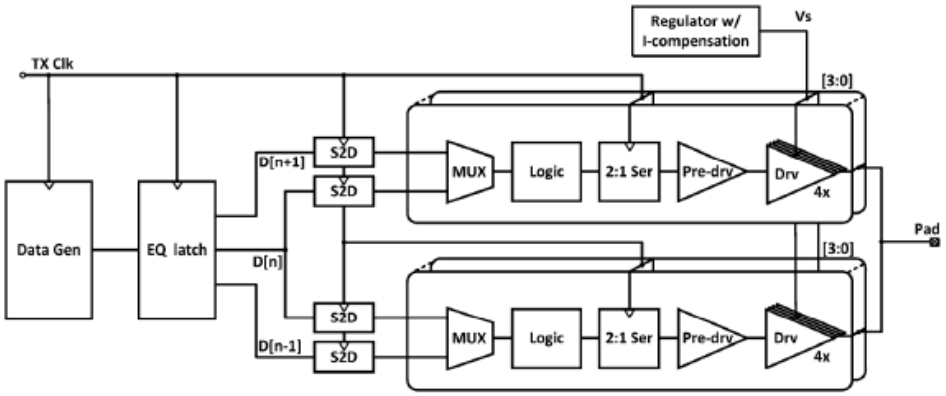
285

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.

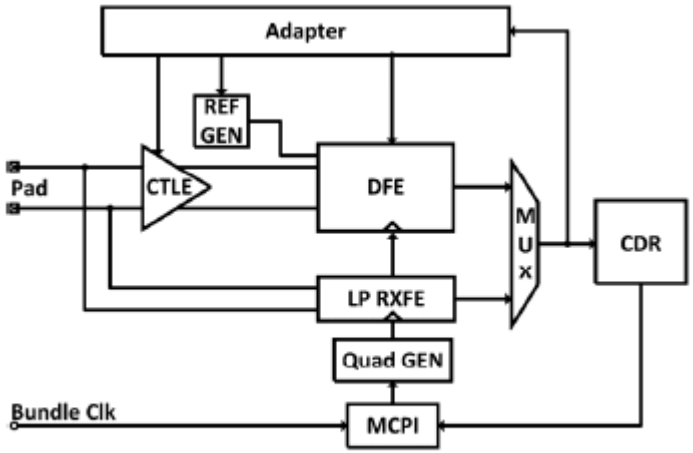
Claim 2	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

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.

288

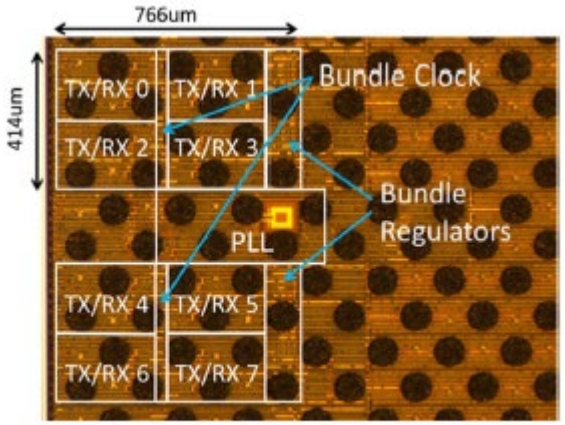
Claim 2	Disclosure
	 Fig. 6. Half-rate transmitter architecture. Musah, 3082. The RX block diagram is shown in Fig. 7. The RXFE is configured to operate in HP and LP modes. The CDR, MCPI and quadrature generator are shared in both modes. A half-rate front-end with 2× oversampling phase detection is used 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

289

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.

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.

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.

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. 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

294

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.

296

Claim 2	Disclosure

297

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.

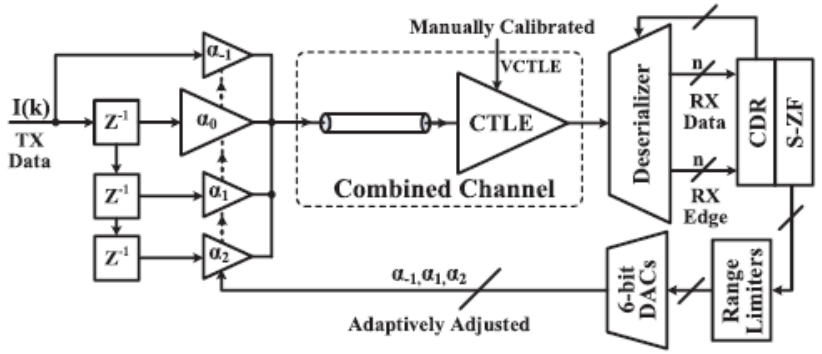
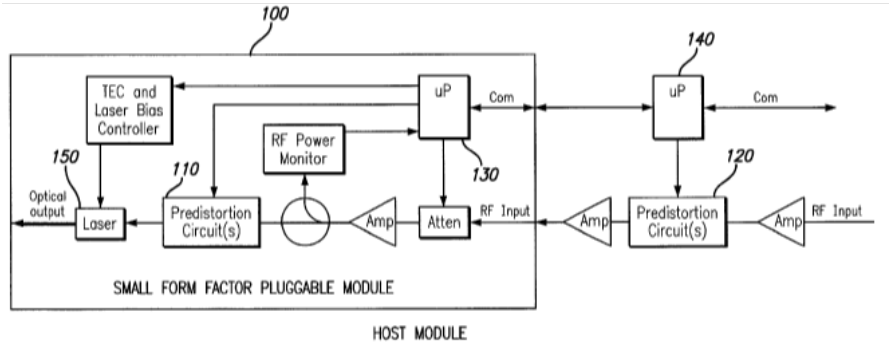


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 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.

298

Claim 2	Disclosure
299	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. 7,233,617 to Gorecki* Gorecki teaches "[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM . In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up," which is a power-on event. *See* Gorecki at 7:18-33, 16:47-53. In addition, Gorecki describes a "controller" that "distribute[s] to the transmitter" the filter coefficients thereby configuring a filter for pre-emphasis at this power-on event. *See* Gorecki at 7:4-17; 7:26-33.

Claim 2	Disclosure
300	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for

Claim 2	Disclosure
	each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4. **[0033]** In example embodiments a low speed digital connection or interface **152** is provided between host SERDES **110** and the interface module **100** to allow back channel communications between transmitting elements of the SERDES module **110** and corresponding receiving elements of the interface module **100** and between transmitting elements of the interface module **100** and receiving components of the SERDES Module **110**. In some implementations, the interface module **100** includes a uC (microcontroller) module **150** that facilitates communications over the low speed digital interface **152**. By way of non-limiting example, the low speed digital interface **152** could include an I2C ("Inter-Integrated Circuit") interface, or SPI ("Serial Peripheral Interface"), MDIO ("Management Data Input/Output") interface, or similar interface. In some example embodiments, the low speed data connection **152** is used for data rates under 10 MB/s. Tremblay, paragraph 33.
2[b] the first controller retrieving the transmit filter coefficient values	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:

301

Claim 2	Disclosure
from the nonvolatile memories as part of said configuring.	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.

302

Claim 2	Disclosure
303	**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 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.

304

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.

305

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.

306

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.

307

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.

308

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.

309

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

310

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.

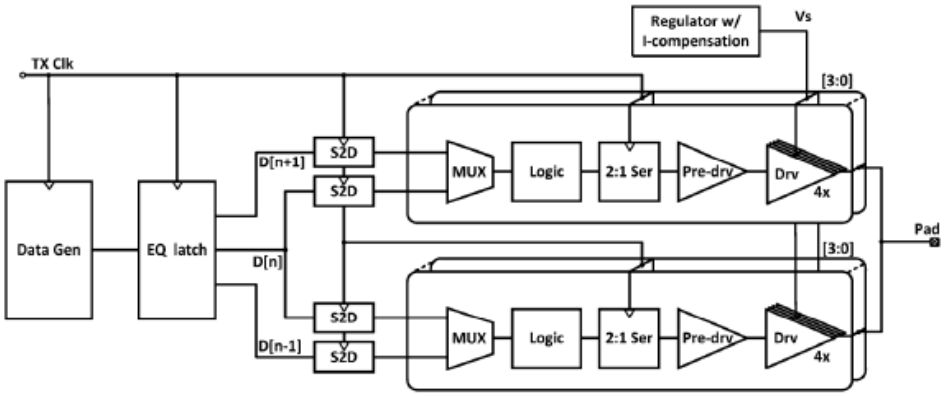
311

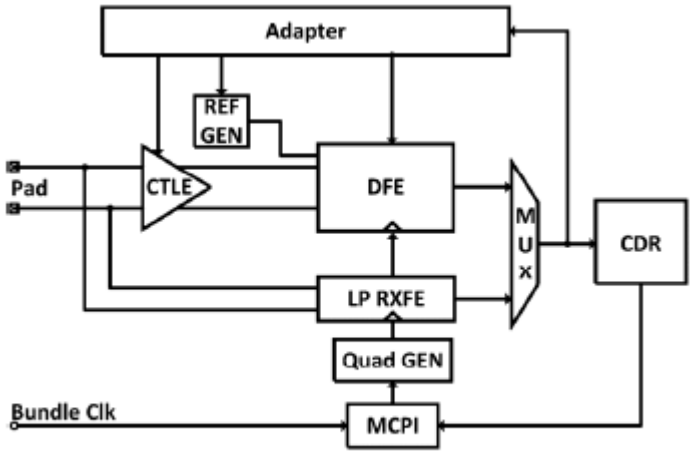
Claim 2	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

312

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.

313

Claim 2	Disclosure
314	 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

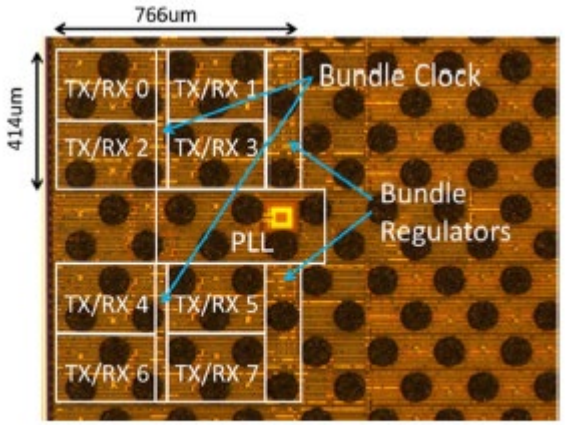
Claim 2	Disclosure
315	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

316

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.

317

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. 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

319

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.

320

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.

321

Claim 2	Disclosure

322

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.

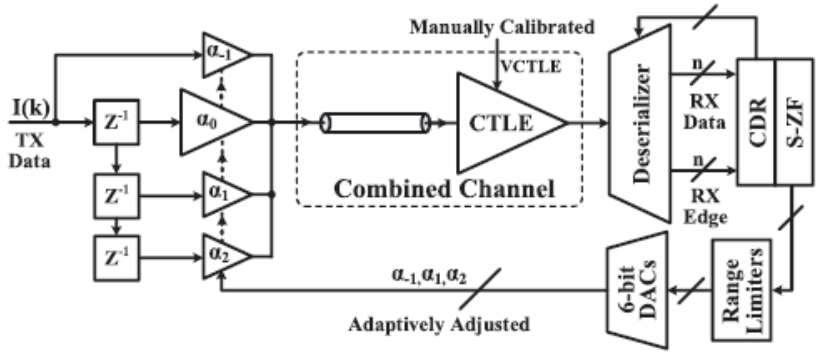
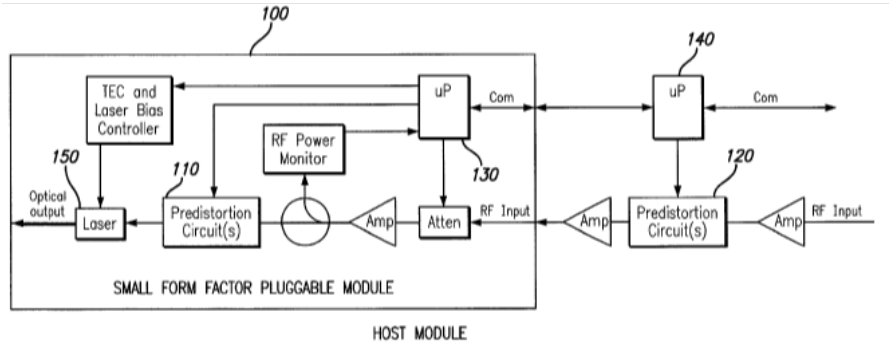


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 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.

323

Claim 2	Disclosure
324	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. 7,233,617 to Gorecki* Gorecki discloses a "controller" that "distribute[s] to the transmitter" the filter coefficients thereby configuring a filter for pre-emphasis at this power-on event. *See* Gorecki at 7:4-17; 7:26-33.

Claim 2	Disclosure
325	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for

Claim 2	Disclosure
	each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4. **[0033]** In example embodiments a low speed digital connection or interface **152** is provided between host SERDES **110** and the interface module **100** to allow back channel communications between transmitting elements of the SERDES module **110** and corresponding receiving elements of the interface module **100** and between transmitting elements of the interface module **100** and receiving components of the SERDES Module **110**. In some implementations, the interface module **100** includes a uC (microcontroller) module **150** that facilitates communications over the low speed digital interface **152**. By way of non-limiting example, the low speed digital interface **152** could include an I2C ("Inter-Integrated Circuit") interface, or SPI ("Serial Peripheral Interface"), MDIO ("Management Data Input/Output") interface, or similar interface. In some example embodiments, the low speed data connection **152** is used for data rates under 10 MB/s. Tremblay, paragraph 33.

C. DEPENDENT CLAIM 3

Claim 3	Disclosure
3. The cable of claim 1, wherein the first and second DRR devices are programmed to use the transmit coefficient values each time power is supplied to the first and second end connector plugs.	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.

327

Claim 3	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.

328

Claim 3	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.

329

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.

330

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.

331

Claim 3	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.

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.

333

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.

334

Claim 3	Disclosure
335	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 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 mm².** Musah, Abstract. The link architecture block diagram is shown in Fig. 2, which includes 8 bidirectional transceiver lanes [5]. The bidirectional

336

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.

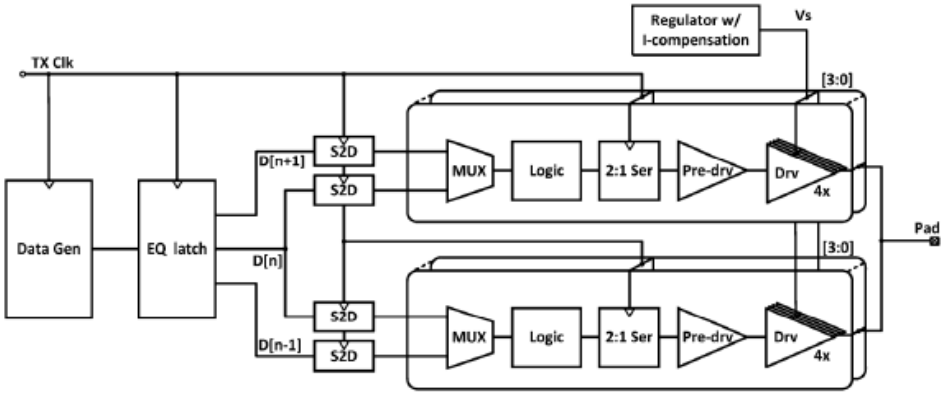
337

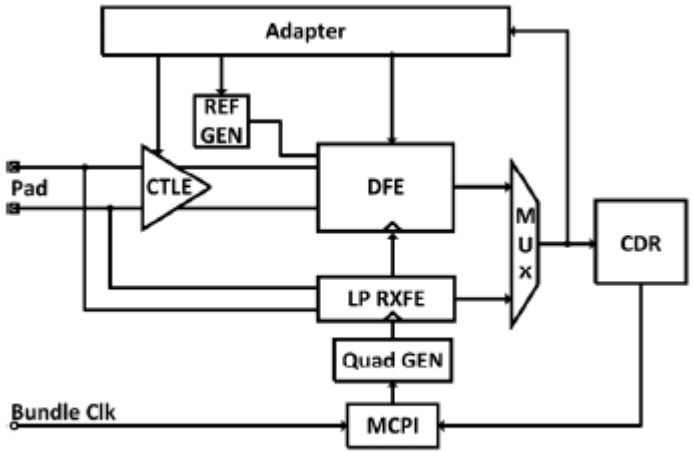
Claim 3	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

338

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.

339

Claim 3	Disclosure
340	 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

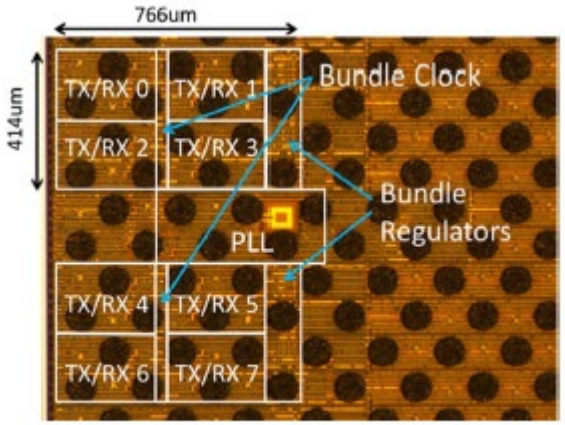
Claim 3	Disclosure
341	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

342

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.

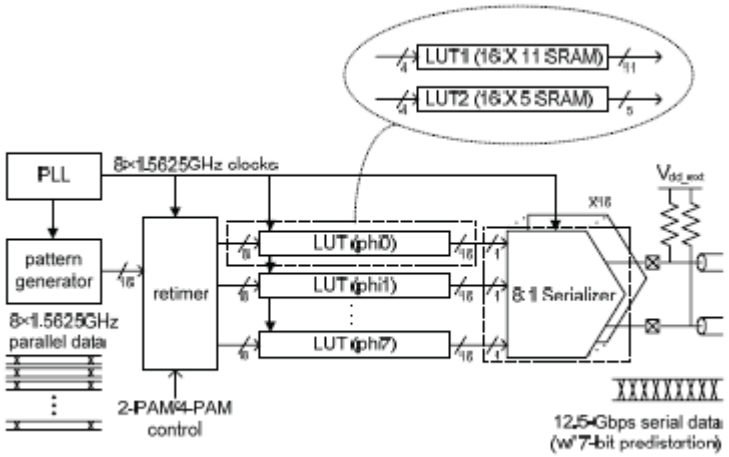
343

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.

345

Claim 3	Disclosure
346	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 3	Disclosure
347	 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.

349

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

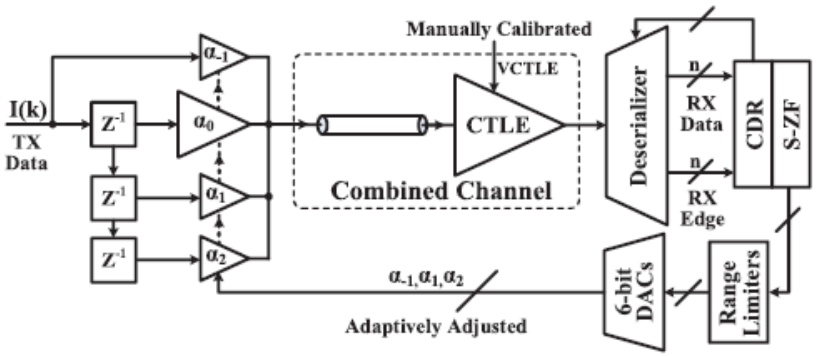
350

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.

351

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.

352

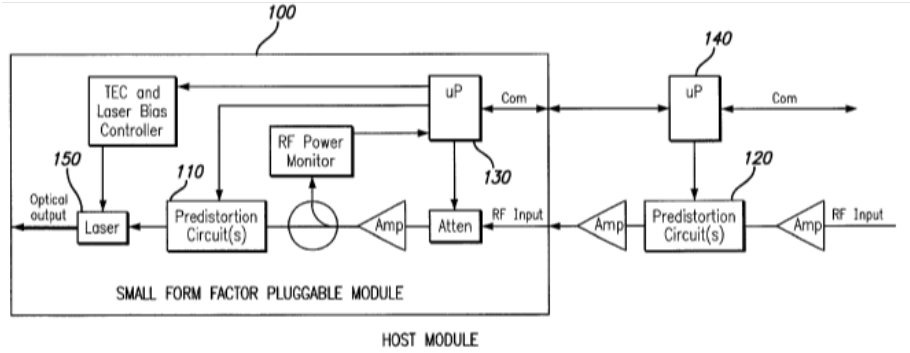
Claim 3	Disclosure
353	IV. CHANNEL EQUALIZATION The equalization scheme consisting of a TX-FFE and an RX-CTLE is utilized to compensate for the channel loss. As shown in Fig. 17, the RX-CTLE is manually calibrated while the tap weights of the TX-FFE are adaptively adjusted by an edge-data correlation-based 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.

354

Claim 3	Disclosure
355	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. 7,233,617 to Gorecki* For the limited purpose of this analysis, the phrase "transmit coefficient values" is interpreted as referring to "transmit filter coefficient values," (*see, e.g.*, claim limitation 1[f]) as the term otherwise has no antecedent basis. Like the '233 Patent's specification, Gorecki discloses storing of coefficient values in memory as programming. *See* Gorecki at 7:1-33 ("the adjustment or control of the taps of the equalization circuitry may be by the user via programming… coefficients of the tap(s), and/or the pulse durations of the tap(s) may be pre-programmed"); *see also*'233 Patent at 11:3-6 ("the automated tester equipment 'burns' the flash memory or otherwise programs the initial default values of the filter coefficients[.]").

Claim 3	Disclosure
	Gorecki teaches that the programmed memory "may reside on the integrated circuit containing the transmitter, receiver or transceiver." *See* Gorecki at 7:43-49. This programming is done to use the transmit coefficient values when power is supplied to the first and second end connector plugs. Gorecki discloses that stored filter coefficients are retrieved from the non-volatile memory at start-up/power-up. *See*, *e.g.*, Gorecki at 7:18-33; *see also* claim 2. Power-up indicates that power is supplied to the first and second end connector plugs. An ordinarily skilled artisan would have understood that power is supplied from a host device when the enclosing connector's connector plug was electrically connected to the host interface port in a corresponding electronic device 101a/b as described claim limitation 1[a] because this was conventionally how standardized connectors like QSFP or SFP worked . *See, e.g.*, SFF-8436 at 21 ("A host board together with the QSFP+ module(s) forms an integrated power system. The host supplies stable power to the module."), 22 Fig. 4 ("Recommended Host Board Power Supply Filtering"). Regardless, an ordinarily skilled artisan would have found it conventional and obvious to implement the QSFP connectors powered by a connected host device. An ordinarily skilled artisan would have understood that Gorecki uses the transmit coefficient values stored in memory each time power is supplied to the first and second end connector plugs, and teaches that the programmed values can be stored "permanently, semi-permanently or temporarily[.]" *See* Gorecki at 7:21. When stored permanently, an ordinarily skilled artisan would have understood them to be used each time power is supplied to the first and second end connector plugs.

Claim 3	Disclosure
	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for

357

Claim 3	Disclosure
	each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4: 51–5:4.

D. DEPENDENT CLAIM 4

Claim 4	Disclosure
4. The cable of claim 3, wherein the transmit filter coefficient values are determined and stored in the nonvolatile memories after assembly of the 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.

359

Claim 4	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.

360

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.

361

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.

362

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.

363

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.

364

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.

365

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.

366

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.

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

368

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.

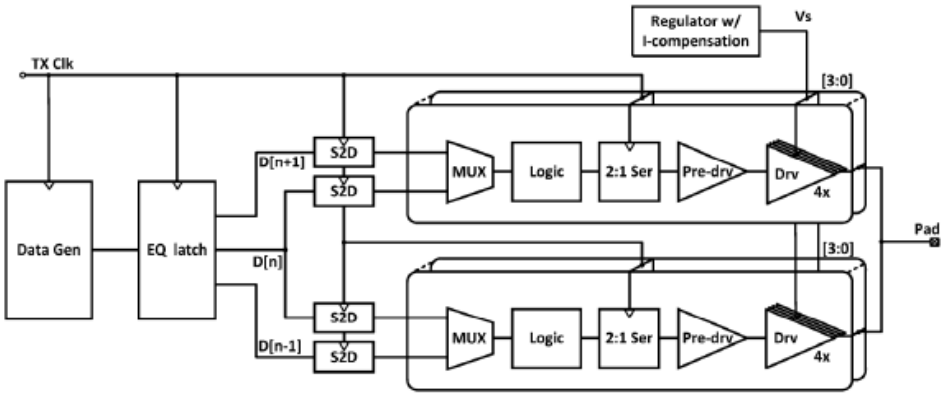
369

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.

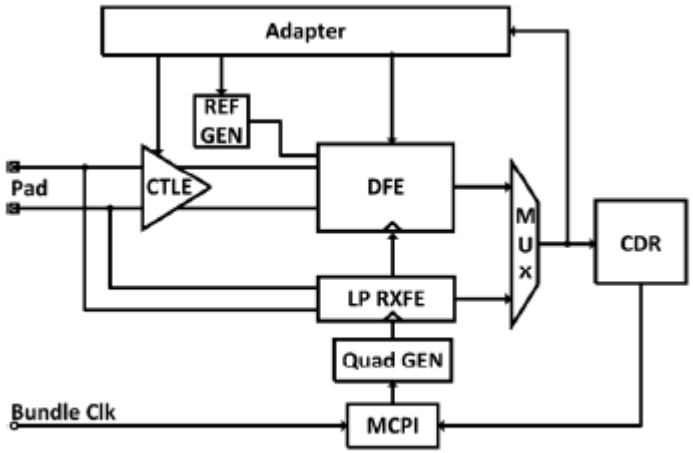
370

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.

371

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

372

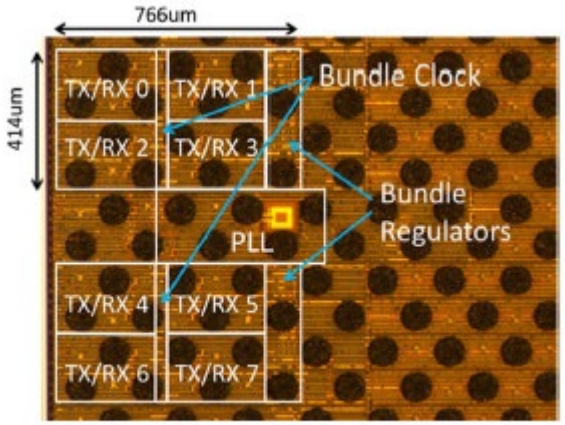
Claim 4	Disclosure
373	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

374

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.

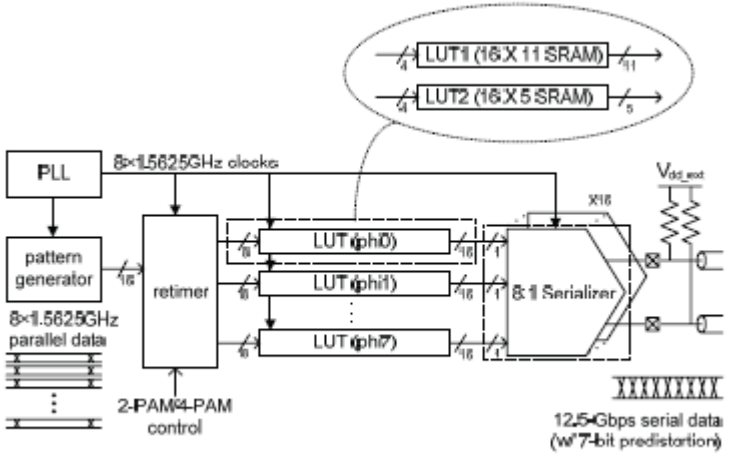
375

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.

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.

377

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.

Claim 4	Disclosure
379	 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

380

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.

381

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

382

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.

383

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.

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.

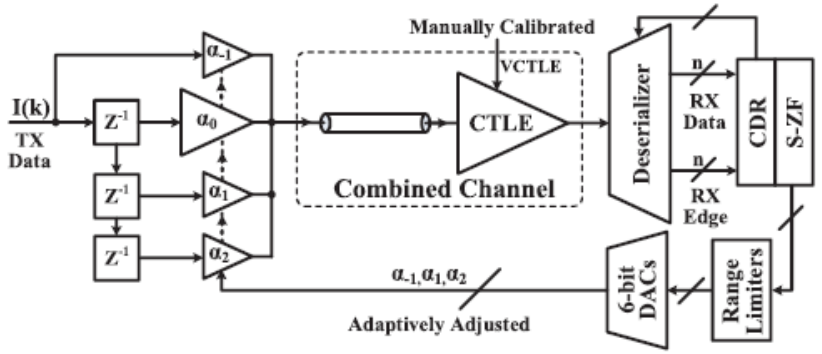


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

385

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.

386

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. 7,233,617 to Gorecki* Gorecki discloses determining and then storing the transmit filter coefficient values in the non-volatile memories after assembly of the cable. The values should be selected to offset effects of a communication channel. *See* Gorecki at 9:8-10 (The equalizers are "configured to introduce compensation that is ideally the inverse of the effects caused by the communications channel."). An ordinarily skilled artisan would have found it obvious to determine and then store filter coefficients, such as by measuring the effects in that communication channel after assembly. Such an approach was known in the art . *See, e.g.*, U.S. Patent App. Pub. No. 2014/0119425 at [0033] ("The programming is preferably performed with the cable installed between the actual source and actual sink it will be connecting."). Gorecki describes "the coefficients of the tap(s)… are adjusted or controlled during operation of the transmitter, for example using an adaptive algorithm." *See* Gorecki at 10:27-51, 6:50-53, 17:41-51.

387

Claim 4	Disclosure
	Determining values "during operation" happens after cable assembly. The adaptive algorithm—which determines transmit filter coefficient values—is performed "after (or during) the performance of an initialization or re-initialization process." *See* Gorecki at 7:50-61. Gorecki discloses that stored coefficient values can be "fine-tuned to enhance the system performance" (Gorecki at 7:51-53), and that the values can be "re-programmed" and, for that reason, are stored in writable non-volatile memory. *See, e.g.*, Gorecki at 7:18-33, 10:37-51. An ordinarily skilled artisan would have had reasons to use NVM including writable NVM ("EEPROM or the like"). *See* Gorecki at 7:24-25, 10:46-47. EEPROM is electrically erasable programmable read only memory, e.g., rewritable NVM. An ordinarily skilled artisan would have known that flash memory was another well-known rewritable non-volatile memory (e.g., "EEPROM or the like") conventionally used with transmit FIR filters. *See, e.g.*, U.S. Patent No. 6,975,140 at 3:42-4:2 ("Any type of NVRAM can be used as the rewritable non-volatile storage, so long as [it] is of the type suitable for integration on an integrated circuit."). An ordinarily skilled artisan would have understood such memories would be used when the content of those memories is both re-programmed and preserved through "startup/power-up, initialization or re-initialization," as described in Gorecki at 7:18-33. Goreki discloses rewritable non-volatile memory that could be updated with new coefficient values that fine-tuned pre-equalization or equalization with new coefficient values learned through training at manufacture or during usage with a particular device in a particular environment . *See, e.g.*, Gorecki '716 at 7:50-61, 10:37-51; *see also* U.S. Patent No. 6,975,140 at 3:42-4:2, 9:21-28. Preserving such adapted pre-equalization filter coefficients in rewritable non-volatile memory like an EEPROM would have ensured that the fine-tuned coefficients remained available after an AEC transceiver—or the host powering it—was power cycled, initialized, or re-initialized. *See* Gorecki at 7:31-33 ("In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up, initialization or re-initialization."). Regardless, an ordinarily skilled artisan would have known that storing updated filter coefficients in the L-G combination's non-volatile memory was conventional and desirable so that the adapted coefficient values would be available after power cycling the active electric cable transceiver by unplugging it from a host device powering it, or power cycling the host device powering the active element cable transceiver. *See* U.S. Patent No. 6,975,140 at Abstract ("[D]ata transmitter includes a rewriteable non-volatile storage,

Claim 4	Disclosure
	operable to be rewritten with control information representing the values of coefficients updated during operation off the FIR driver."), 2:46-49 (same), 3:45-52 ("Updated values of the coefficients used in the taps of the FIR transmitter 30 are determined and stored to the flash memory 20 during the operation of the FIR transmitter . When the FIR transmitter is powered off and back on again, the coefficient values as last updated are available to be retrieved from the flash memory 20 and applied to the taps of the FIR transmitter again.").

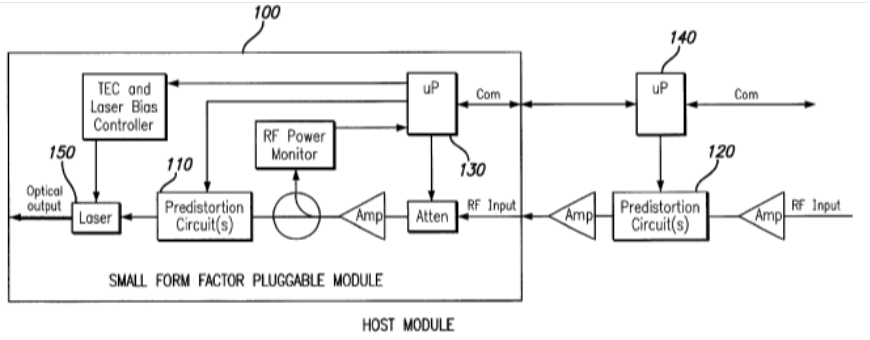


FIG. 1

Claim 4	Disclosure
	Ishaug, Figure 1. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:63–5:4. **[0043]** In some example embodiments, the interface module **100** includes a printed circuit board to which the CDR module **128**, LD module **130**, laser diode module **140**, photo diode module **148**, TIA module **146**. LA/CDR module **144** and uC module **150** are each mounted. At least some the modules mounted to the printed circuit board may be implemented as monolithic integrated circuits on respective silicon chips, including for example CDR module **126** and LA/CDR module **144**. Tremblay, paragraph 43.

E. DEPENDENT CLAIM 5

Claim 5	Disclosure
5. The cable of claim 4, wherein the first and second DRR devices employ receiver-based equalization of the electrical transit signals using coefficient values stored in the nonvolatile memories.	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.

391

Claim 5	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.

392

Claim 5	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.

393

Claim 5	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.

394

Claim 5	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.

395

Claim 5	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.

Claim 5	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.

397

Claim 5	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.

398

Claim 5	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.

399

Claim 5	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

400

Claim 5	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.

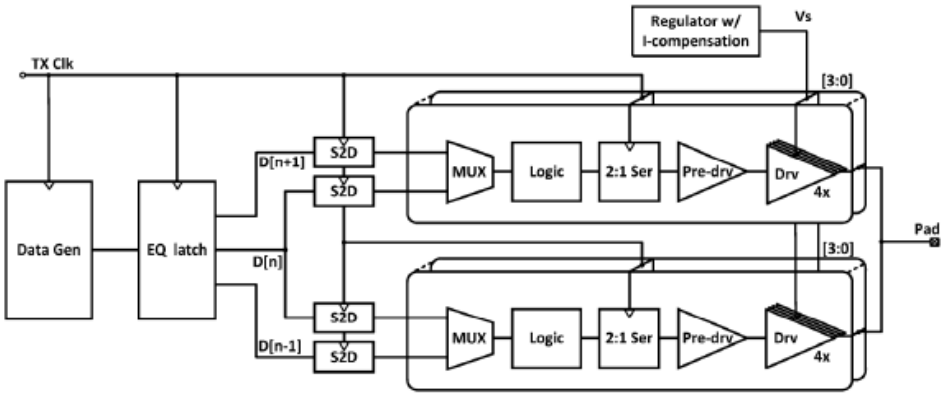
401

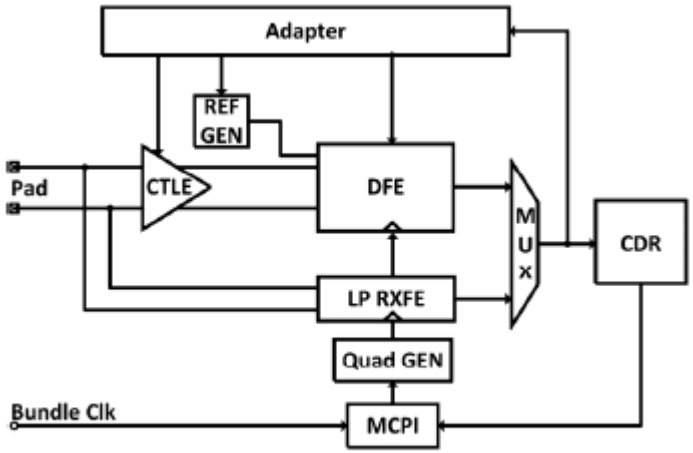
Claim 5	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

402

Claim 5	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.

403

Claim 5	Disclosure
404	 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

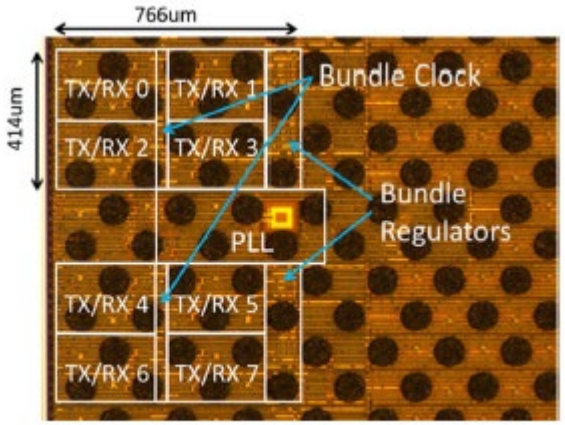
Claim 5	Disclosure
405	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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 5	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.

406

Claim 5	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.

407

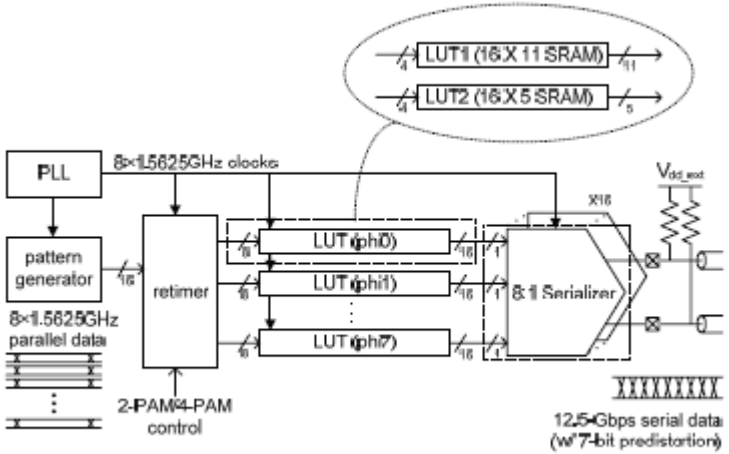
Claim 5	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 5	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.

409

Claim 5	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.

410

Claim 5	Disclosure
411	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

Claim 5	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

412

Claim 5	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.

413

Claim 5	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

414

Claim 5	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.

415

Claim 5	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.

416

Claim 5	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.

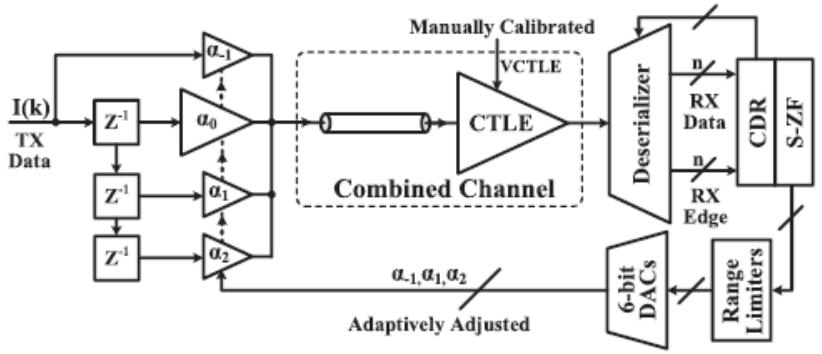


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

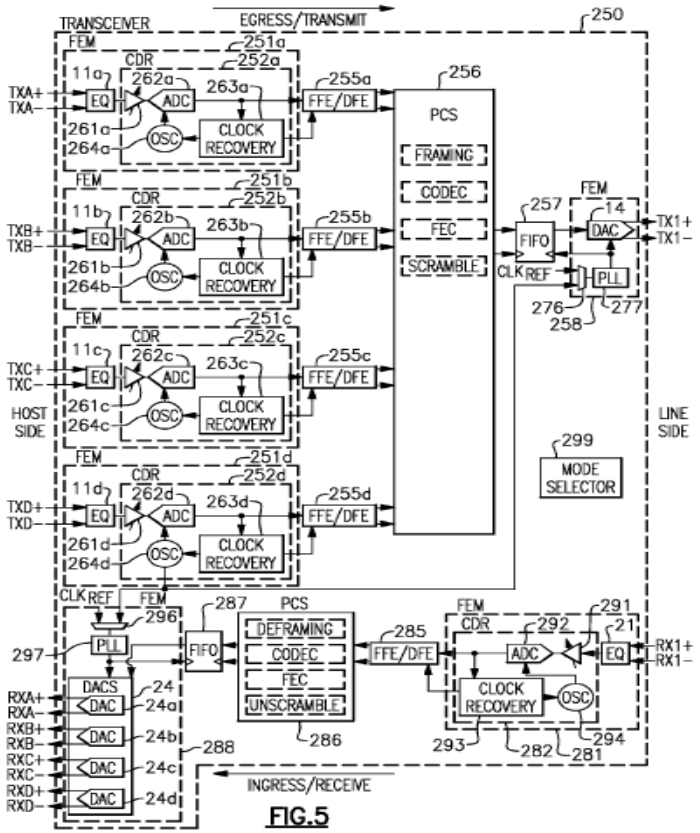
417

Claim 5	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.

418

Claim 5	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. 7,233,617 to Gorecki* The '233 Patent specification repeats the "receiver-based equalization" language of dependent claim 5, but fails to provide any explanation of what the phrase means. *See, e.g.*, '233 Patent at 2:42-45, 2:61-63. Nevertheless, "receiver-based equalization" is met by equalizing a received signal to compensate for distortion incurred during transmission. *See, e.g.*, U.S. Patent No. 9,178,542 at 10:22-35; U.S. Patent App. Pub. No. 2005/0078758 at [0004]; U.S. Patent No. 7,570,708 at 2:17-21; Pavan Kumar Hanumolu, *et al*., "Equalizers for High-Speed Serial Links," International Journal of High Speed Electronics and Systems, vol. 15, no. 2, 2005 ("Hanumolu") at p.185 (describing "receive-side equalization" suppressing channel loss). Gorecki suggest that equalizers are, like filters, configured to compensate for signal distortion over the cable. *See* Gorecki at 1:59-67 ("[C]onventional high speed digital baseband communications systems

Claim 5	Disclosure
	often employ circuitry, for example, a finite impulse response filter ('FIR filter'), in the receiver…. Such circuitry typically includes [sic] one or more taps having fixed or pre-programmed 'positions' and coefficients."); *see id.* at 2:17-21 ("Regardless of where the equalization circuitry is implemented, the duration of the equalization signal and the relative position or placement of the tap(s) of the equalization circuitry are selected or designed to avoid interference with the signal representative of the transmitted information.").
	The equalization functionality depends on filter coefficient settings. *See, e.g.,* Hanumolu at pp. 185-198 (describing different receive-side equalization architectures configured with coefficients [C-1, C0, C1]); U.S. Patent App. Pub. No. 2014/0281068 at [0069] (describing applying coefficients to a receiver's settings to minimize communication data loss).
	Though Gorecki describes implementing its equalization structure and techniques at a transmitter "may have an advantage," (*see, e.g.*, Gorecki at 18:47-53), an ordinarily skilled artisan would have recognized from this statement that implementation in the receiver was also an option, particularly as the combination explains implementation more generally in connection with equalization circuitry more generally. *See* Gorecki at 2:17-21, 6:14-23 ("The circuitry and techniques include leading and/or trailing taps to reduce, minimize, mitigate or effectively eliminate precursor and/or post-cursor intersymbol interference due to, for example, bandwidth limitations and reflections in high-speed digital communication systems. In this way, the equalization circuitry and techniques may reduce, minimize or eliminate non optimum (e.g., over-equalization) at the boundaries of the data signal (i.e., the symbol)."). An ordinarily skilled artisan would have considered such an approach would have been applicable to equalizers 21 as they provide equalization in the same way as transmit filter 17.
	Additionally, an ordinarily skilled artisan would have understood that the equalization performed at both a transmit filter and a receive path equalizer should collectively compensate for the effects introduced the communication channel and would have had a reason to set equalization parameters for both in the same way.

Claim 5	Disclosure
	 FIG.5 Lugthart, Figure 5. The illustrated transmit path line-side FEM 258 includes a transmit path DAC 14, a multiplexer 276 and a phase-locked loop (PLL) 277. Additionally, the receive path line-side FEM 281 of FIG. 5 includes a receive path equalizer 21 and a receive path CDR circuit 282, which includes a thresholder 291, an ADC 292, a clock recovery circuit 293, and a controllable oscillator 294. Furthermore, the illustrated receive

421

Claim 5	Disclosure
	path host-side FEM 288 includes a multiplexor 296, a PLL 297, and a plurality of receive path DACs 24 including a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. Lugthart, 24:24–34. 10. A transceiver comprising: an equalizer configured to receive a data signal and to generate a clock and data recovery (CDR) input signal based on equalizing the data signal; and a CDR circuit comprising: a data recovery analog-to-digital converter (ADC) configured to receive the CDR input signal and to generate a first digital output signal; a clock recovery ADC configured to receive the CDR input signal and to generate a second digital output signal, wherein the data recovery ADC has a higher resolution than the clock recovery ADC; a timing control circuit configured to receive the second digital output signal and to generate a clock control signal based on the second digital output signal; and a clock generation unit configured to control a timing of the data recovery ADC and a timing of the clock recovery ADC based on the clock control signal. Lugthart, 58:64–59:13.

422

Claim 5	Disclosure
	[0029] Example embodiments will now be described with reference to FIG. **1** which illustrates a host **102** that is connected through a host channel **103** to a pluggable network transceiver interface module **100**, which in turn is connected to a transmission medium **104**. In the illustrated embodiment, the transmission medium **104** includes fibre optic cable **124** for transmitting high speed data from the interface module **100** and fibre optic cable **126** for transmitting high speed data to the interface module **100**, however embodiments described herein can also be applied to non-fibre transmission mediums such as copper for example. In the illustrated embodiment, the host **102** includes a network interface card or board **106** that allows data to be exchanged between the host **102** and the network interface module **100**. In at least one example embodiment, the host interface board **106** includes data processing circuitry, which could for example be implemented an ASIC **108** ("Application Specific Integrated Circuit"), and a SERDES ("SERializer-DESerializer") chip or module **110**. SERDES module **110** receives parallel data from the host ASIC **108** over a parallel data channel **170** and converts it into high speed serial data for transmission over the host channel **103** to the network interface module **100**, and receives high speed serial data from the network interface module **100** through the host channel **103** which it then converts into parallel data and passes on to the ASIC **108** over the parallel data channel **170**. Tremblay, paragraph 29. **[0038]** With reference to the drawings, FIG. **2** is a flow chart of the operation of an example embodiment, wherein the LA/CDR unit **144** in conjunction with the microcontroller

423

Claim 5	Disclosure
	150 performs pre-equalization and optimizes power consumption of the transmission over host channel **103** via physical conductors **114** based on the feedback received from the SERDES module **110** via the low-speed digital interface **152**. In FIG. **2**, the operational flow of the LA/CDR unit **144** and the microcontroller **150** begins at step **202** when the microcontroller receives feedback via the low speed digital interface **152**. At step **204**, the microcontroller **150** instructs the LA/CDR unit **144** to equalize its transmitted data signal in response to the feedback, applying sufficient transmit pre-emphasis and other pre-equalization effects to the transmitted data signal to achieve predetermined target eye characteristics at the SERDES module **110**. Once the feedback from the SERDES module **110** indicates that the predetermined target eye characteristics have been achieved, the microcontroller **150** at step **206** instructs the LA/CDR unit **144** to begin reducing the launch amplitude of the transmitted data signal as much as it can without letting the received signal at the SERDES module **110** fall short of the predetermined target eye characteristics. The process is iterative, with feedback being monitored at step **202** and equalization being adjusted at step **204** while the launch amplitude is being reduced at step **206**. In alternative embodiments, the process may operate according to a different sequence of steps while achieving the same result, namely the optimization of signal quality and power consumption over the host channel **103**. Tremblay, paragraph 38.

F. DEPENDENT CLAIM 6

Claim 6	Disclosure
6. The cable of claim 4, wherein the electrical conductors comprise twin-axial conductors that carry each of the electrical transit signals in differential form.	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

425

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.

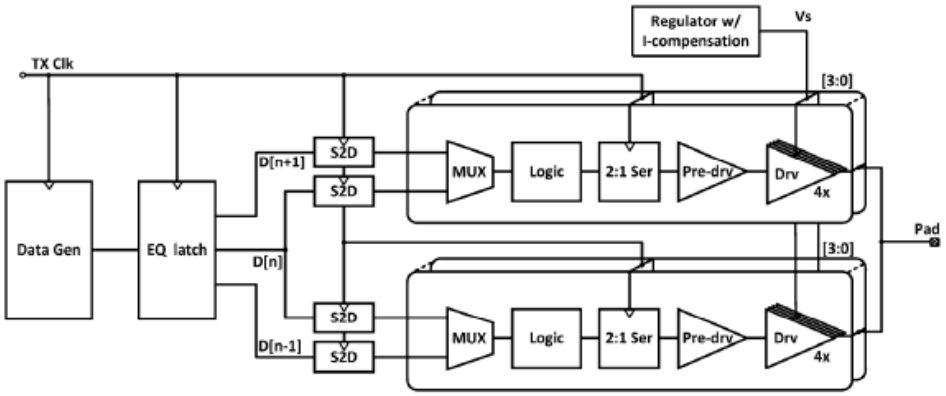
426

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.

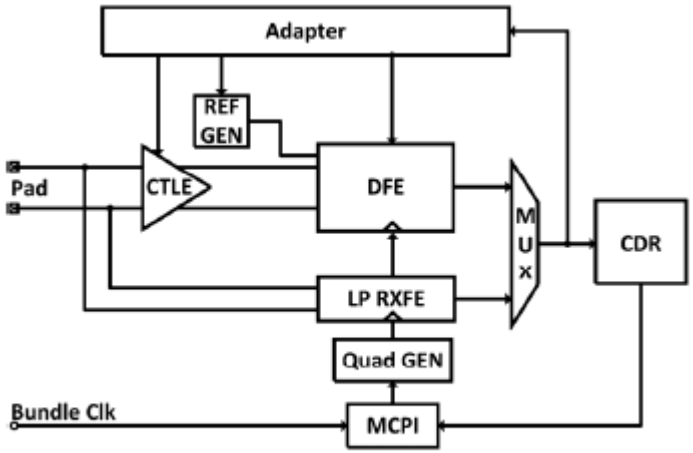
427

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.

428

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

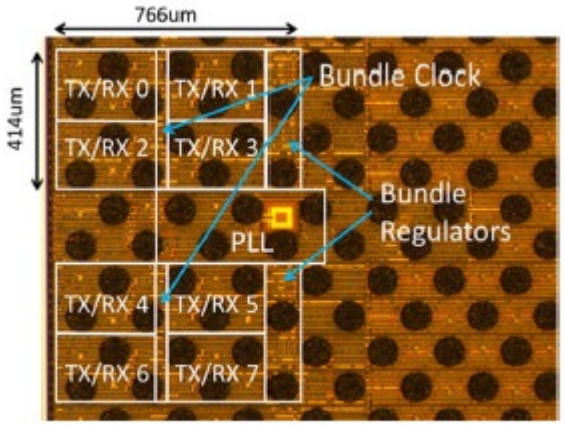
429

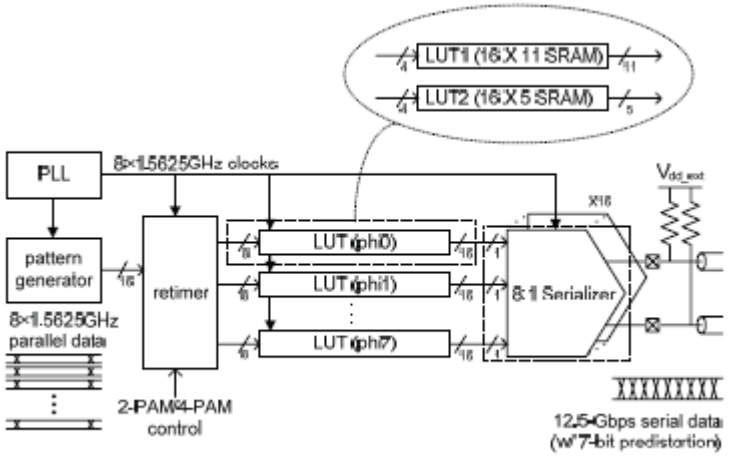
Claim 6	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

430

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.

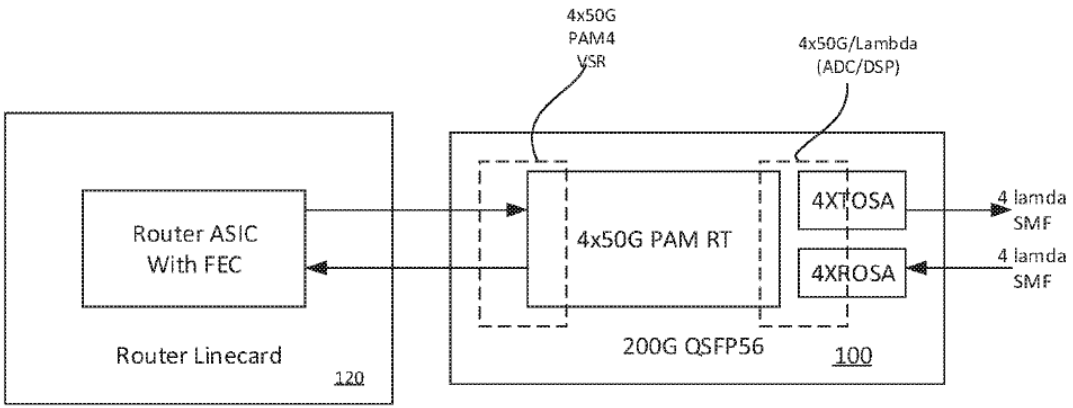
Claim 6	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses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

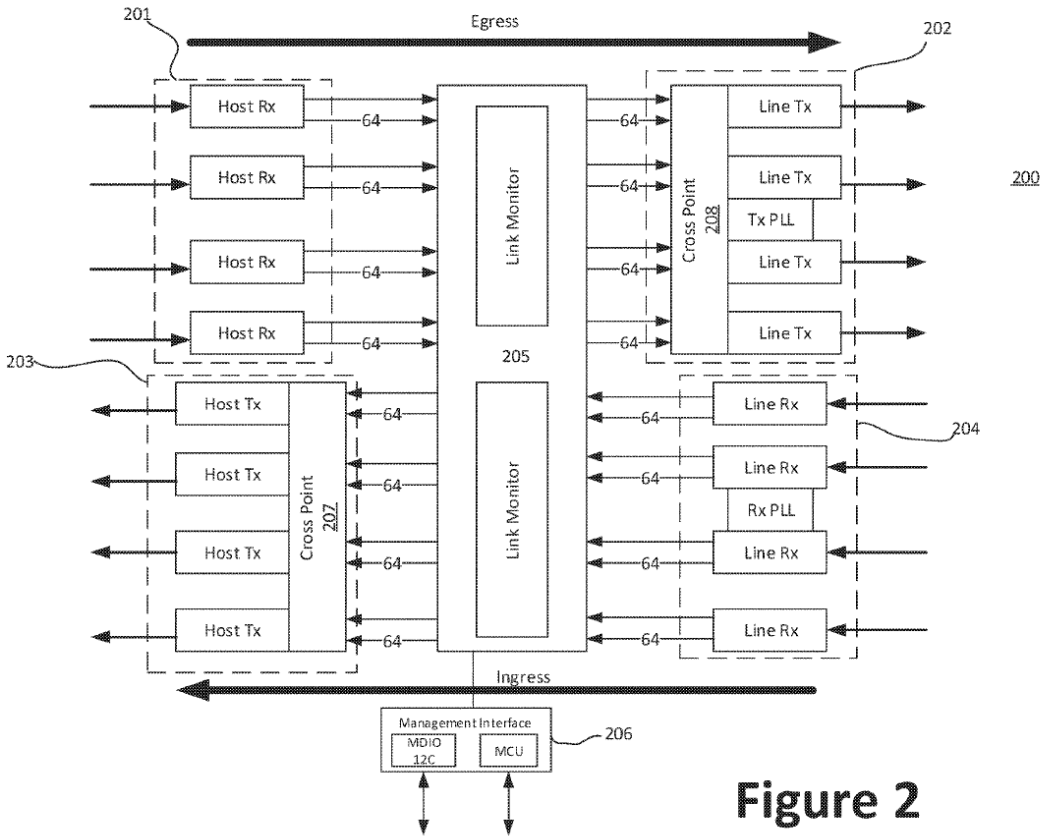
Claim 6	Disclosure
	We have described a scalable bidirectional serial link that targets a variety of applications with data rate requirements within the range of 4 to 32 Gb/s. Supply voltage and TX output swing scaling, regulated CMOS clocking and clock bundling were employed to maintain the per lane 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.  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

435

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.

Claim 6	Disclosure
437	 **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.

439

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.

441

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.

442

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.

443

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.

444

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.

445

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.

446

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.

447

	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.	

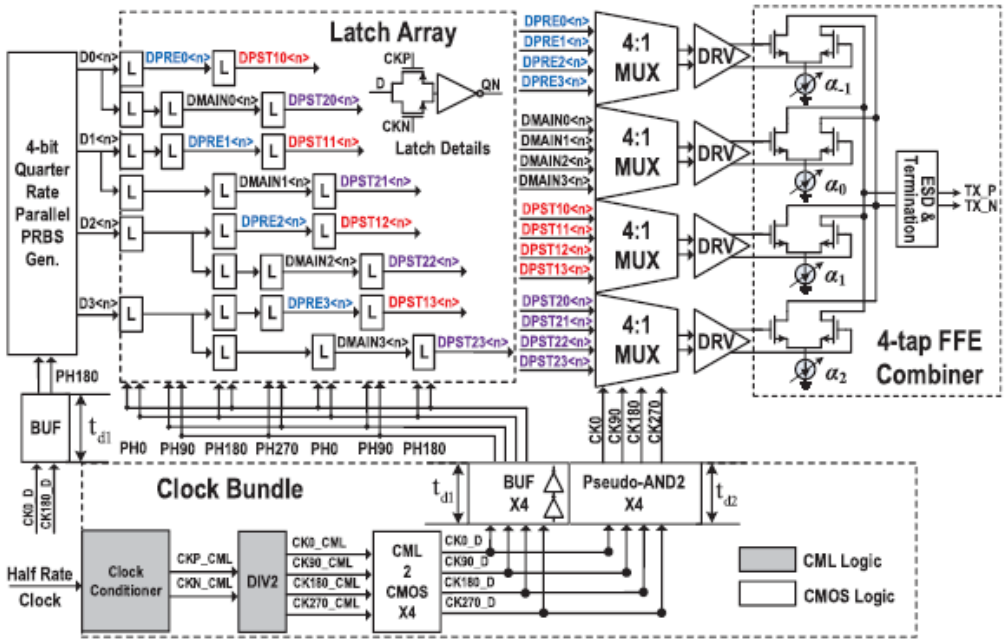
448

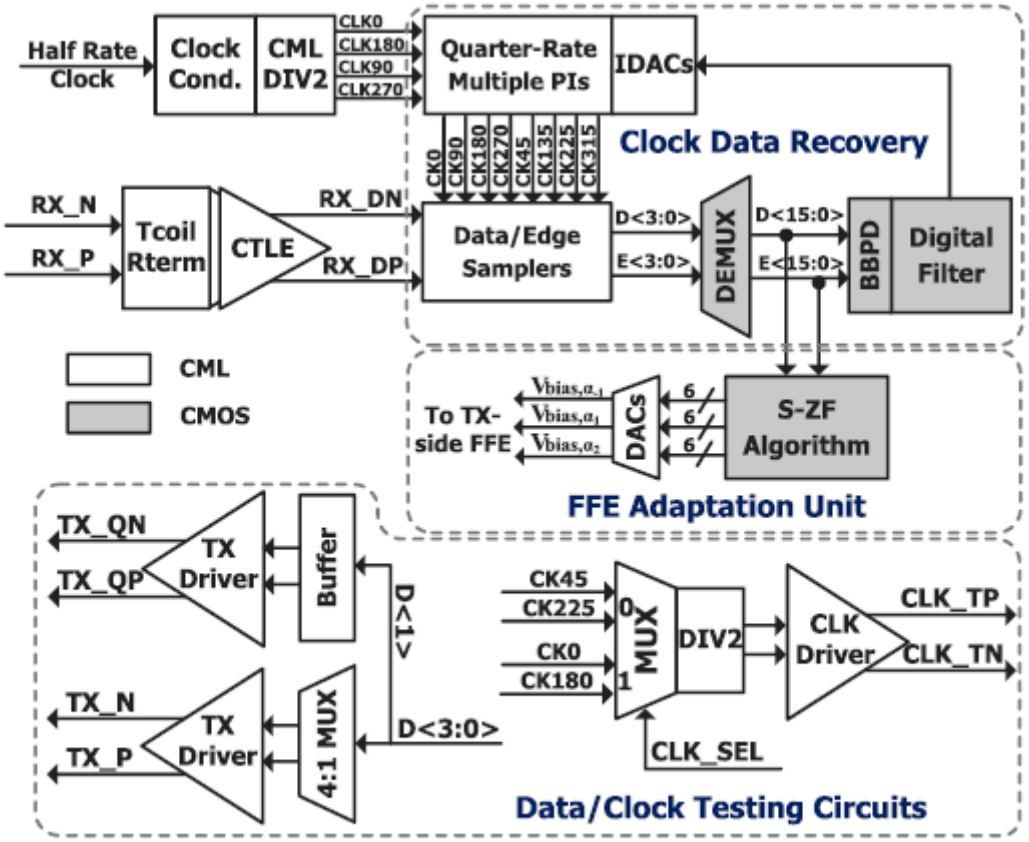
Claim 6	Disclosure
	Riani, 11:4-45. It is to be appreciated that embodiments of the present invention implemented in various ways. For example, re-timers according to embodiments of the present invention operate in low power and are compatible with PAM communication protocols. In various implementations, 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.

449

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.

450

Claim 6	Disclosure
	 Fig. 1. Block diagram of the TX chip. Zheng, 2964.

Claim 6	Disclosure
	 Fig. 8. Block diagram of the RX chip. Zheng, 2966.

452

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

453

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.

454

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.

455

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.

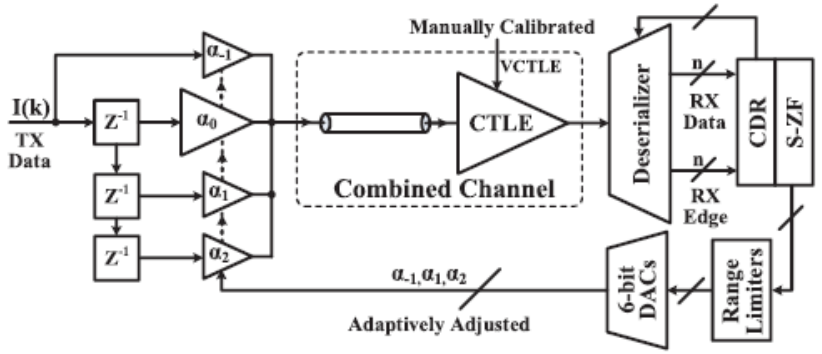
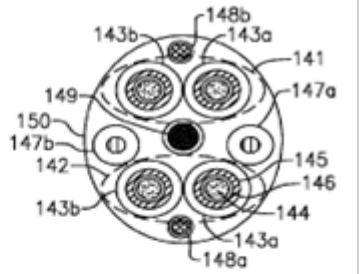
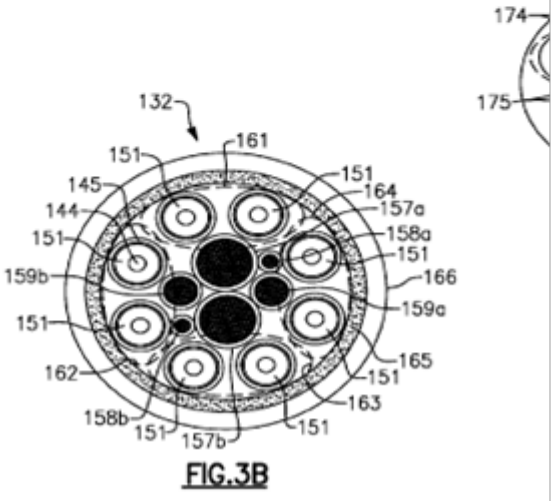


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

456

Claim 6	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.

Claim 6	Disclosure
	 FIG.3A  FIG.3B Lugthart, Figures 3A and 3B. FIG. 3A illustrates a cross-sectional view of one embodiment of a cable 131.

458

Claim 6	Disclosure
	The illustrated cable 131 includes a first pair of conducting lines 141 and a second pair of conducting lines 142. Each of the pairs of conducting lines 141-142 can be configured for differential signal transmission. For instance, the individual conducting lines in the pair can be implemented coaxially in certain embodiments. In the illustrated embodiment, each pair of conducting lines 141-142 includes a first micro coaxial cable 143 a and a second micro coaxial cable 143 b. In the depicted embodiment, each of the micro coaxial cables includes an inner conductor 144 and an outer conductor 145. An insulating material or a dielectric 146 is disposed between the inner conductor 144 and the outer conductor 145. In various implementations, the insulating material or dielectric 146 can include a polymer or a fluoroplastic material. Lugthart, 17:9–24. For example, FIG. 3B illustrates a cross-sectional view of another embodiment of a cable 132. The cable 132 includes eight individual conducting lines 151, which can be, for example, micro coaxial cables. As depicted by the dashed lines, according to some embodiments, each of the conducting lines 151 can be paired with another of the conducting lines 151 to form a differential pair, and a single corresponding logical channel. For example, the cable 132 can include a first pair of conducting lines 161, a second pair of conducting lines 162, a third pair of conducting lines 163, and a fourth pair of conducting lines 164. Lugthart, 17:63–18:6.

459

G. DEPENDENT CLAIM 7

Claim 7	Disclosure
7. The cable of claim 1, wherein the first and second DRR devices do not perform pre-equalization of the multi-lane data streams provided to the first and second host interface ports.	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.

460

Claim 7	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.

461

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.

462

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.

463

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.

464

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.

465

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.

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.

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.

468

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

469

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.

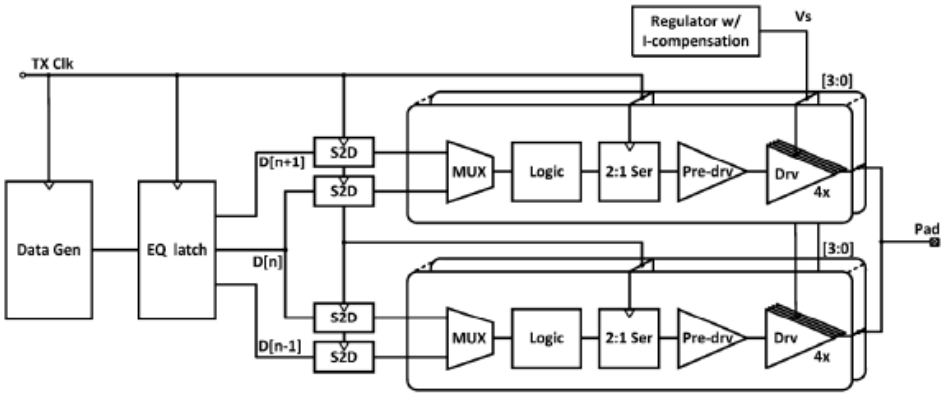
470

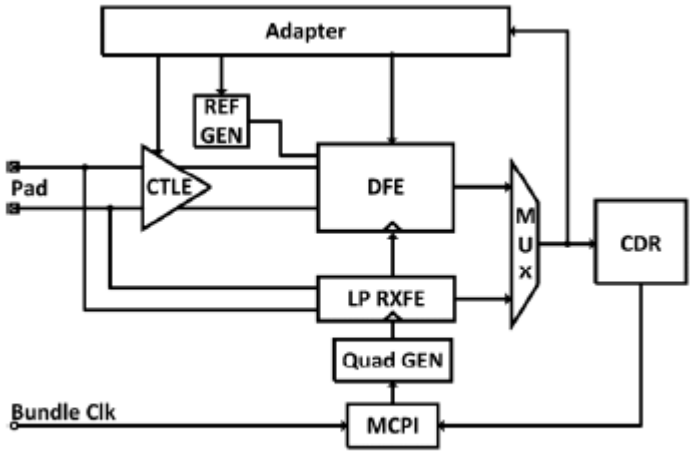
Claim 7	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

471

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.

472

Claim 7	Disclosure
473	 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

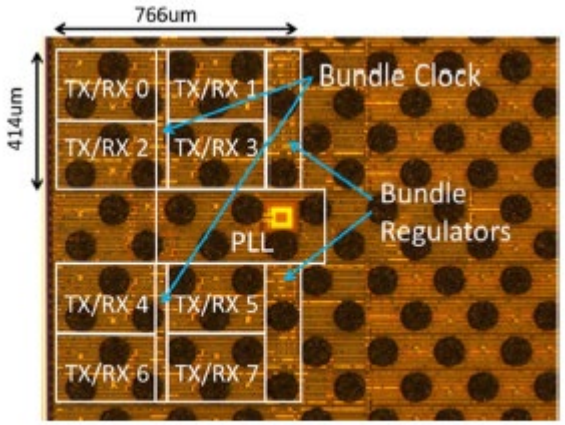
Claim 7	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

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.

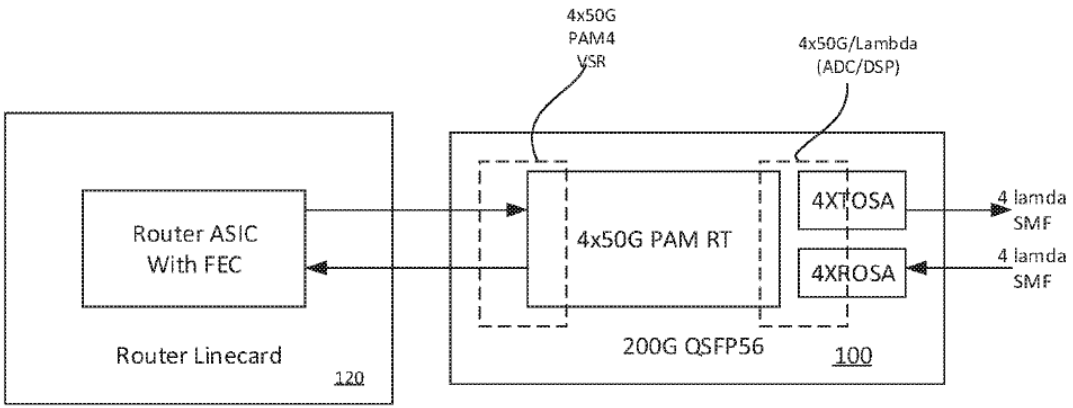
475

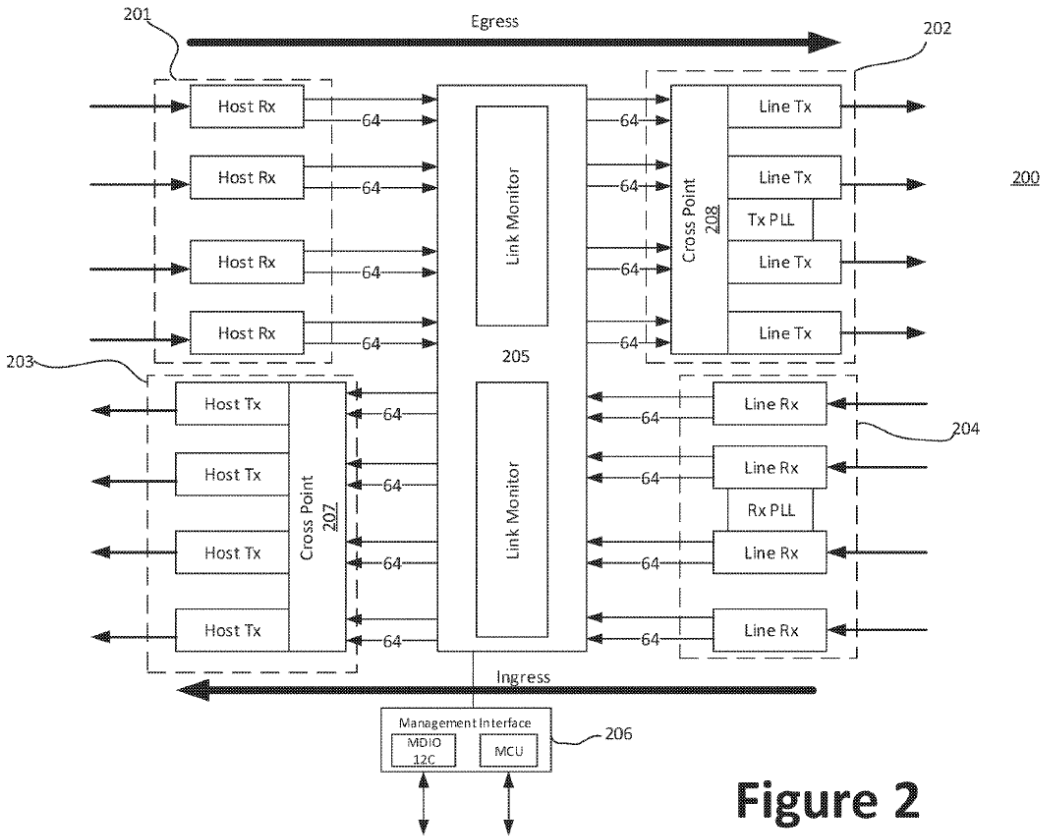
Claim 7	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed 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.

476

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.

Claim 7	Disclosure
	 **Figure 1** Riani, Figure 1.

Claim 7	Disclosure
480	 Figure 2 Riani, Figure 2.

Claim 7	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.

481

Claim 7	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 7	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.

483

Claim 7	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.

484

Claim 7	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.

485

Claim 7	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.

Claim 7	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.

487

Claim 7	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.

488

Claim 7	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.

489

	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.	

490

Claim 7	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.

491

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

492

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.

493

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

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.

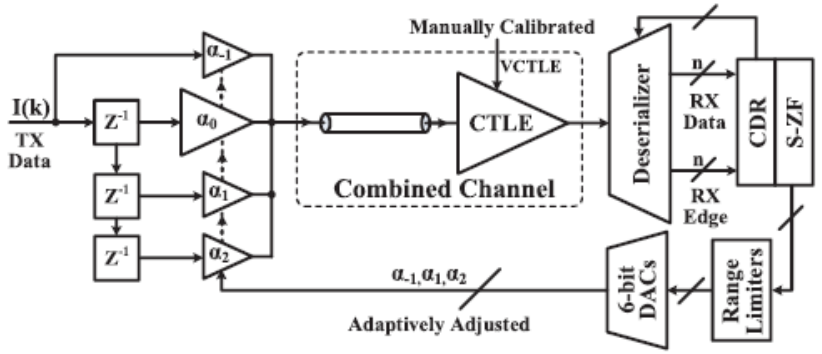


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

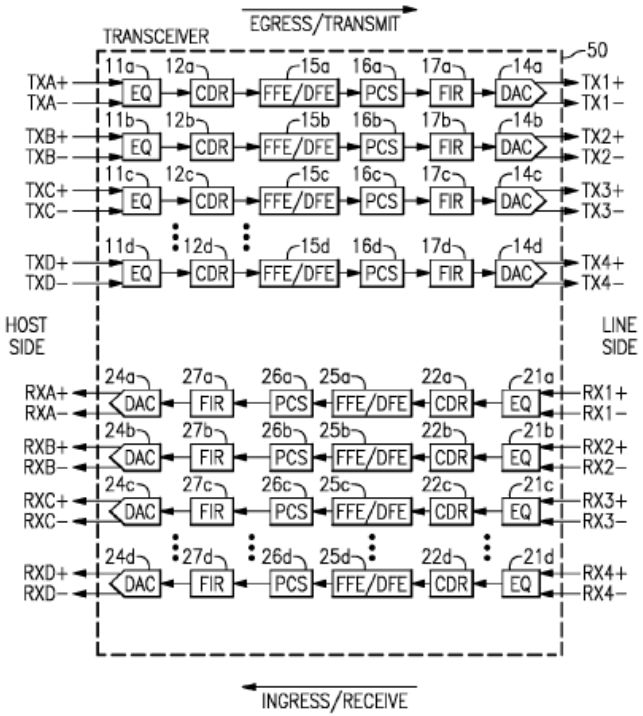
Claim 7	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

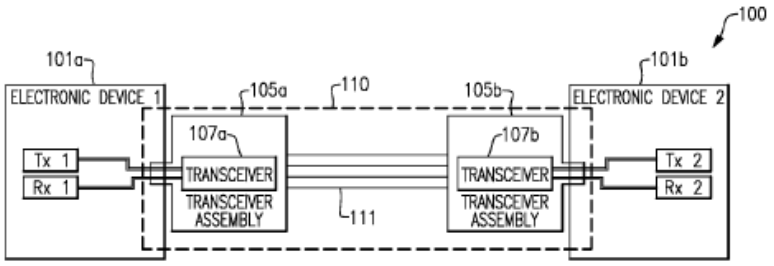
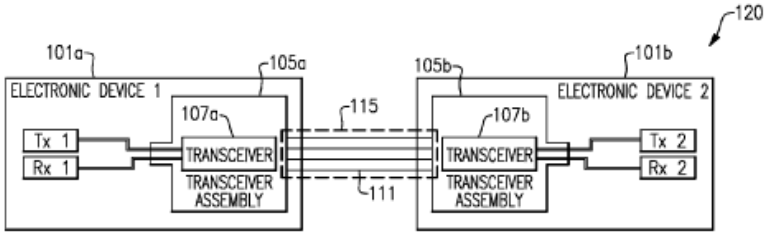
496

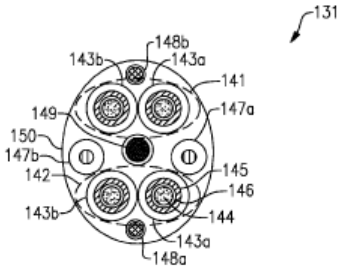
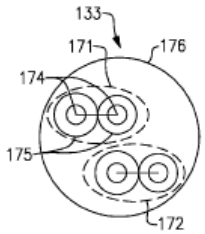
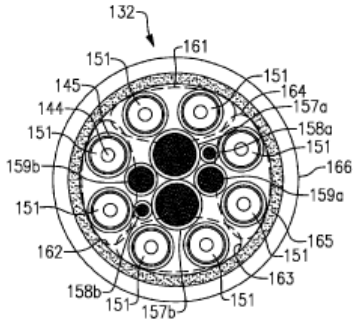
Claim 7	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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. *IEEE 802.3-2015* The IEEE 802.3-2015 standard discloses setting transmit filter coefficients to disable pre-equalization. As was known, an ordinarily skilled artisan would have disabled pre-equalization "if not needed to consume minimum power while satisfying the performance or margin objective" for a system. *See* U.S. Patent No. 9,137,063 to Zerbe at 27:24-37. Thus, for instances where signal transmission between host devices using the active cable meets performance objectives without pre-equalization of the multi-lane data streams provided to the first and second host interface ports, a user would disable those pre-equalization based on the IEEE 802.3-2015 standard.

497

Claim 7	Disclosure
	72.6.10.3.4 Functions COEFF_UPDATE(coefficient, preset, initialize, inc, dec) Returns an updated coefficient based on the contents of the coefficient update field in the training frame. Sets a fixed coefficient value, or adds, or subtracts from the current coefficient value to create the updated coefficient. If multiple actions are requested in the coefficient update field, then the priority is: 1) preset 2) initialize 3) inc/dec Values: preset; If preset is TRUE then the function returns the coefficient value equivalent to no equalization [$c(-1)$ and $c(1)$ coefficients are set to zero, $c(0)$ set to maximum]. initialize; If initialize is TRUE, then the function returns the coefficient value such that the transmit output meets the conditions defined in 72.6.10.4.2. inc; If inc is TRUE then the function returns (coefficient + step). dec; If dec is TRUE then the function returns (coefficient – step). The requirements for the value of step are defined in 72.7.1.10 and Table 72–7. IEEE 802.3-2015, Section 5 at 480.

Claim 7	Disclosure
	 FIG.1D Lugthart, Figure 1D.

Claim 7	Disclosure
500	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

Claim 7	Disclosure
501	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3C.

Claim 7	Disclosure
	According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:66–4:5. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second

502

Claim 7	Disclosure
	receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1–53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver

503

Claim 7	Disclosure
	assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:3–38. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

504

Claim 7	Disclosure
	[0037] Accordingly, the system of FIG. **1** provides a back-channel using a low speed interface to set the operating parameters of the serializer section **160** of SERDES module **110** and the operating parameters of the LA/CDR module **144** of the transceiver interface module **100**. In some applications, this allows the operating parameters to be change over time which may in turn facilitate adaptive optimization of high data speed transmission characteristics and power consumption used for the communications link. For example, some embodiments may optimize power consumption used for the communications link by reducing the launch amplitude of the signals transmitted across the host channel **103**. By monitoring the eye characteristics of the signal at the SERDES eye monitor **164** and the CDR eye monitor **142** and modifying the transmit equalization parameters used by the LA/CDR module **144** and the serializer section **160** of the SERDES module **110**, respectively, the predetermined target eye characteristics can be maintained through transmit equalization even while launch amplitude is decreased. Some embodiments may also optimize power consumption by deactivating the circuits that perform transmit equalization when the monitored eye characteristics indicate that pre-equalization is unnecessary for maintaining the predetermined target eye characteristics. Tremblay, paragraph 37.

505

H. INDEPENDENT CLAIM 8

Claim 8	Disclosure
8[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.

506

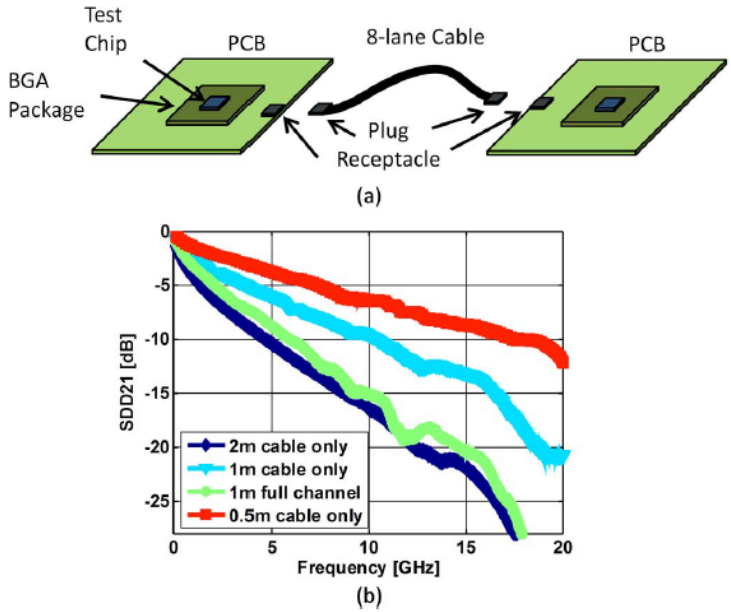
Claim 8	Disclosure
	Abstract: This document defines the Common Management Interface Specification (CMIS) for pluggable modules such as QSFP Double Density (QSFP-DD), OSFP and COBO. This document provides a common specification for systems manufacturers, system integrators, and suppliers of modules and 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 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

508

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.

509

Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

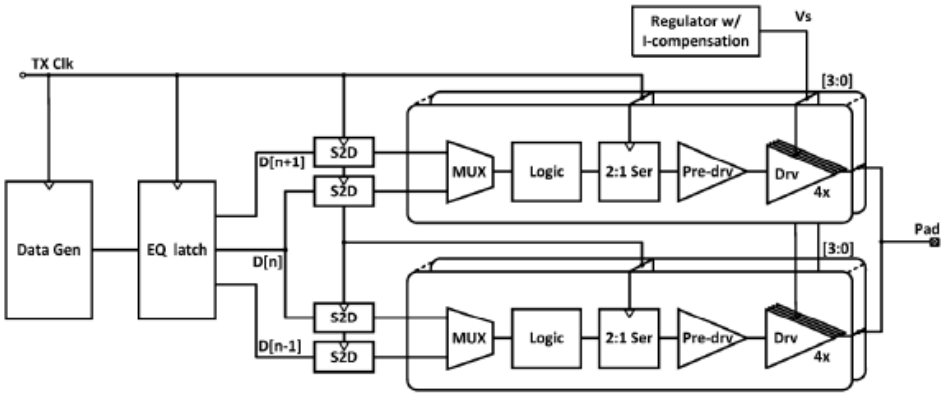
510

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.

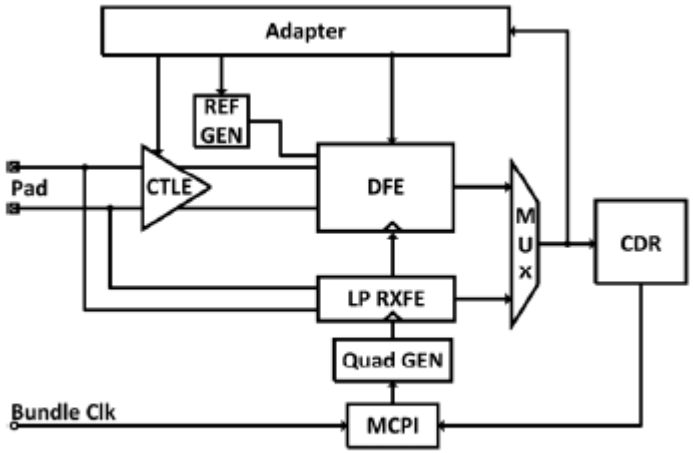
511

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.

512

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

513

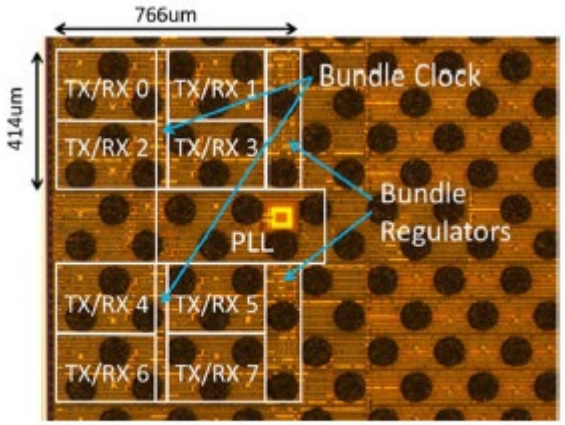
Claim 8	Disclosure
514	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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 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.

515

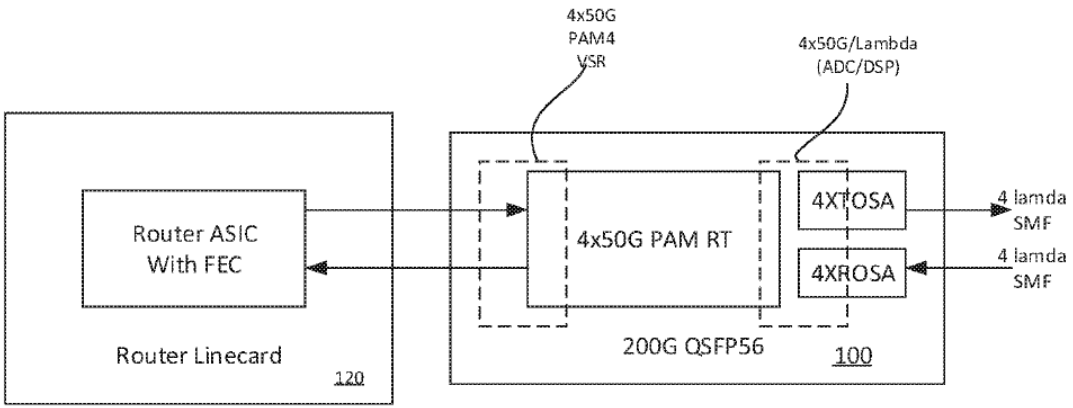
Claim 8	Disclosure
	The CTLE drives 4 slices of the reconfigurable DFE. Fig. 13 shows a simplified schematic of one DFE slice. The quadrature sampling clocks for the interleaved track and hold circuit ensure that only two DFE slices load the CTLE at any instance, thereby reducing the CTLE power needed 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.

516

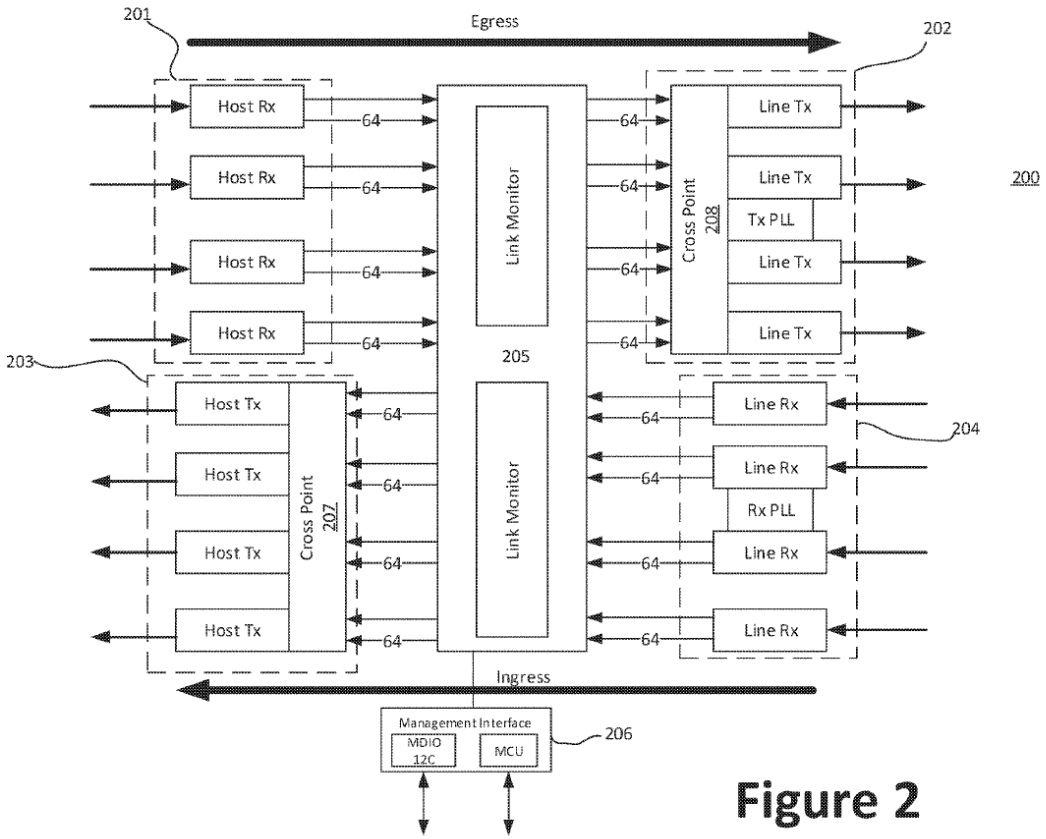
Claim 8	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

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.

518

Claim 8	Disclosure
	 **Figure 1** Riani, Figure 1.

519

Claim 8	Disclosure
	 Figure 2 Riani, Figure 2.

520

Claim 8	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.

521

Claim 8	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 8	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.

523

Claim 8	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.

524

Claim 8	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.

525

Claim 8	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.

526

Claim 8	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.

527

Claim 8	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.

528

Claim 8	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.

529

	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.	

530

Claim 8	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.

531

Claim 8	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.

532

Claim 8	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.

533

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

534

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.

535

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.

536

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.

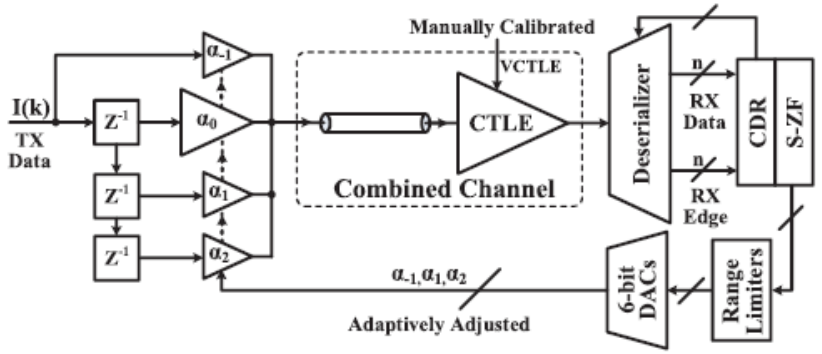
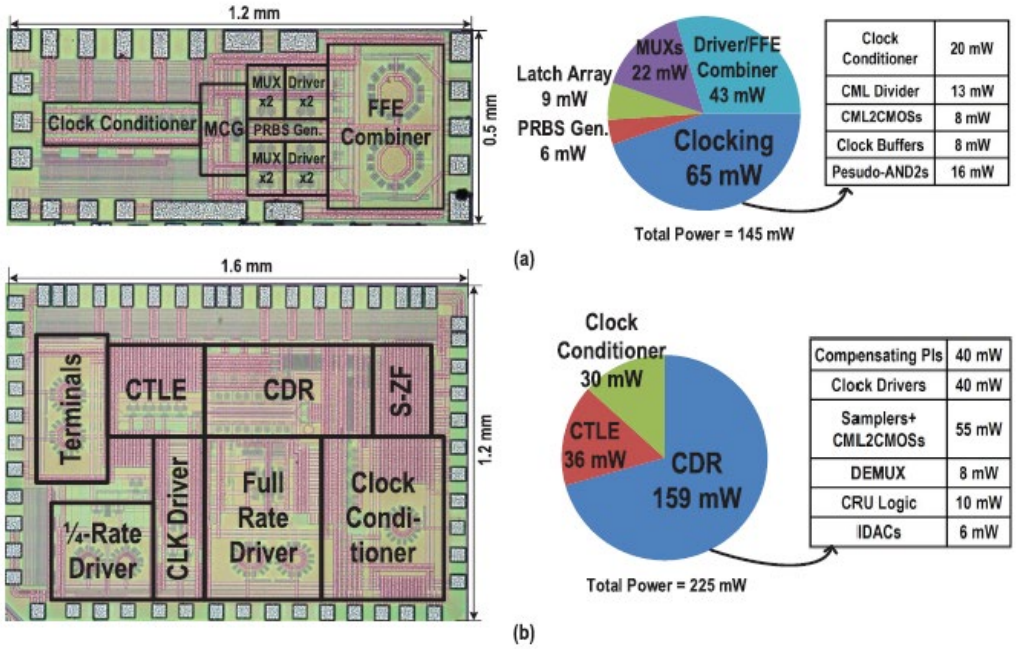


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

537

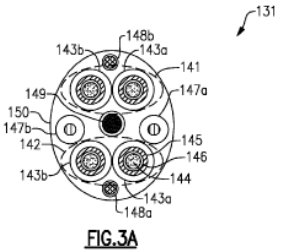
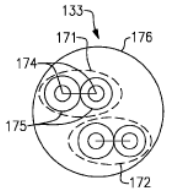
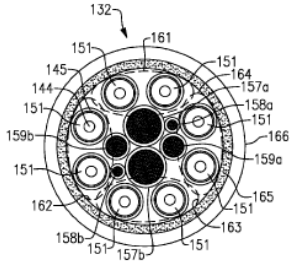
Claim 8	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

538

Claim 8	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

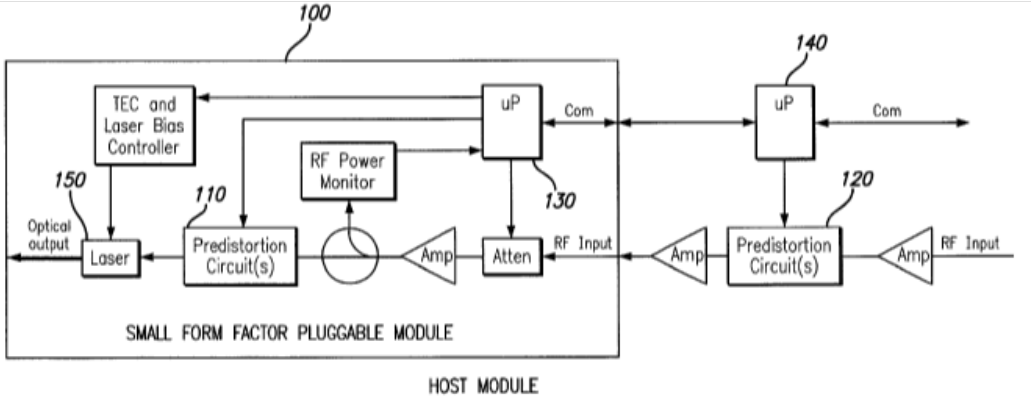
539

Claim 8	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.

Claim 8	Disclosure
541	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3C. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described,

Claim 8	Disclosure
	other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. Circuits, methods, and apparatus that provide for the powering of active components in connector inserts at each end of a cable may in various ways. For example, where a host is coupled to a device that is not self-powered, the host may provide power for circuitry at each end of the cable. In various embodiments of the present invention, the device may request higher voltage from the host, such that more power can be delivered. In these cases, the device may regulate the voltage received from the host to a lower voltage, and then provide the lower voltage to circuitry at one or both ends of the cable. Where the host is connected to a device that is self-powered, the host and the self-powered device may power their respective connector insert circuits. Baker, Abstract.

542

Claim 8	Disclosure
	 FIG. 1 Ishaug, Figure 1.

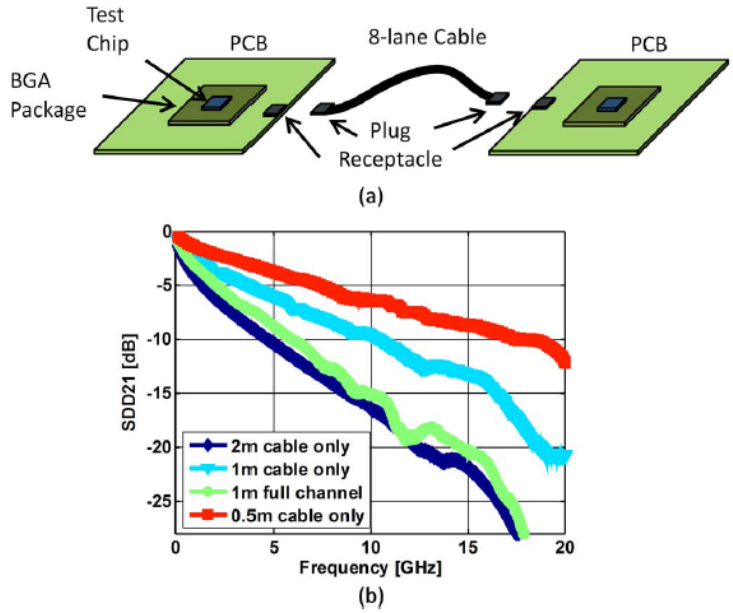
Claim 8	Disclosure
	In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. Ishaug 4:51–63. 0027. In high speed data networks a host network element Such as a Switch or router (also referred to as a host) can be connected to a transmission medium (for example fibre cable, copper cable or other physical medium) through an interface module that processes data coming in from the transmission medium to place it into a condition suitable for the host and that processes data from the host to place it into a condition suitable for the transmission medium. Tremblay, paragraph 27. 0029. Example embodiments will now be described with reference to FIG. 1 which illustrates a host 102 that is con nected through a host channel 103 to a pluggable network transceiver interface module 100, which in turn is connected to a transmission medium 104. In the illustrated embodiment, the transmission medium 104 includes fibre optic cable 124 for transmitting high speed data from the interface module 100 and fibre optic cable 126 for transmitting high speed data to the interface module 100, however embodiments described herein can also be applied to non-fibre transmission mediums Such as copper for example. Tremblay, paragraph 29.

Claim 8	Disclosure
8[a] connecting a first connector plug to a first data recovery and remodulation (DRR) device that exchanges multi-lane data streams with a first host interface port via the first connector plug;	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

545

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.

546

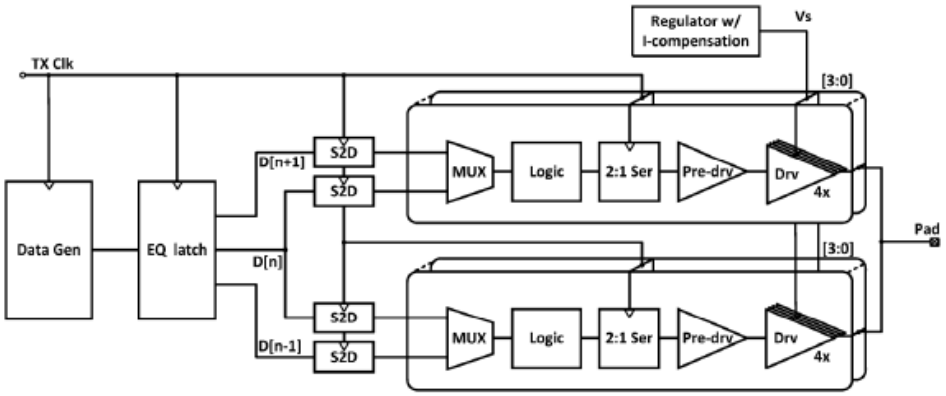
Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

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.

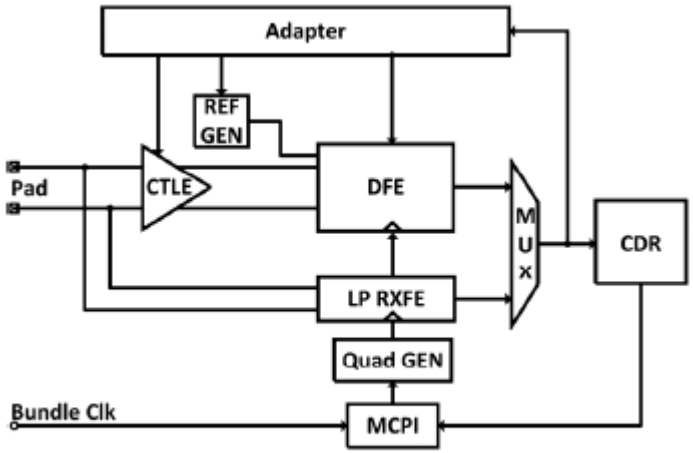
548

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.

549

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

550

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.

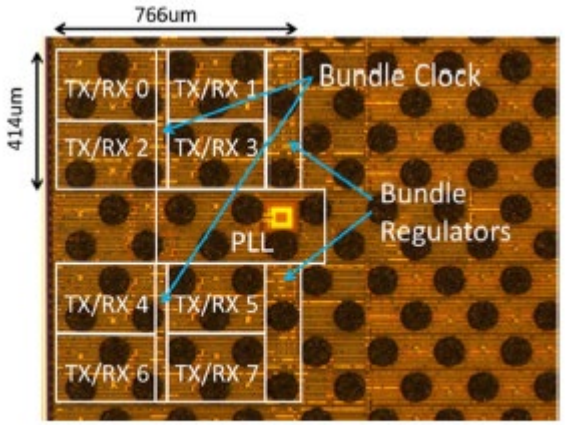
551

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.

552

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.

553

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.

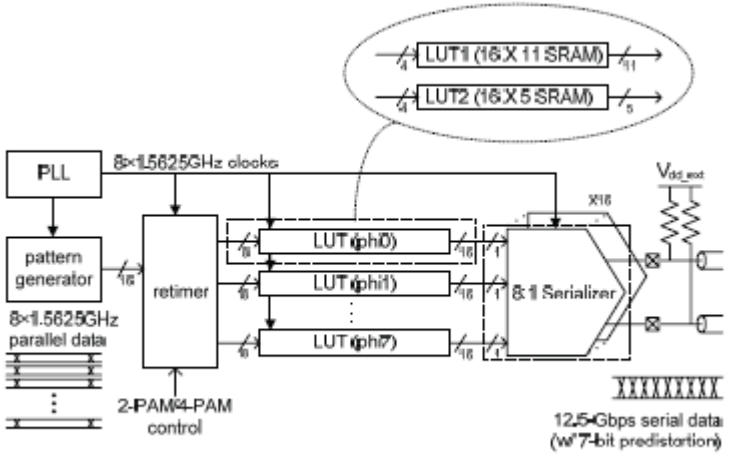
554

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.

555

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.

556

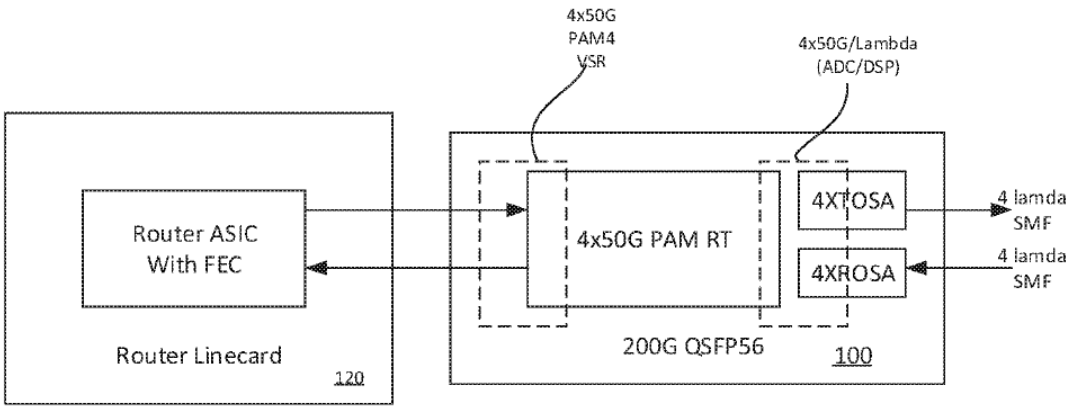
Claim 8	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

557

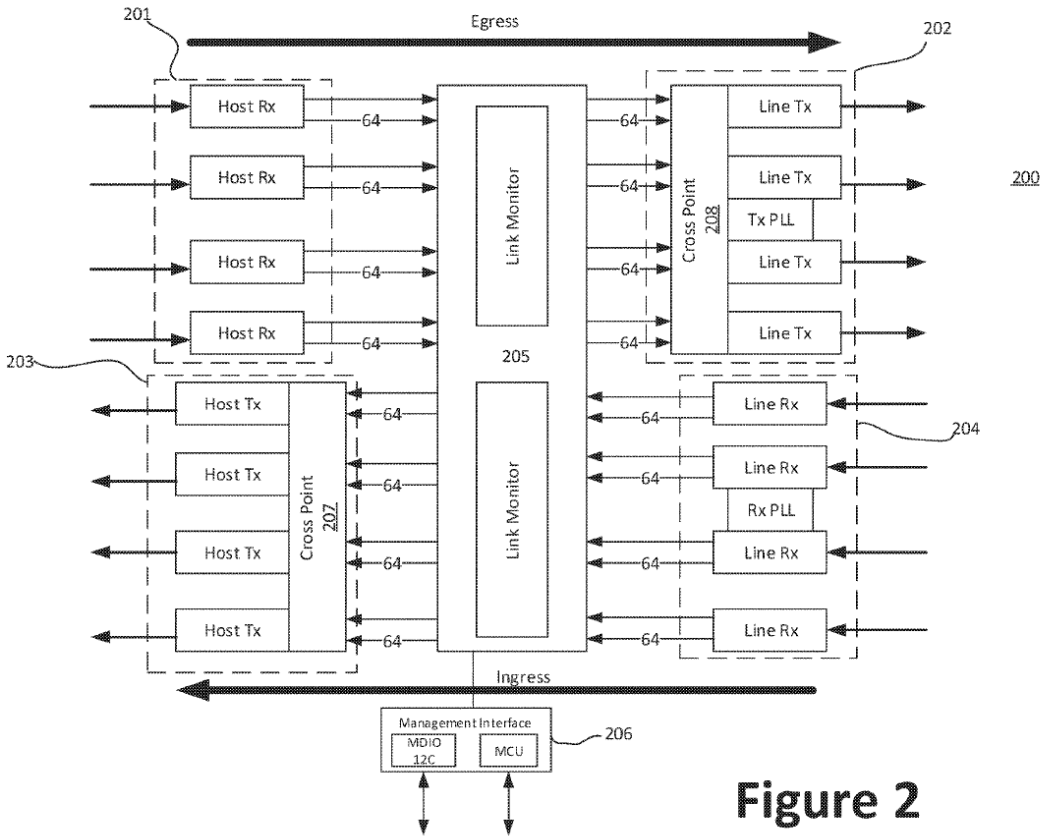
Claim 8	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

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.

559

Claim 8	Disclosure
	 **Figure 1** Riani, Figure 1.

560

Claim 8	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 8	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.

562

Claim 8	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.

563

Claim 8	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.

564

Claim 8	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.

565

Claim 8	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.

566

Claim 8	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.

567

Claim 8	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.

568

Claim 8	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.

569

Claim 8	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.

570

	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.	

571

Claim 8	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.

572

Claim 8	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.

573

Claim 8	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.

574

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

575

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.

576

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.

577

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.

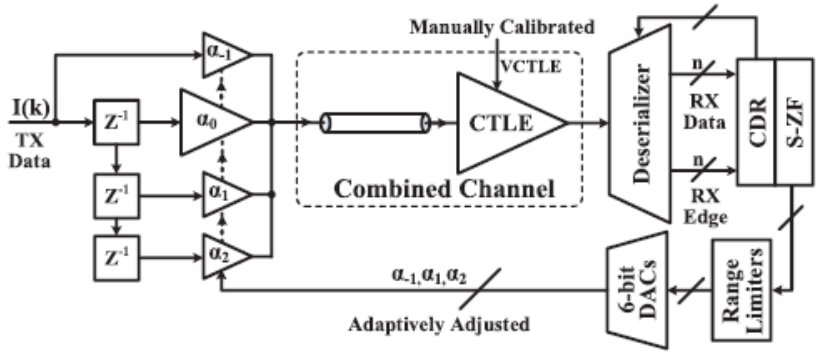
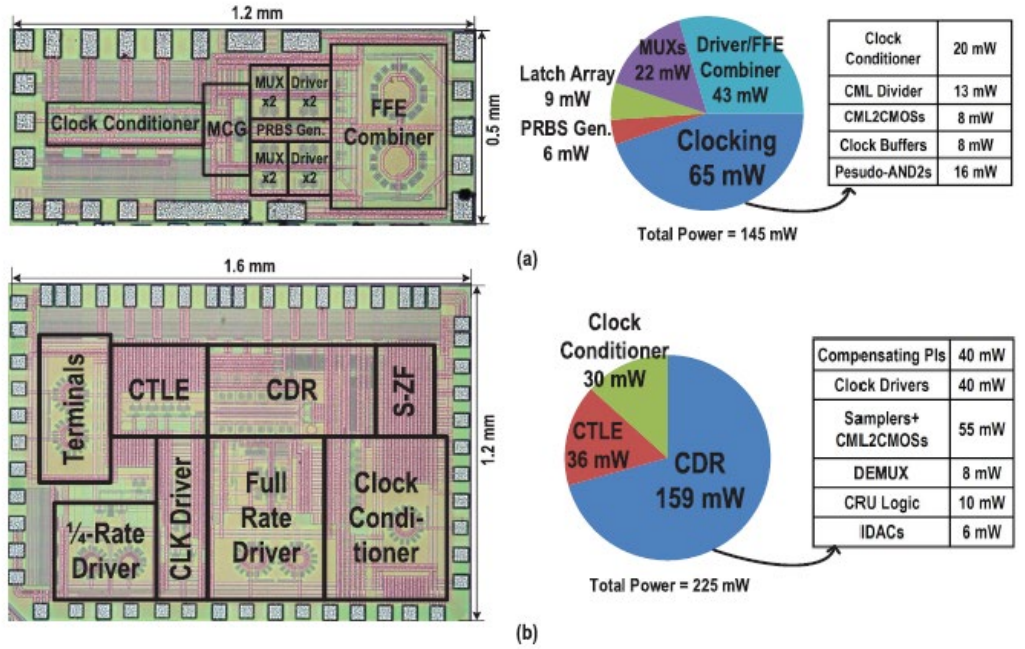


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

578

Claim 8	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.

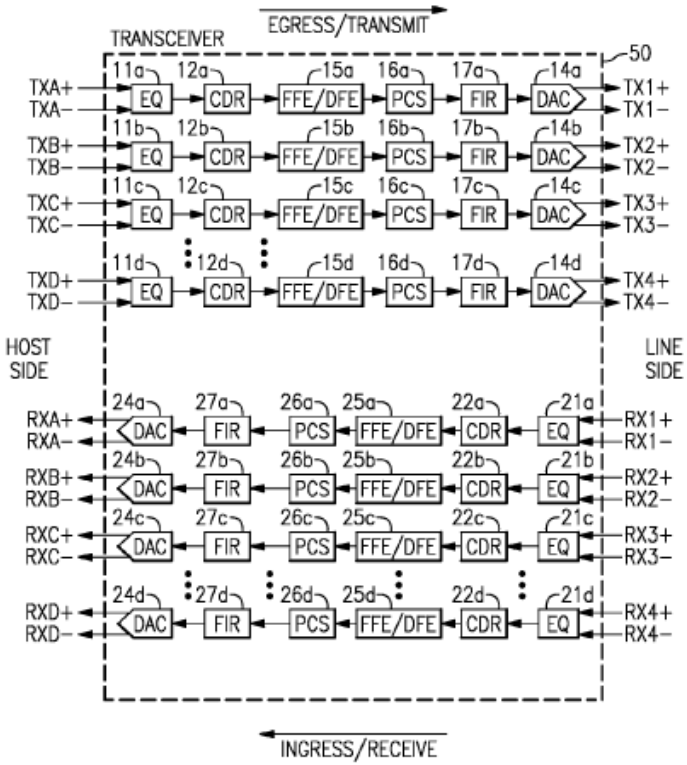
579

Claim 8	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

580

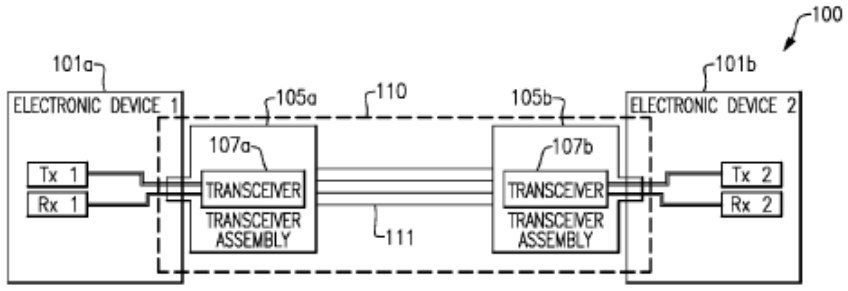
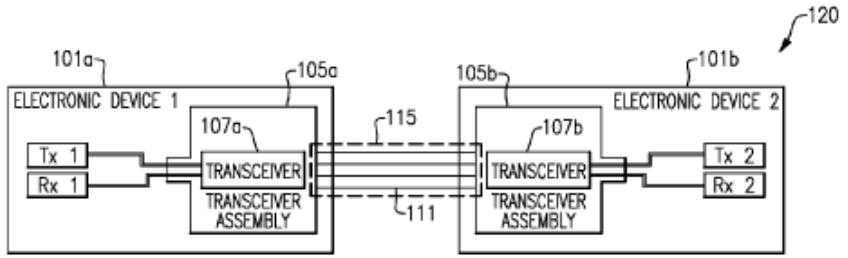
Claim 8	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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.

581

Claim 8	Disclosure
	 FIG.1D Lugthart, Figure 1D.

582

583

Claim 8	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figure 2A–2B. According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver

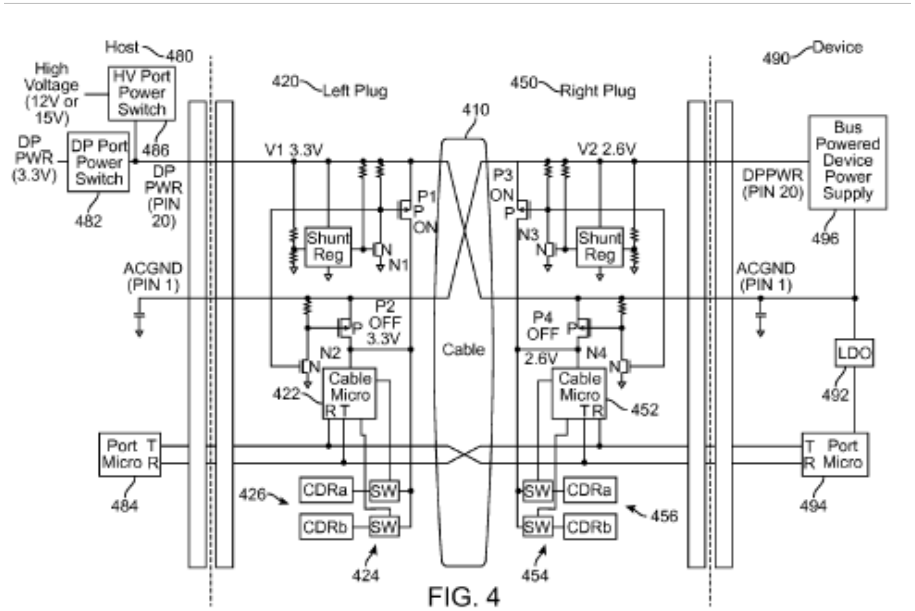
Claim 8	Disclosure
	including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:66 –4:5. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances.

584

Claim 8	Disclosure
	In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1 –53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b.

585

Claim 8	Disclosure
	Lugthart, 14:3–38. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:19–23.

Claim 8	Disclosure
	 FIG. 4 Baker, Figure 4.

Claim 8	Disclosure
	FIG. 4 illustrates an electronic system according to an embodiment of the present invention. This figure includes host 480 coupled to device 490 via cable 410. Cable 410 includes left plug 424 connected to host 480, and right plug 454 connected to device 490. Host 480 may be capable of providing one or more voltages to both device 490 and cable circuitry in cable 410. Host 480 may provide a lower voltage of 3.3 V, or a higher voltage of 12 or 15 volts. In other embodiments of the present invention, host 480 may provide various voltage levels to host 490 and cable circuitry in cable 410. Baker, 5:26 –36. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12. **[0031]** In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**.

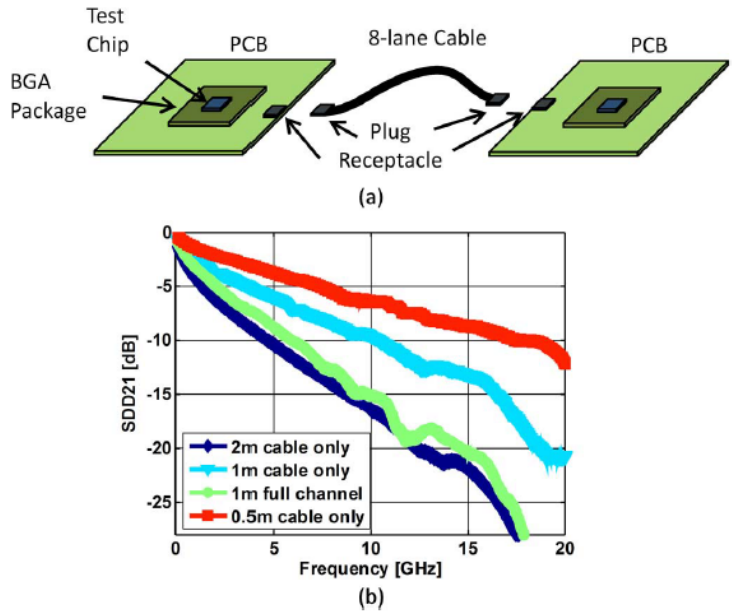
588

Claim 8	Disclosure
	Tremblay, paragraph 31.
8[b] connecting a second connector plug to a second DRR device that exchanges multi-lane data streams with a second host interface port via the second connector plug;	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

589

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.

590

Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

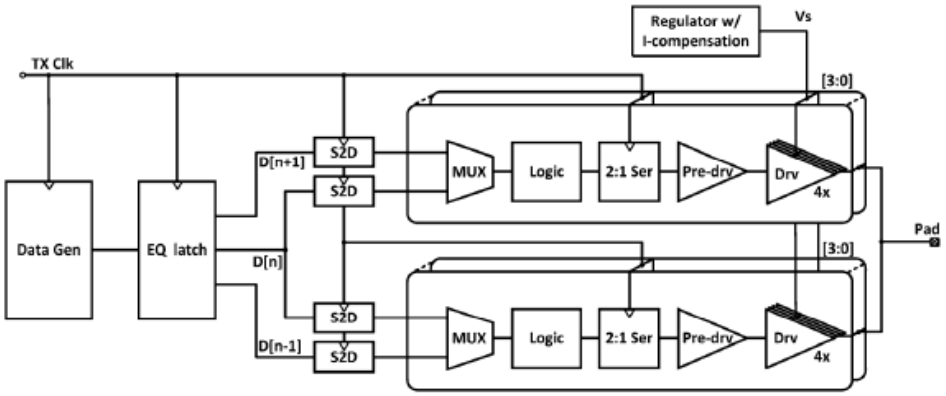
591

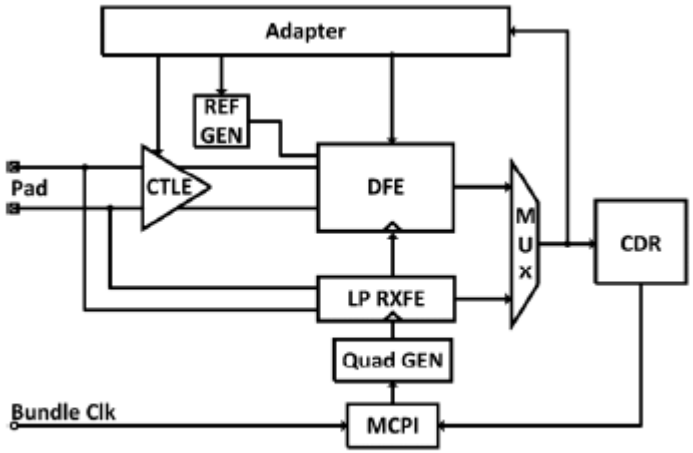
Claim 8	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

592

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.

593

Claim 8	Disclosure
594	 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 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.

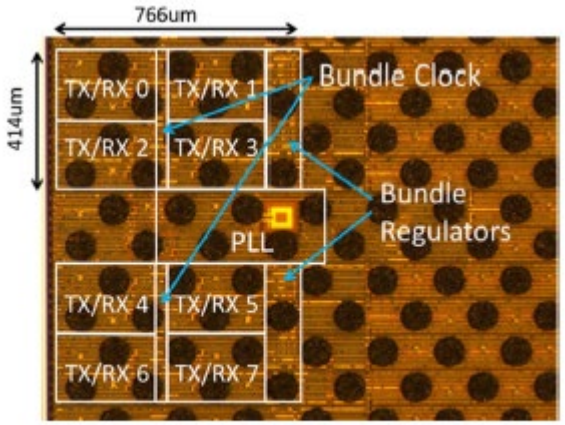
595

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.

596

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.

597

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.

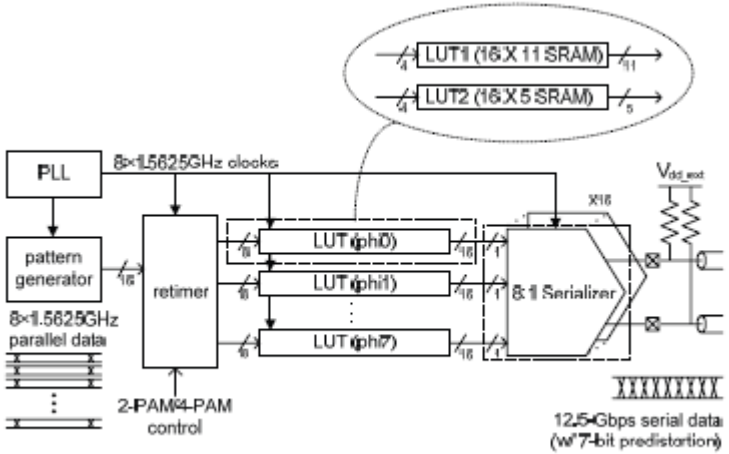
598

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.

599

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.

600

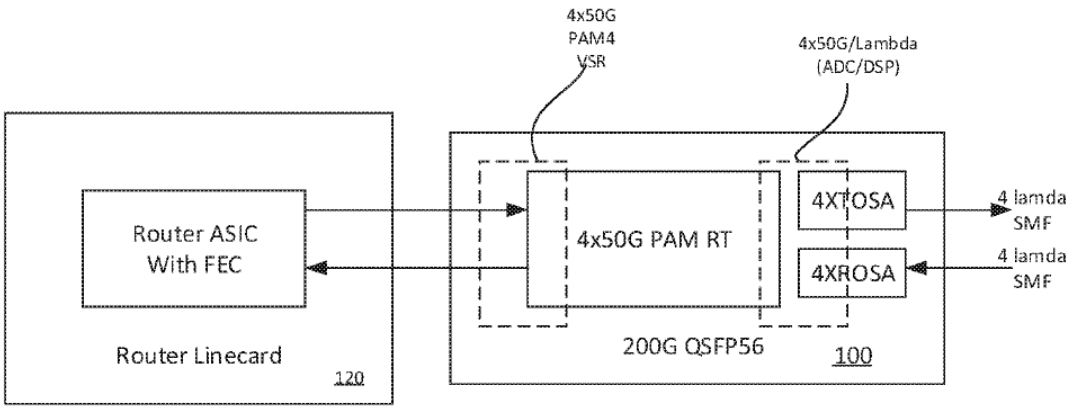
Claim 8	Disclosure
601	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

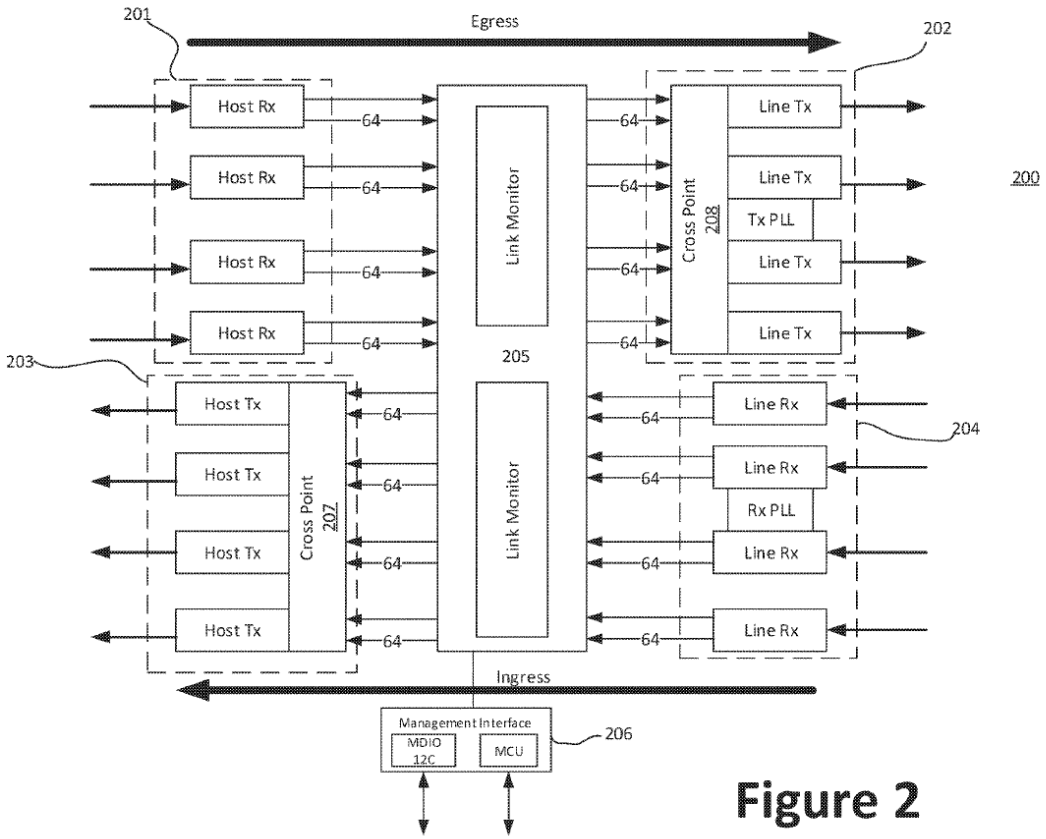
Claim 8	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

602

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.

603

Claim 8	Disclosure
604	 **Figure 1** Riani, Figure 1.

Claim 8	Disclosure
605	 Figure 2 Riani, Figure 2.

Claim 8	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.

609

Claim 8	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.

607

Claim 8	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.

608

Claim 8	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.

609

Claim 8	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.

610

Claim 8	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.

611

Claim 8	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.

612

Claim 8	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.

613

Claim 8	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.

614

	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.

615

Claim 8	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.

616

Claim 8	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.

617

Claim 8	Disclosure
618	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 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

619

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.

620

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.

621

Claim 8	Disclosure

622

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.

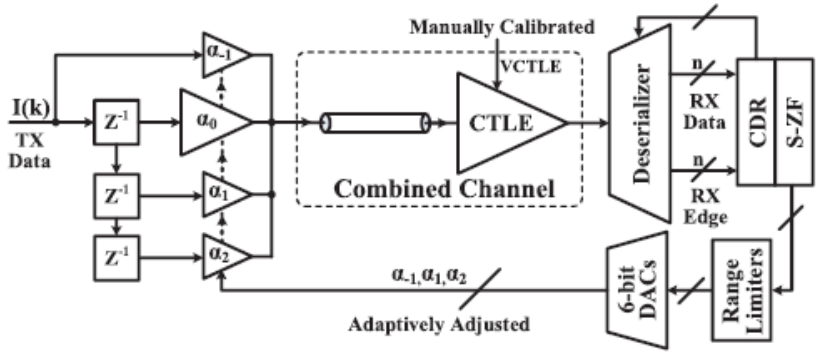
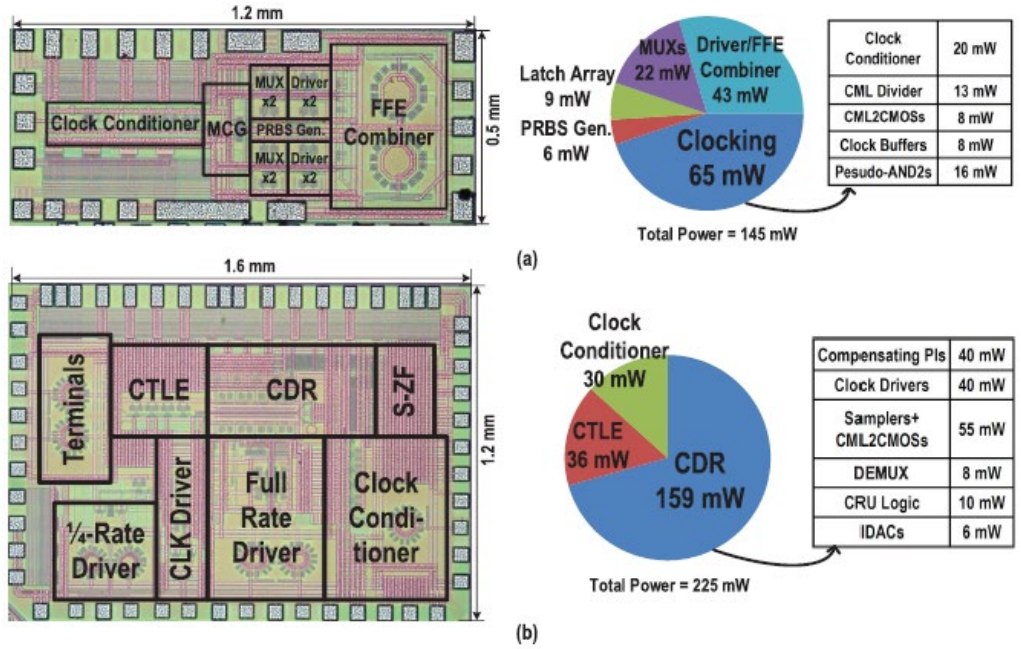


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

Claim 8	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.

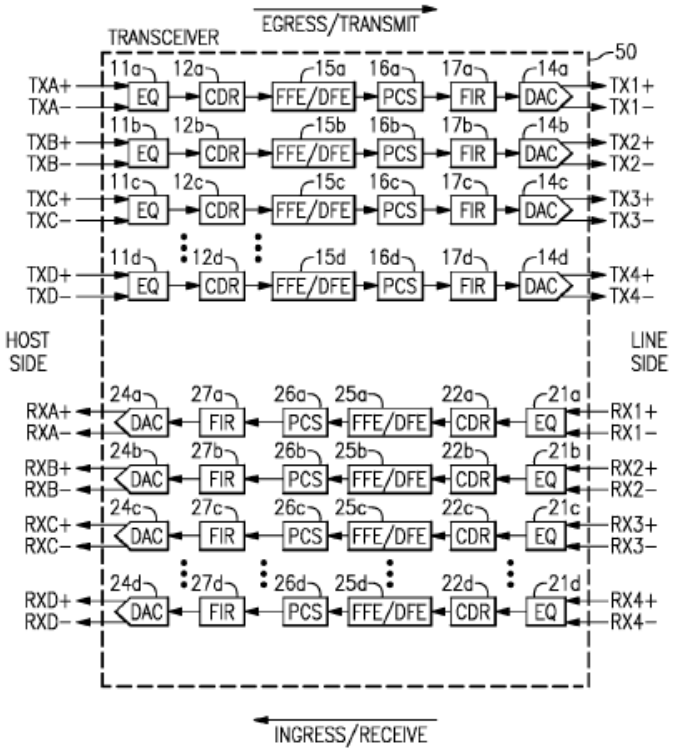
623

Claim 8	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

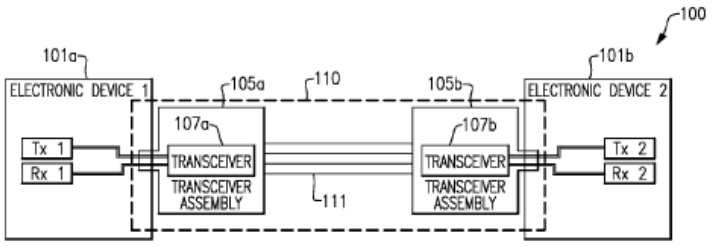
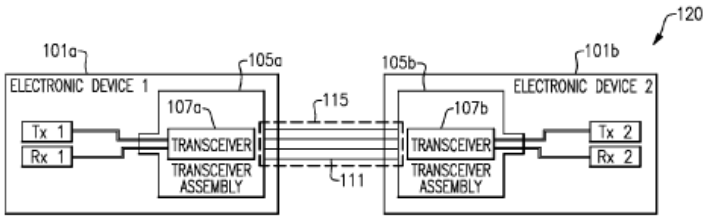
624

Claim 8	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.

625

Claim 8	Disclosure
	 FIG.1D Lugthart, Figure 1D.

626

Claim 8	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B. According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:66–4:5.

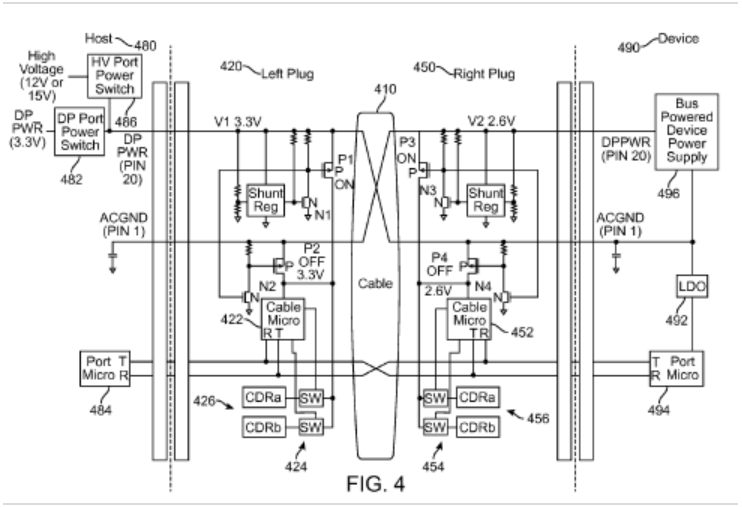
627

Claim 8	Disclosure
	FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. Lugthart, 12:1–53.

628

Claim 8	Disclosure
	In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:3–38.

629

Claim 8	Disclosure
	The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:19–23.  FIG. 4 Baker, Figure 4.

630

Claim 8	Disclosure
	FIG. 4 illustrates an electronic system according to an embodiment of the present invention. This figure includes host 480 coupled to device 490 via cable 410. Cable 410 includes left plug 424 connected to host 480, and right plug 454 connected to device 490. Host 480 may be capable of providing one or more voltages to both device 490 and cable circuitry in cable 410. Host 480 may provide a lower voltage of 3.3 V, or a higher voltage of 12 or 15 volts. In other embodiments of the present invention, host 480 may provide various voltage levels to host 490 and cable circuitry in cable 410. Baker, 5:26–36. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12. **[0031]** In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**.

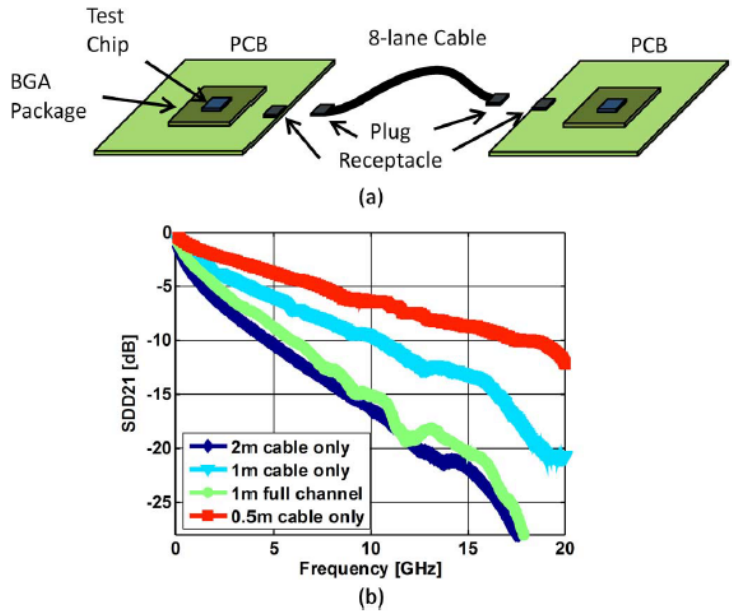
631

Claim 8	Disclosure
	Tremblay, paragraph 31.
8[c] connecting electrical conductors to the first and second DRR devices to convey electrical transit signals therebetween,	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

632

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.

633

Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

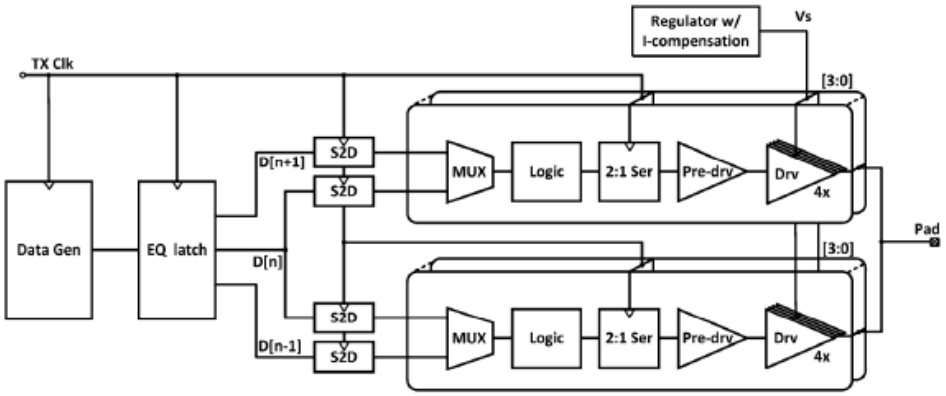
634

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.

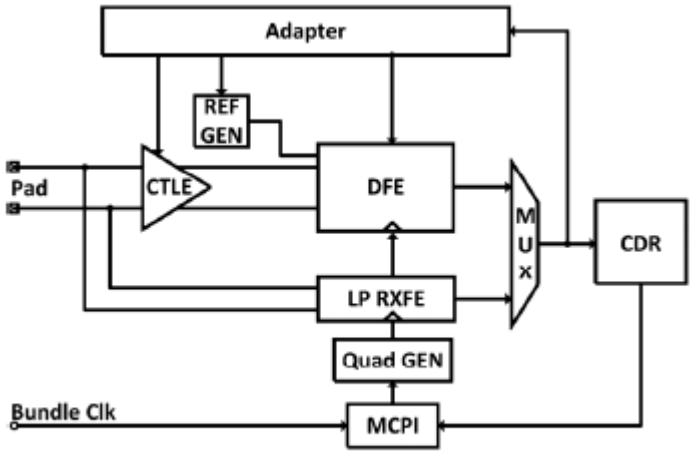
635

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.

636

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

637

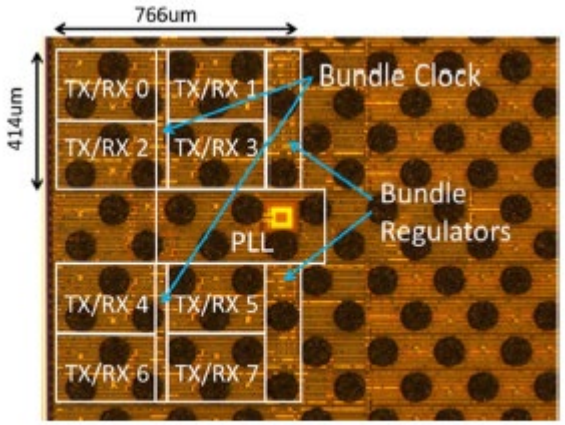
Claim 8	Disclosure
638	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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 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.

639

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.

640

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.

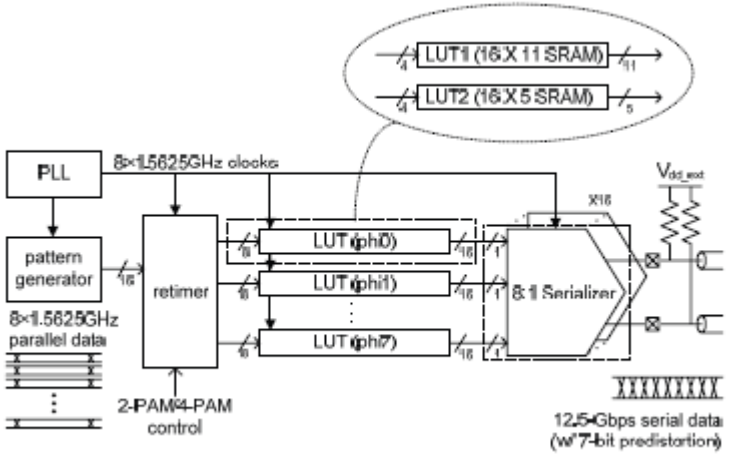
641

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.

642

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.

643

Claim 8	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

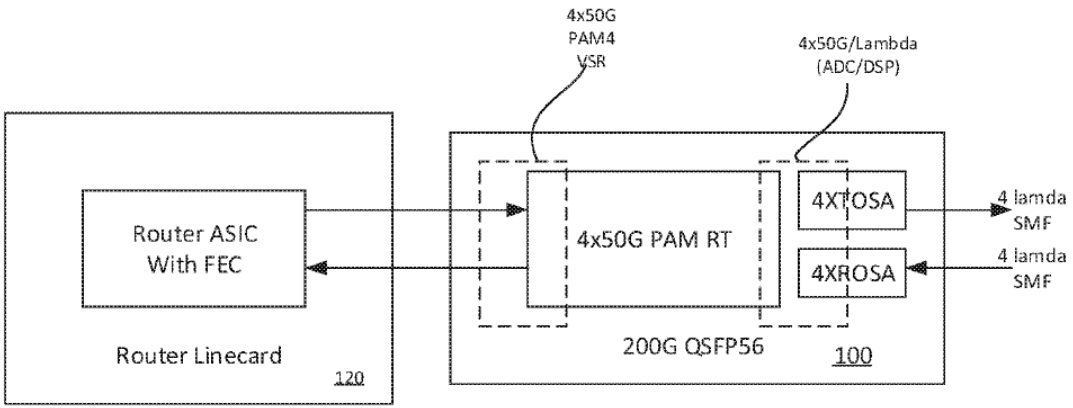
644

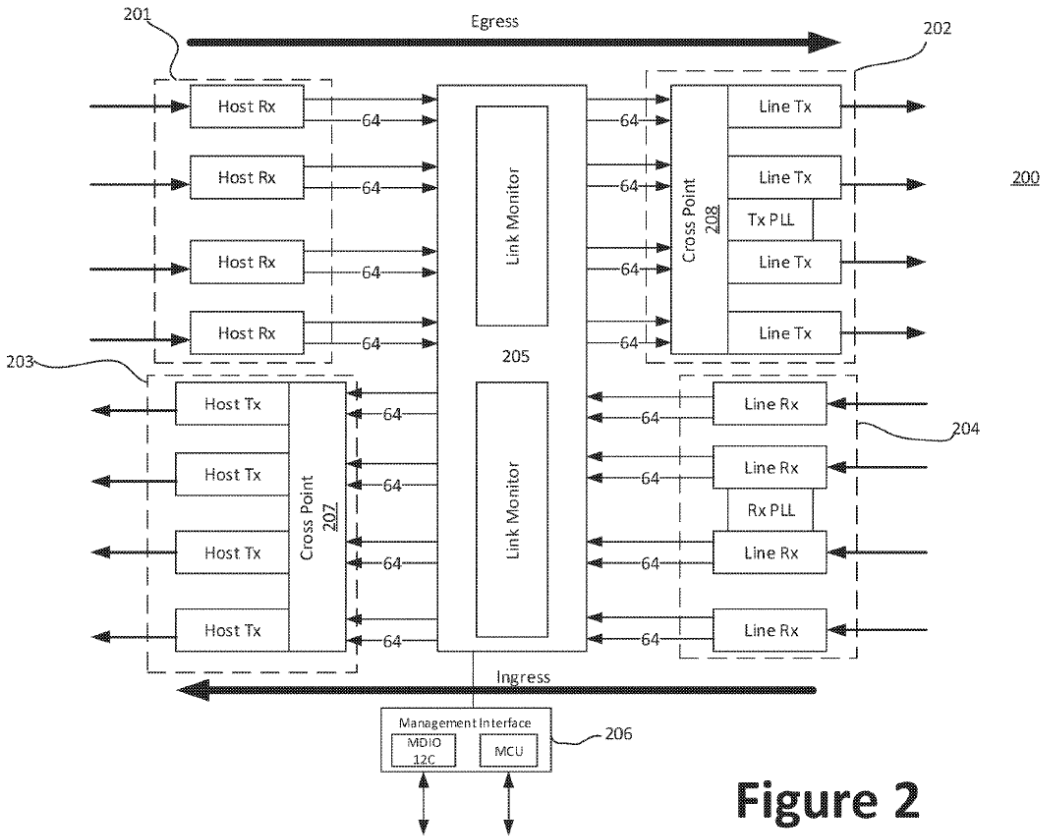
Claim 8	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

645

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.

646

Claim 8	Disclosure
647	 **Figure 1** Riani, Figure 1.

Claim 8	Disclosure
648	 Figure 2 Riani, Figure 2.

Claim 8	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.

649

Claim 8	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.

650

Claim 8	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.

651

Claim 8	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.

652

Claim 8	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.

653

Claim 8	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.

654

Claim 8	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.

655

Claim 8	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.

656

Claim 8	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.

657

	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.	

658

Claim 8	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.

659

Claim 8	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.

660

Claim 8	Disclosure
661	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 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

662

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.

663

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.

664

Claim 8	Disclosure

665

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.

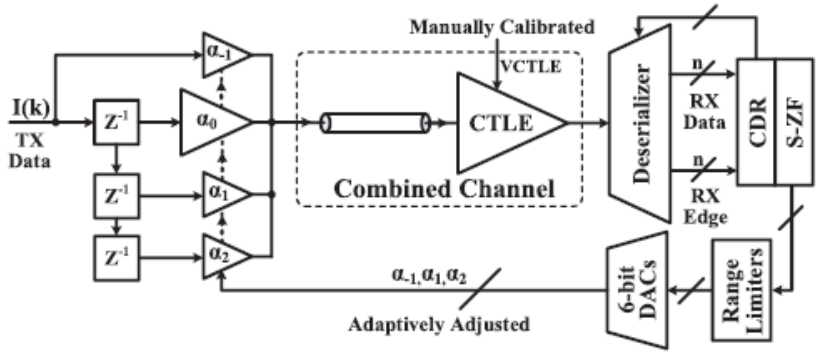
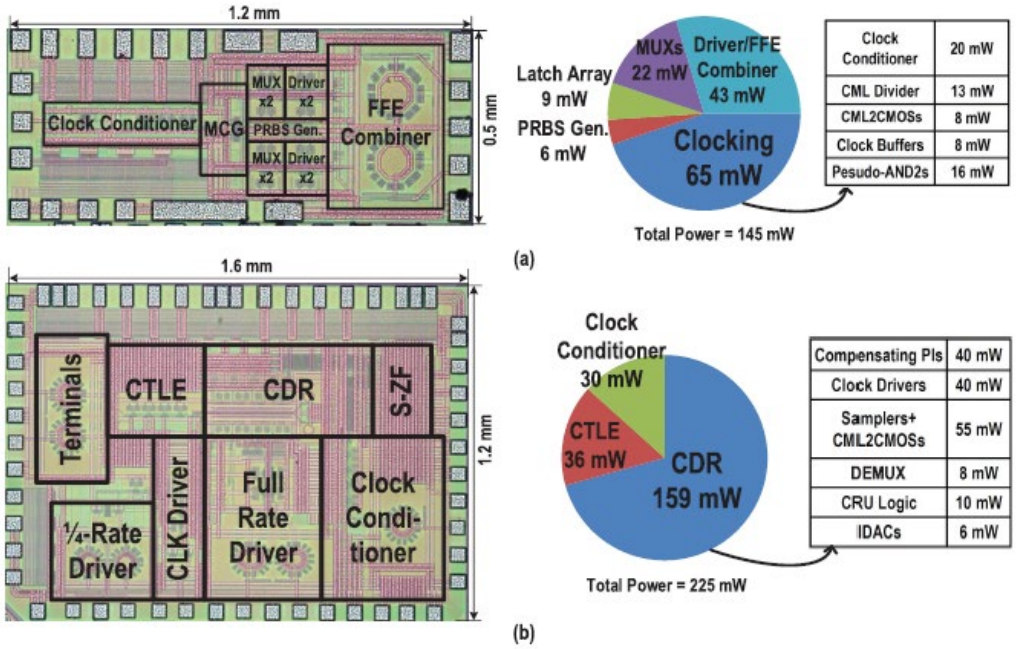


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

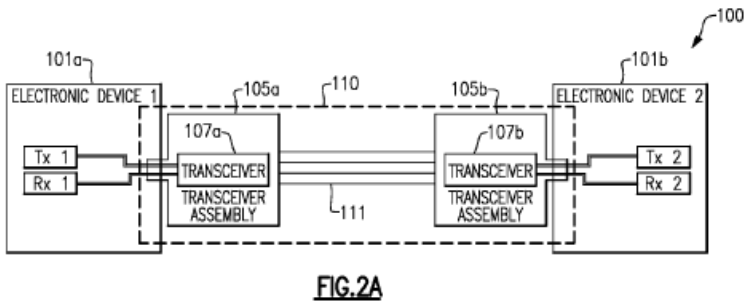
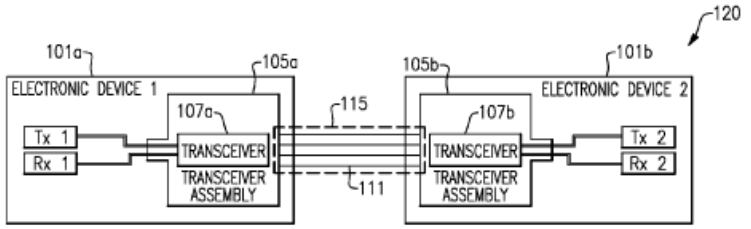
Claim 8	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

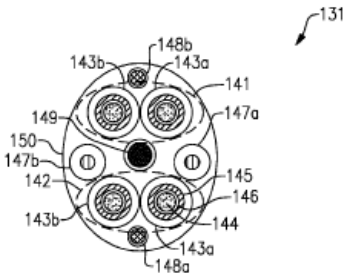
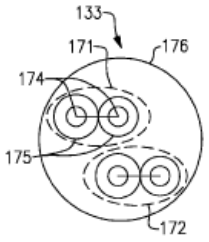
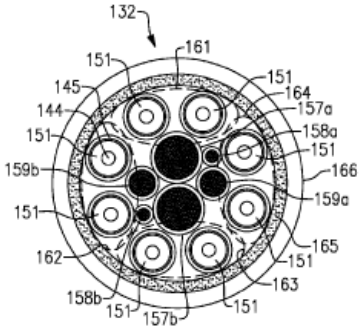
999

Claim 8	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

667

Claim 8	Disclosure
668	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.

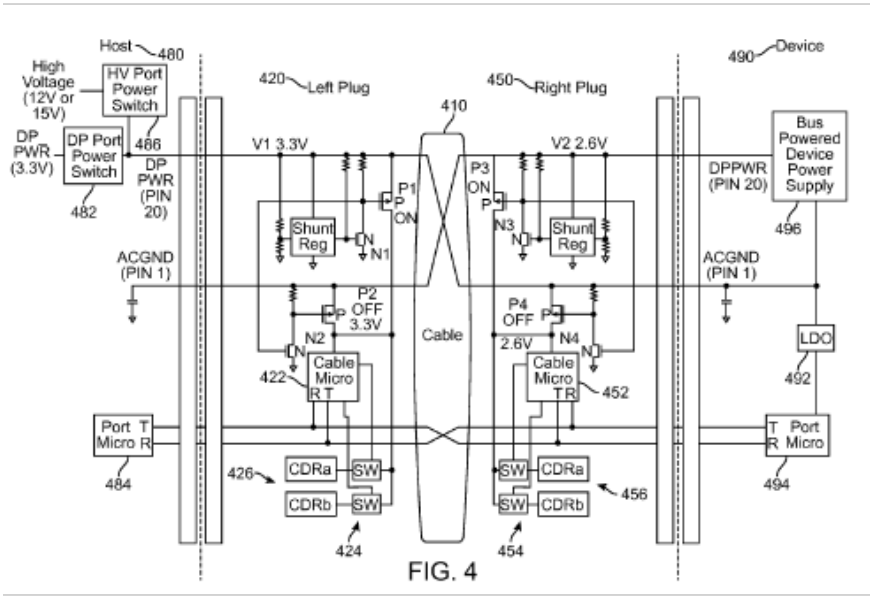
Claim 8	Disclosure
699	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

Claim 8	Disclosure
670	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3C.

Claim 8	Disclosure
	The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. Lugthart 14:3–20.

671

Claim 8	Disclosure
	In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:30–38.

Claim 8	Disclosure
	 FIG. 4 Baker, Figure 4. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry.

673

Claim 8	Disclosure
	In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12. **[0029]** Example embodiments will now be described with reference to FIG. **1** which illustrates a host **102** that is connected through a host channel **103** to a pluggable network transceiver interface module **100**, which in turn is connected to a transmission medium **104**. In the illustrated embodiment, the transmission medium **104** includes fibre optic cable **124** for transmitting high speed data from the interface module **100** and fibre optic cable **126** for transmitting high speed data to the interface module **100**, however embodiments described herein can also be applied to non-fibre transmission mediums such as copper for example. In the illustrated embodiment, the host **102** includes a network interface card or board **106** that allows data to be exchanged between the host **102** and the network interface module **100**. In at least one example embodiment, the host interface board **106** includes data processing circuitry, which could for example be implemented an ASIC **108** ("Application Specific Integrated Circuit"), and a SERDES ("SERializer-DESerializer") chip or module **110**. SERDES module **110** receives parallel data from the host ASIC **108** over a parallel data channel **170** and converts it into high speed serial data for transmission over the host channel **103** to the network interface module **100**, and receives high speed serial data from the network interface module **100** through the host channel **103** which it then converts into parallel data and passes on to the ASIC **108** over the parallel data channel **170**. Tremblay, paragraph 29.

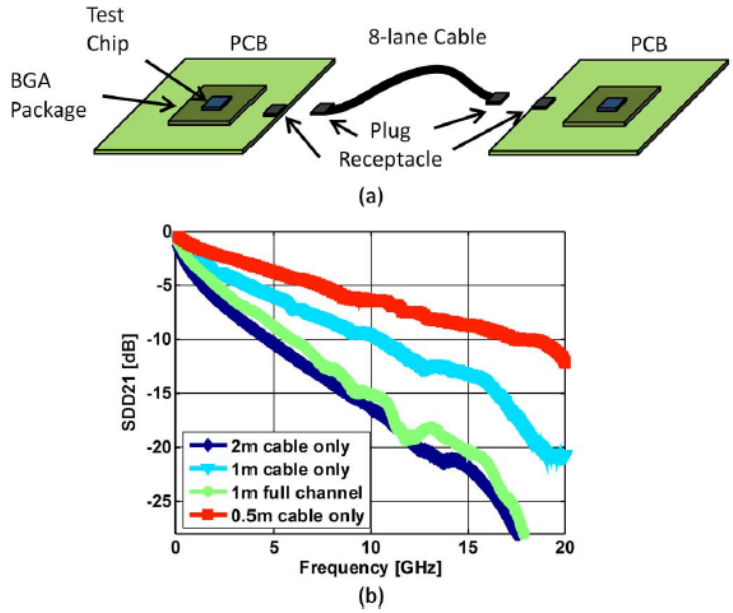
674

Claim 8	Disclosure
8[d] the first DRR device converting between said electrical transit signals and said multi-lane data streams for the first host interface port, 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

675

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.

676

Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

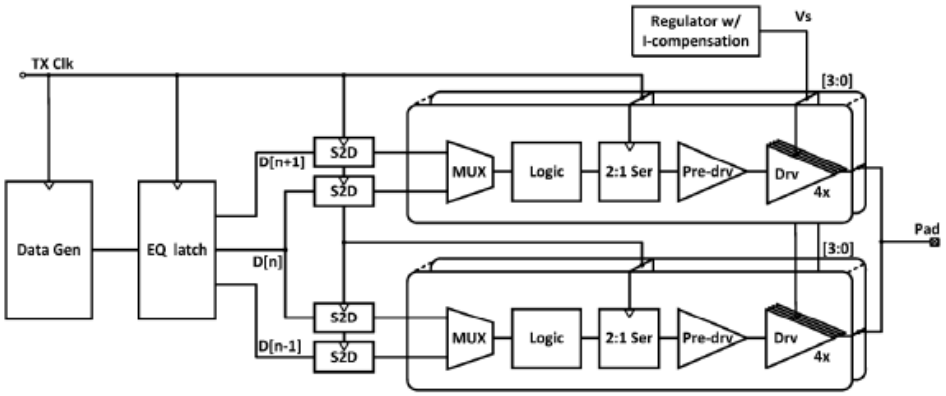
677

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.

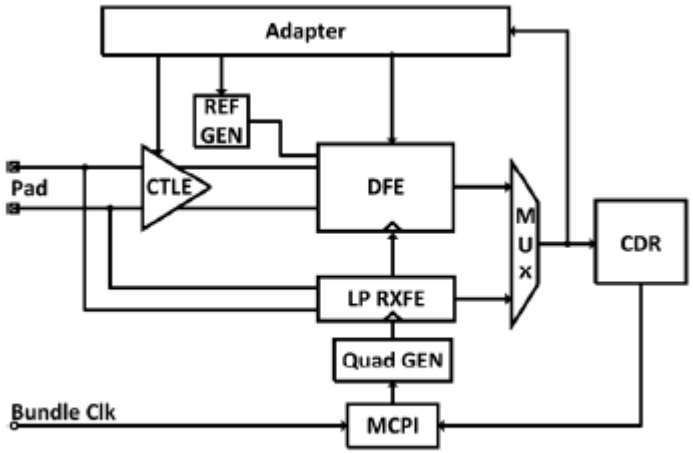
678

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.

679

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

680

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.

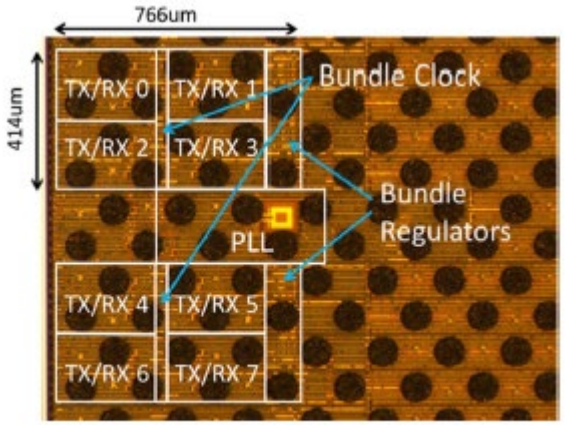
681

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.

682

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.

683

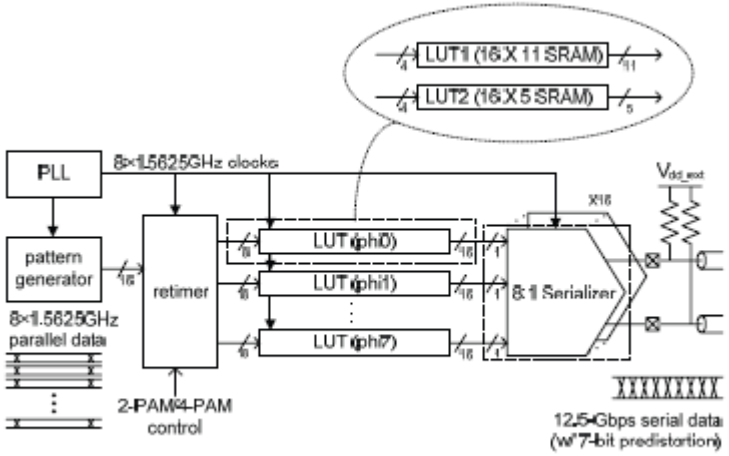
Claim 8	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

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.

685

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.

689

Claim 8	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

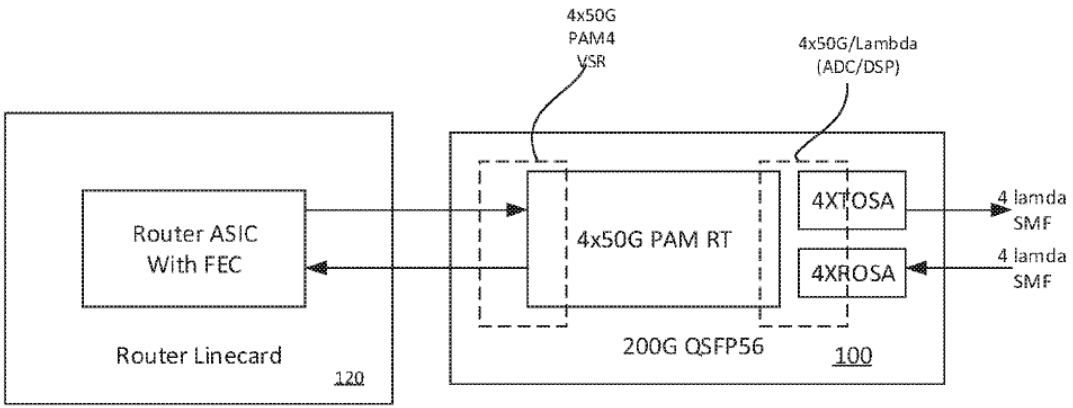
687

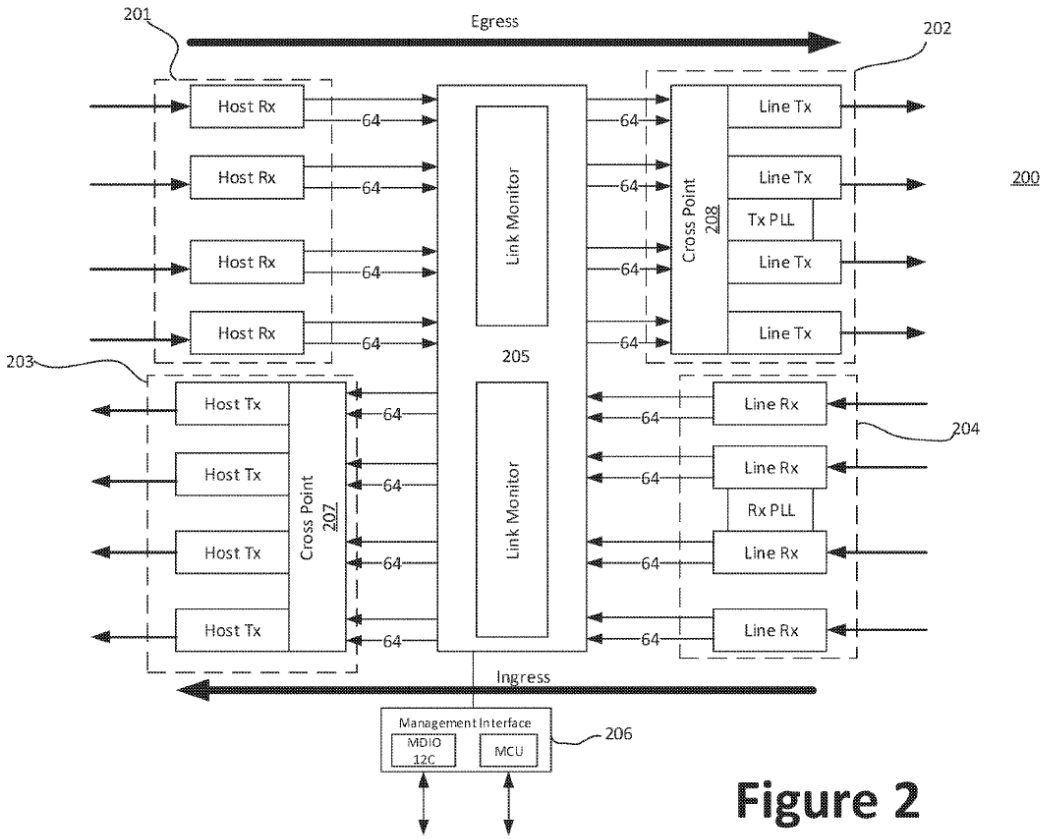
Claim 8	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

688

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.

689

Claim 8	Disclosure
069	 **Figure 1** Riani, Figure 1.

Claim 8	Disclosure
	 Figure 2 Riani, Figure 2.

691

Claim 8	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.

692

Claim 8	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.

693

Claim 8	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.

694

Claim 8	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.

695

Claim 8	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.

969

Claim 8	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.

697

Claim 8	Disclosure
869	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 8	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.

669

Claim 8	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.

700

	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.	

701

Claim 8	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.

Claim 8	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.

703

Claim 8	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.

704

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

705

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.

706

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.

707

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.

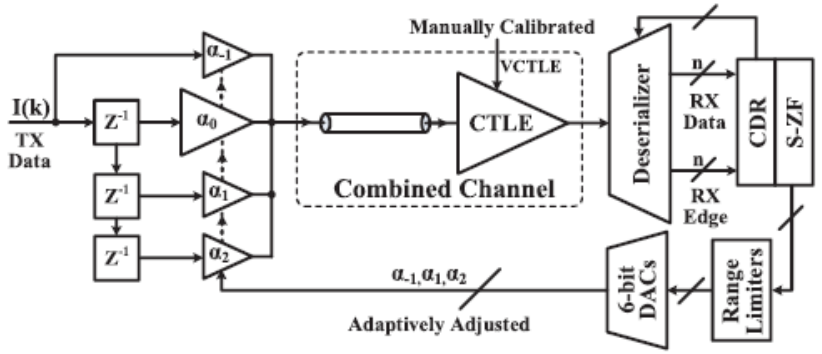
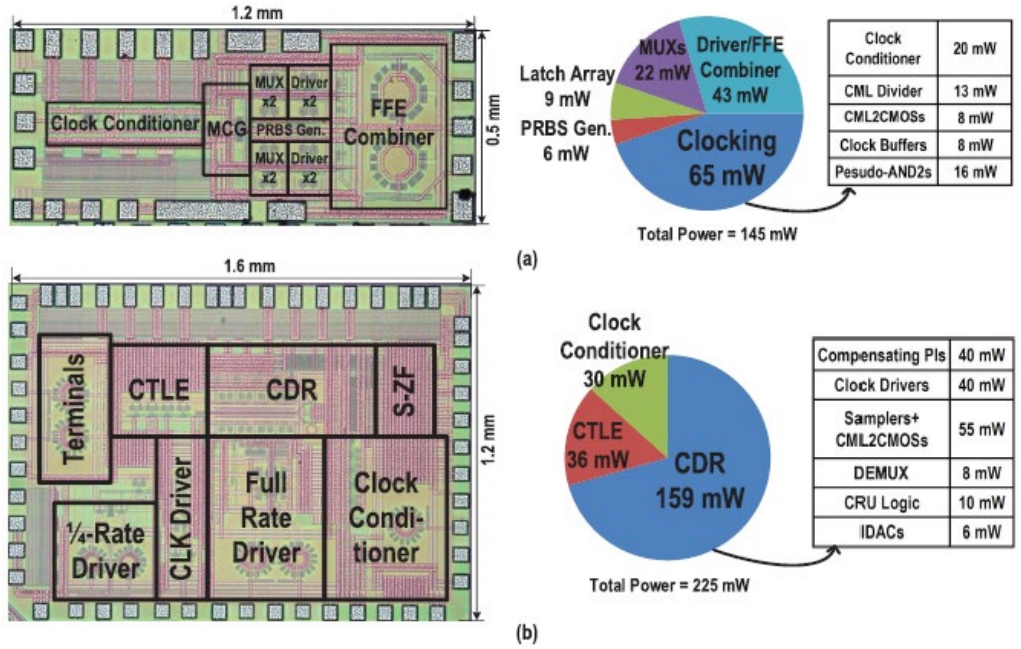


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

708

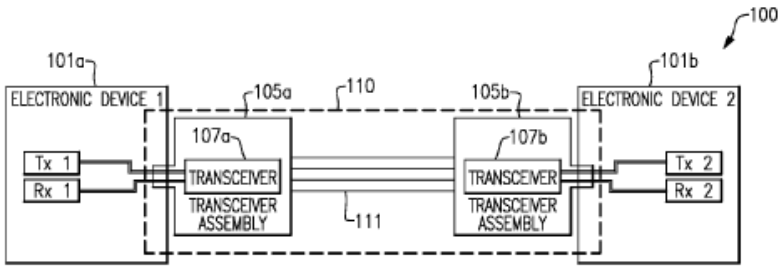
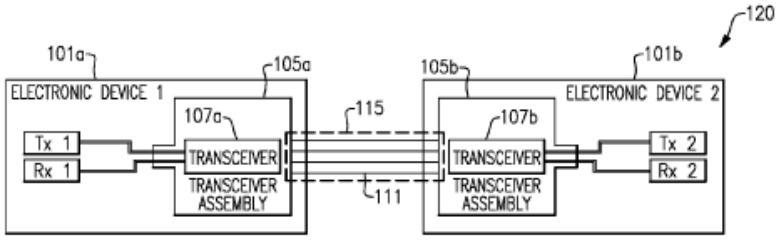
Claim 8	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying a 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.

709

Claim 8	Disclosure
710	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

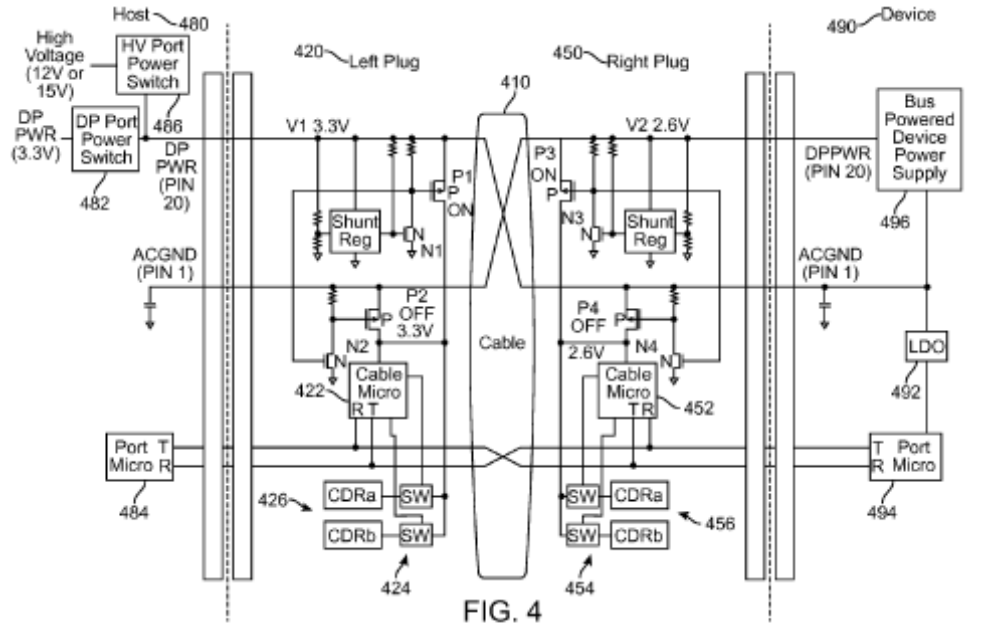
Claim 8	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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.

711

Claim 8	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

712

Claim 8	Disclosure
	In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

Claim 8	Disclosure
	 FIG. 4 Baker, Figure 4.

Claim 8	Disclosure
	Voltage V1 is received by low-drop-out regulator 492 in device 490, which provides power to port microcontroller 494. The host port microcontroller 484 may then communicate with cable microcontrollers s 422 and 452 and port microcontroller 494 to determine proper configuration for the cable. In a specific embodiment of the present invention, the host port microcontroller 484 may check with device port microcontroller 494 to determine whether device 490 requires a higher level of power. If it does, host port microcontroller 484 may check with cable microcontroller 422 to determine whether the cable can support the delivery of this higher level of power. If device 490 requires higher power, and the cable can deliver it, then host 480 may provide the higher level of power. In another embodiment of the present invention, host port microcontroller 484 may determine how much power will be needed by the cable and device 490. In some circumstances, one link including one pair of clock and data recovery circuits, or other circuitry, may need to be powered off. Baker, 5:49–67.

715

Claim 8	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. **[0032]** The receive section **122** includes a photo diode module **148**, a TIA ("TransImpedance Amplifier") module **146** and a LA/CDR ("Limiting Amplifier/Clock and Data Recovery") module **144**. The photo diode module **148** receives high speed optical data from fibre optic cable **126** and converts the optical data into high speed electrical data that is further converted from current signals into voltage signals by TIA module **146**. The high speed data from the TIA module **146** is processed by LA/CDR module **144** and then sent over host channel **103** as high speed serial data to a receiver portion of SERDES **110**. In embodiments having multiple channels in the form of multiple transmitting conductors **112**, the LA/CDR module **144** performs these operations for each channel of the conductors **112**. Tremblay, paragraphs 31-32.
8[e] the second DRR device converting between said electrical transit signals and said multi-lane data streams	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:

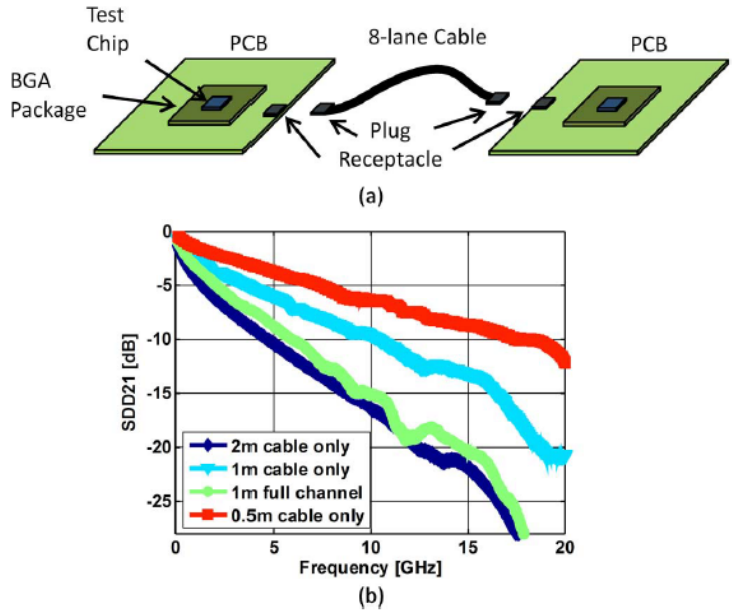
716

Claim 8	Disclosure
for the second host interface port,	***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

717

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.

718

Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

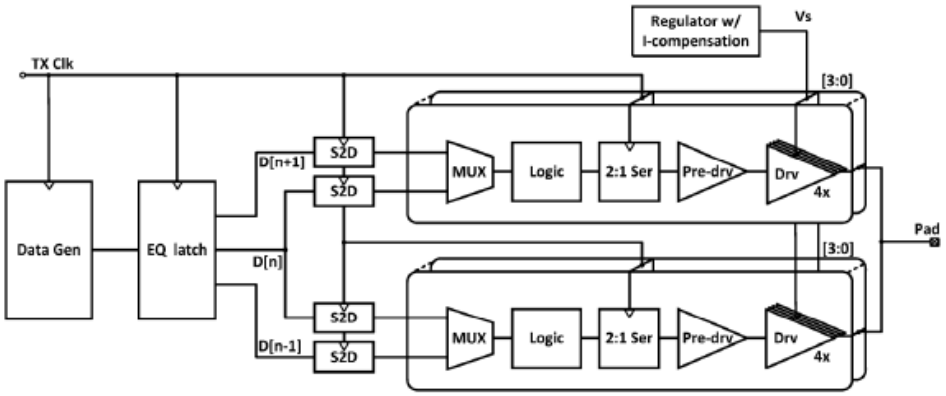
719

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.

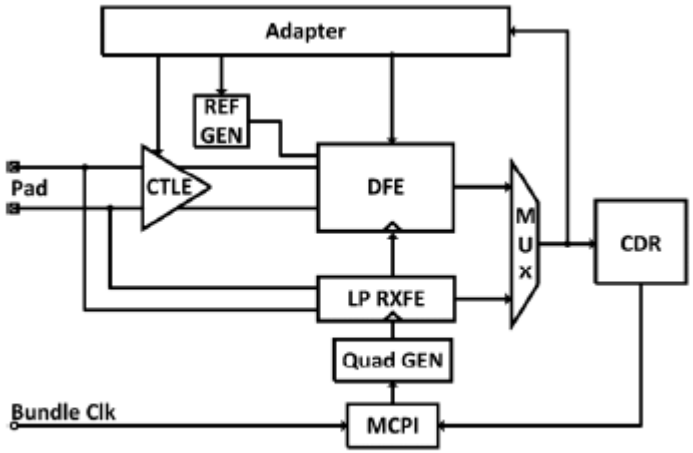
720

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.

721

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

722

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.

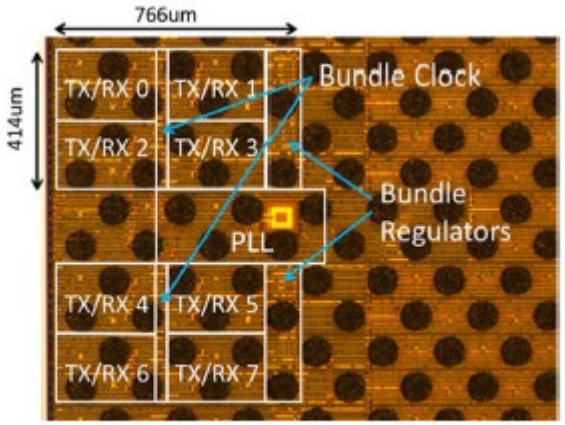
723

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.

724

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.

725

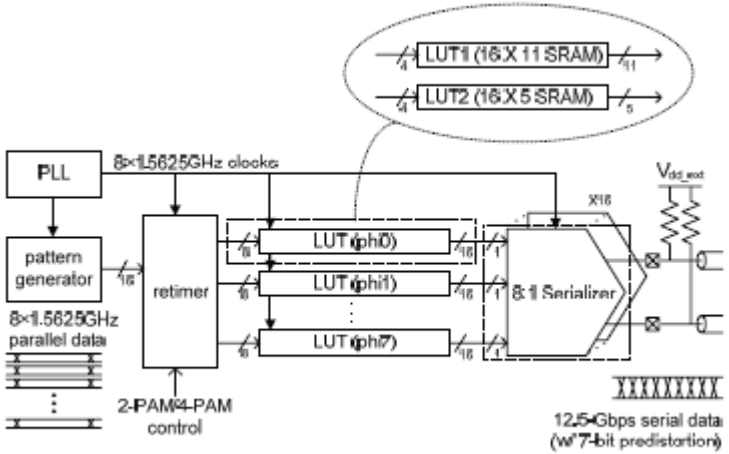
Claim 8	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

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.

727

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.

728

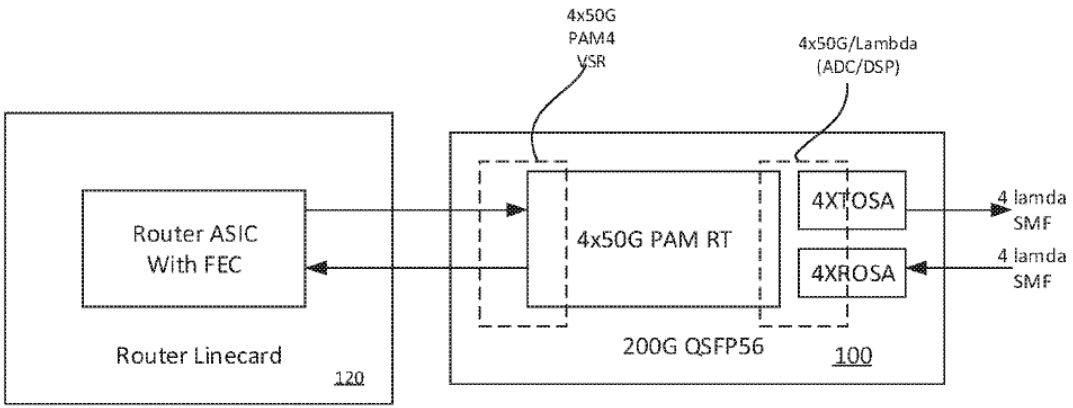
Claim 8	Disclosure
729	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

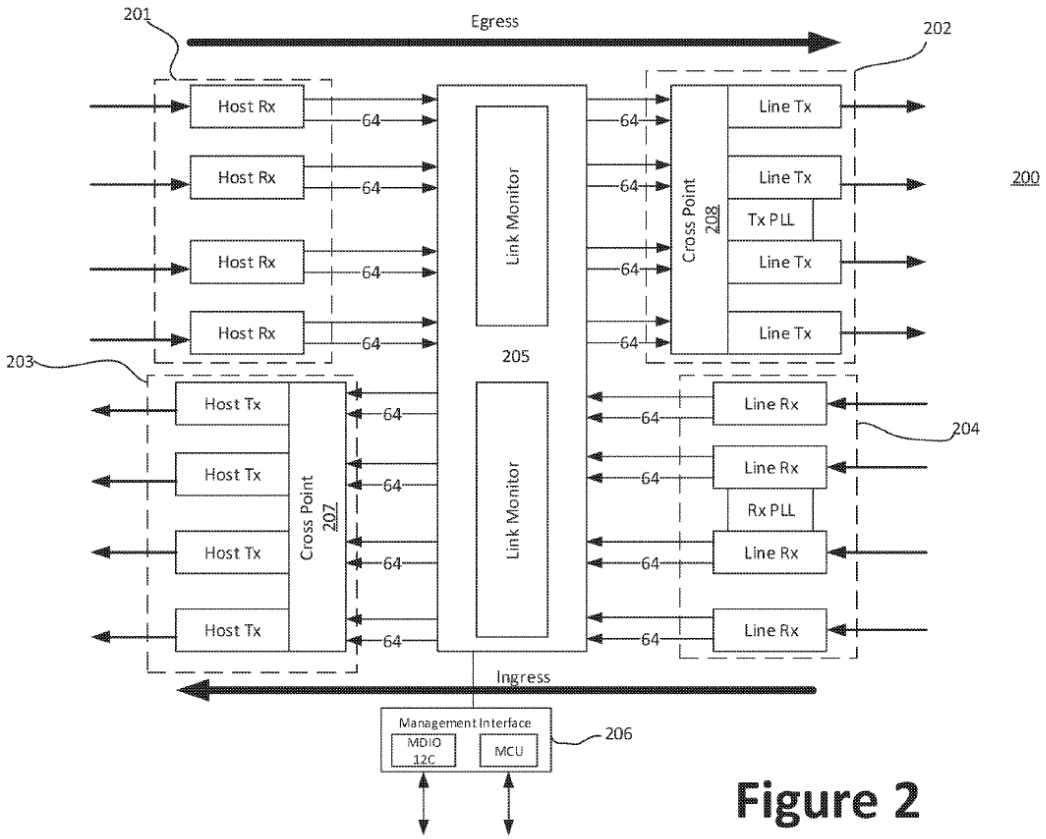
Claim 8	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

730

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.

731

Claim 8	Disclosure
732	 **Figure 1** Riani, Figure 1.

Claim 8	Disclosure
733	 Figure 2 Riani, Figure 2.

Claim 8	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.

734

Claim 8	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.

735

Claim 8	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.

736

Claim 8	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.

737

Claim 8	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.

738

Claim 8	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.

739

Claim 8	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.

740

Claim 8	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.

741

Claim 8	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.

	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.	

743

Claim 8	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.

744

Claim 8	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.

745

Claim 8	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.

746

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

747

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.

748

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.

749

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.

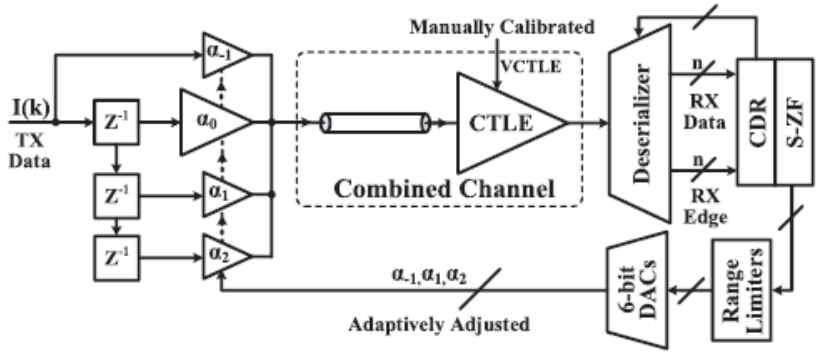
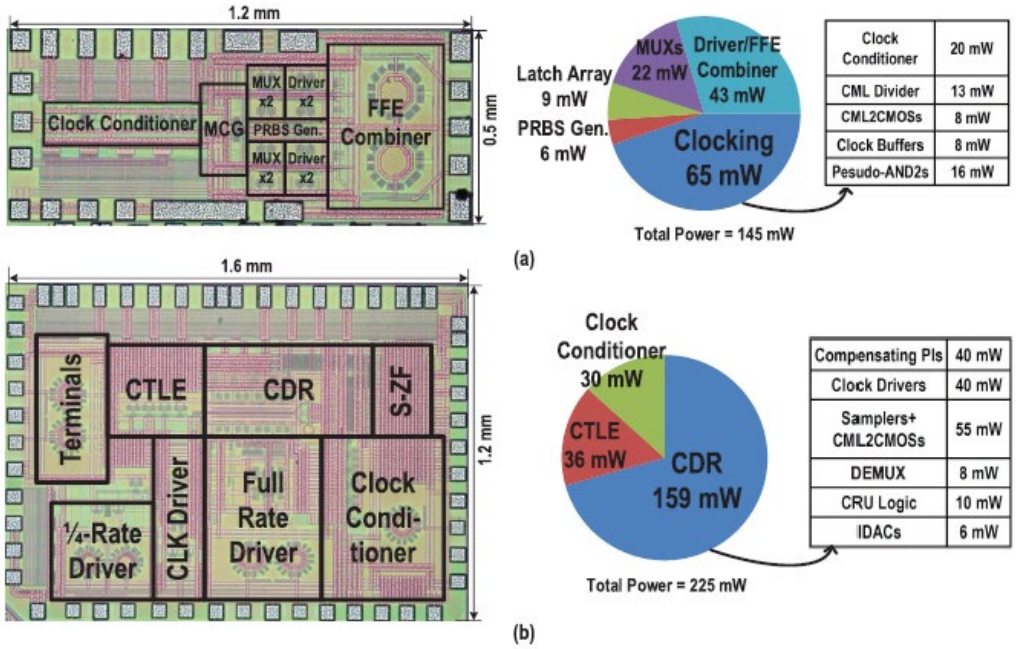


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

750

Claim 8	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.

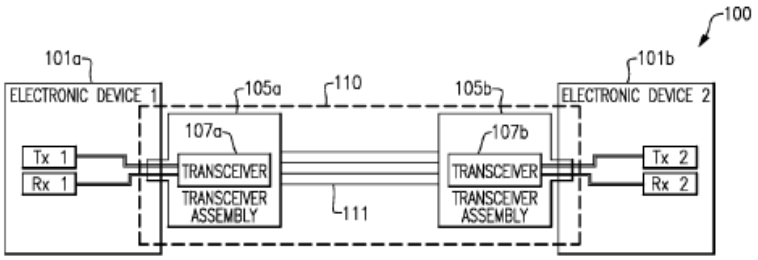
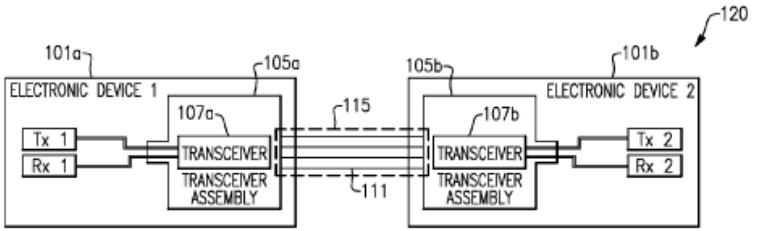
751

Claim 8	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

752

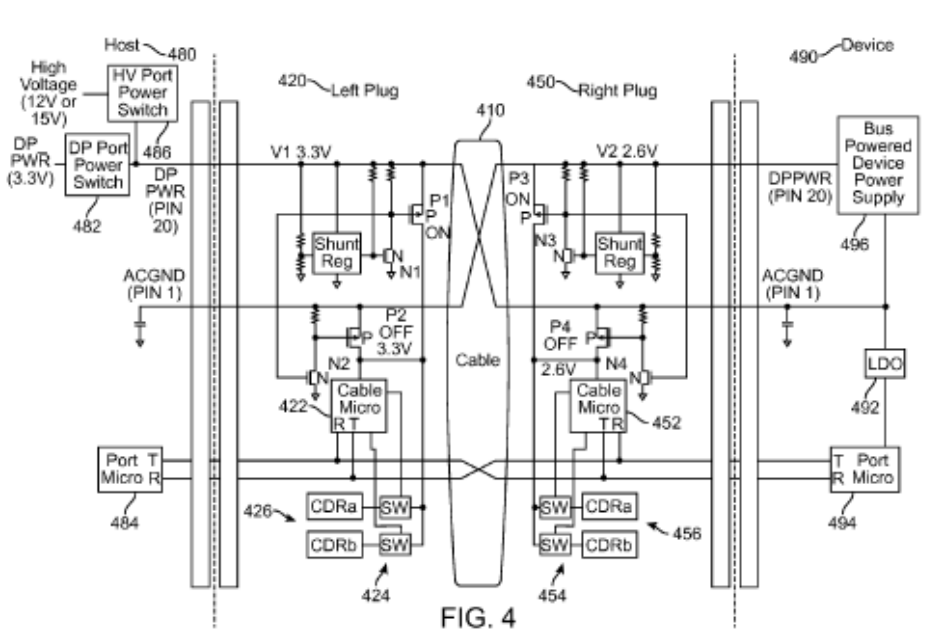
Claim 8	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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.

753

Claim 8	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a

754

Claim 8	Disclosure
	releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

Claim 8	Disclosure
	 FIG. 4 Baker, Figure 4. Voltage V1 is received by low-drop-out regulator 492 in device 490, which provides power to port microcontroller 494. The host port microcontroller 484 may then communicate with cable

Claim 8	Disclosure
	microcontrollers s 422 and 452 and port microcontroller 494 to determine proper configuration for the cable. In a specific embodiment of the present invention, the host port microcontroller 484 may check with device port microcontroller 494 to determine whether device 490 requires a higher level of power. If it does, host port microcontroller 484 may check with cable microcontroller 422 to determine whether the cable can support the delivery of this higher level of power. If device 490 requires higher power, and the cable can deliver it, then host 480 may provide the higher level of power. In another embodiment of the present invention, host port microcontroller 484 may determine how much power will be needed by the cable and device 490. In some circumstances, one link including one pair of clock and data recovery circuits, or other circuitry, may need to be powered off. Baker, 5:49–67.

757

Claim 8	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. **[0032]** The receive section **122** includes a photo diode module **148**, a TIA ("TransImpedance Amplifier") module **146** and a LA/CDR ("Limiting Amplifier/Clock and Data Recovery") module **144**. The photo diode module **148** receives high speed optical data from fibre optic cable **126** and converts the optical data into high speed electrical data that is further converted from current signals into voltage signals by TIA module **146**. The high speed data from the TIA module **146** is processed by LA/CDR module **144** and then sent over host channel **103** as high speed serial data to a receiver portion of SERDES **110**. In embodiments having multiple channels in the form of multiple transmitting conductors **112**, the LA/CDR module **144** performs these operations for each channel of the conductors **112**. Tremblay, paragraphs 31-32.
8[f] the first and second DRR devices providing pre-equalization of the electrical transit signals using transmit filter	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:

758

Claim 8	Disclosure
coefficient values stored in nonvolatile memories.	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.

759

Claim 8	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.

760

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.

761

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.

762

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.

763

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.

764

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.

765

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.

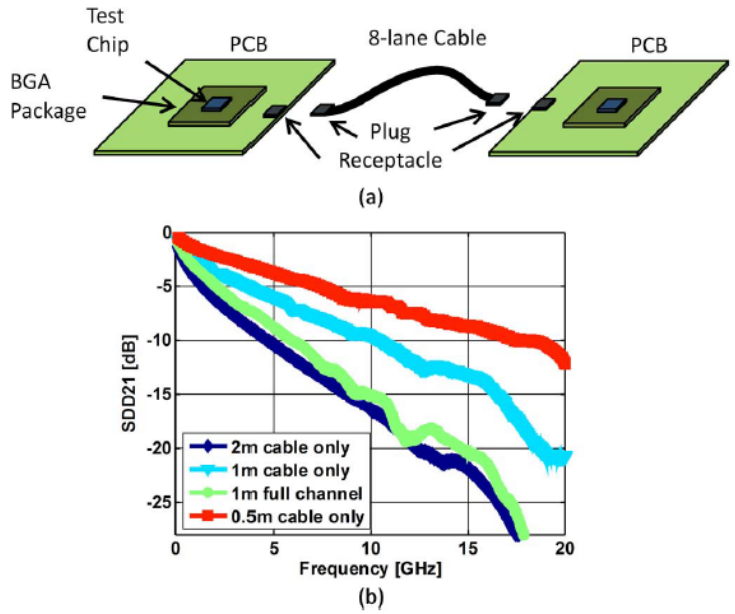
766

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

767

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.

768

Claim 8	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

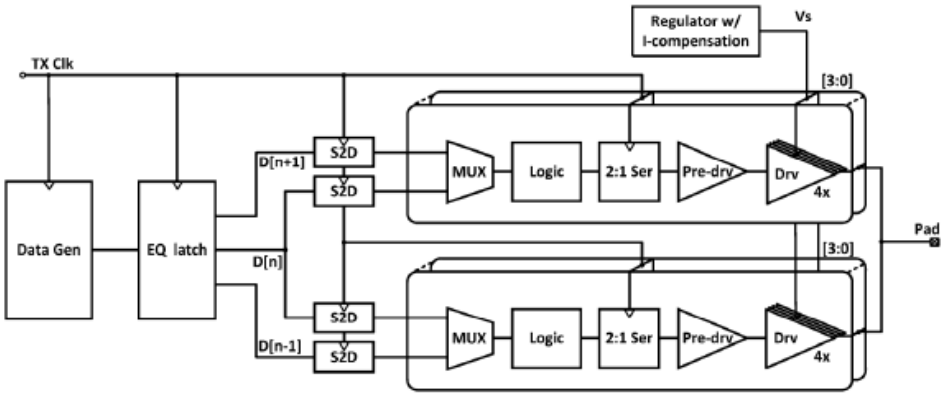
769

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.

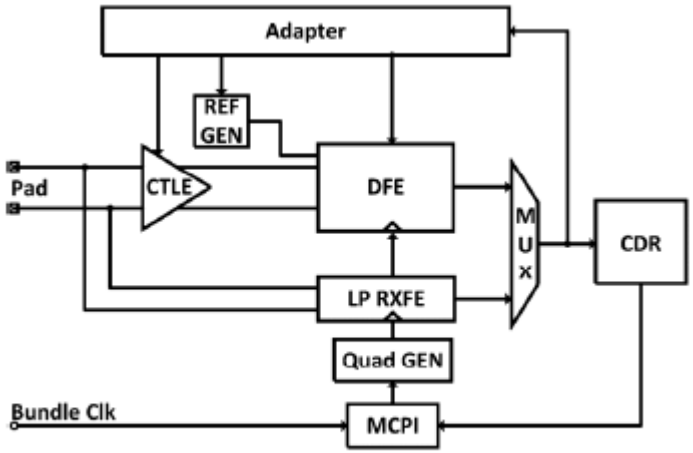
770

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.

771

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

772

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.

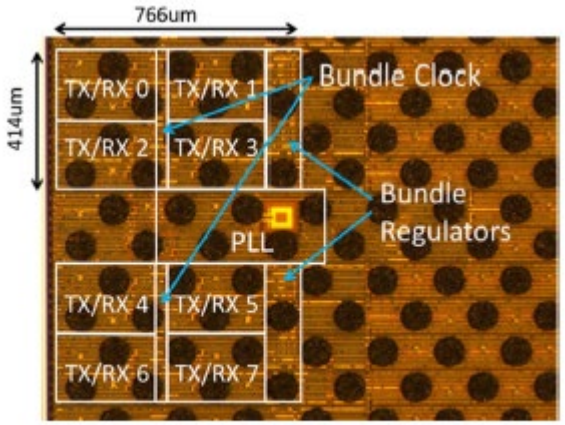
773

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.

774

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.

775

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.

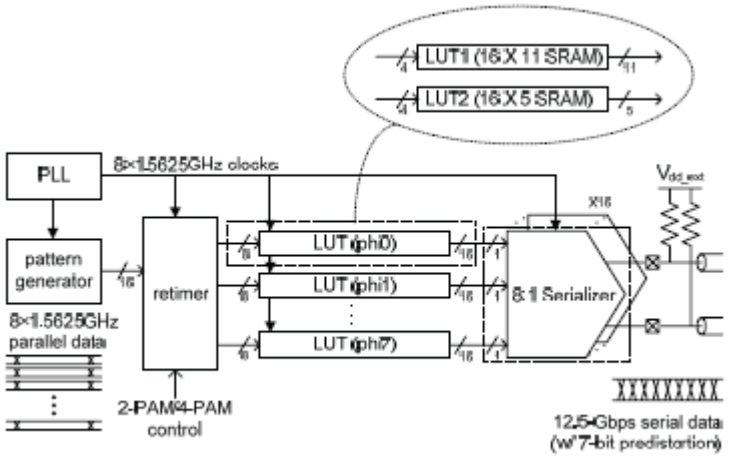
776

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.

777

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.

778

Claim 8	Disclosure
779	 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

780

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.

781

Claim 8	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.

782

Claim 8	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.

783

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

784

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.

785

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.

786

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.

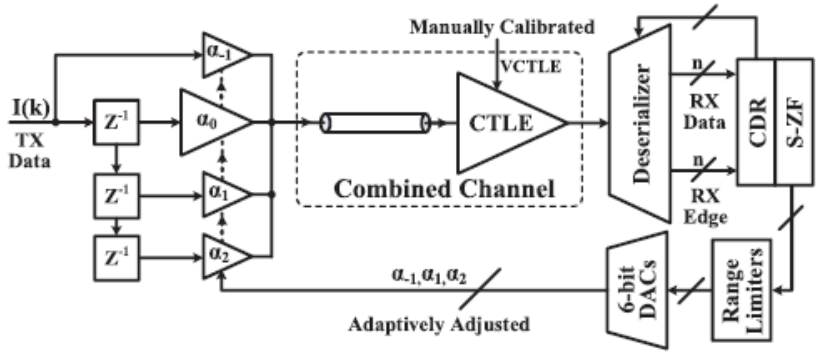
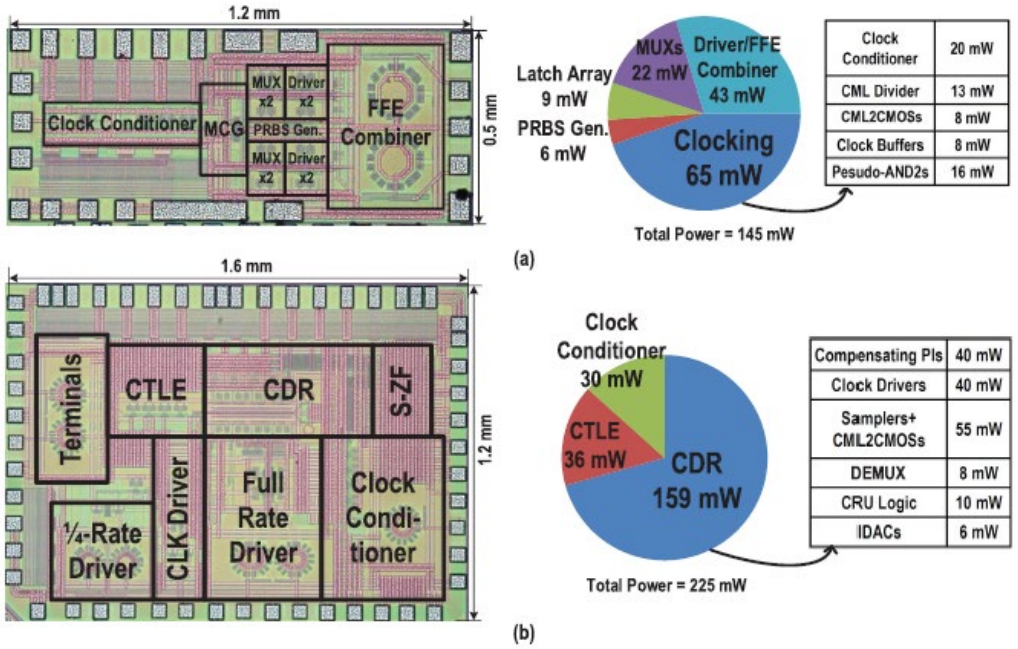


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

787

Claim 8	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.

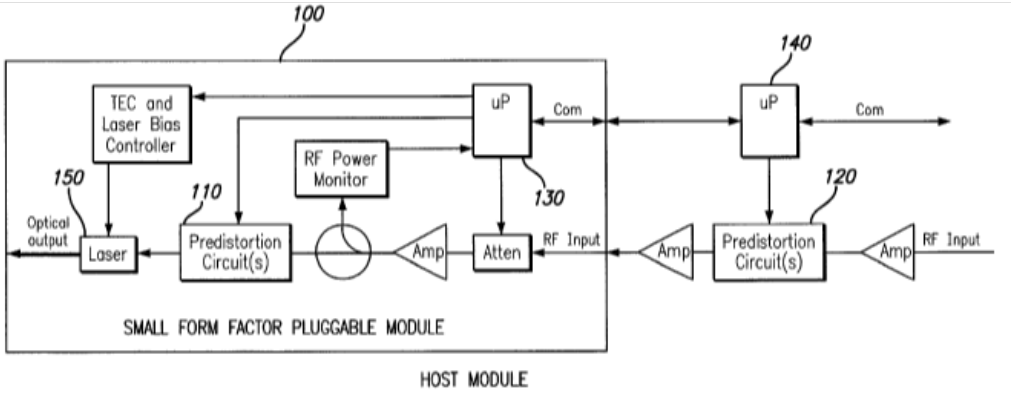
788

Claim 8	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

789

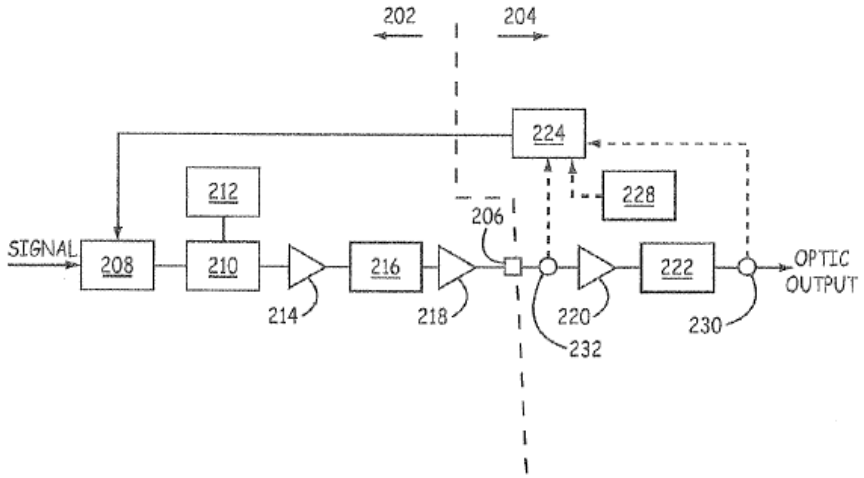
Claim 8	Disclosure
790	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By 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. 7,233,617 to Gorecki* Gorecki discloses that "tap weights or coefficients (values) may be determined or controlled[.]" *See* Gorecki at 6:56-57. Gorecki discloses the storing of filter coefficient values in EEPROM, and use those values to set filter coefficients to provide pre-equalization. *See* Gorecki at 7:18-33 ("[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM. In this way, the transmitter may access the memory to retrieve the necessary information during start-up/power-up, initialization or re-initialization."). An EEPROM [electrically erasable programmable read-only memory] is non-volatile memory. *See, e.g.*, U.S. Patent App. Pub. No. 2015/0106536 at [0081] ("nonvolatile memory can include… electrically erasable PROM (EEPROM)"); U.S. Patent No. 8,000,176 at 2:53-54 ("EEPROM is a non-volatile memory"); U.S. Patent App. Pub. No. 2005/0013317 at [0012] ("non-

Claim 8	Disclosure
	volatile memory may include… an EEPROM"). The L-G combination teaches programming the coefficients of the FIR filter 17 using the values stored in non-volatile EEPROM to pre-equalize the TX signal in transceiver 107a and 107b.

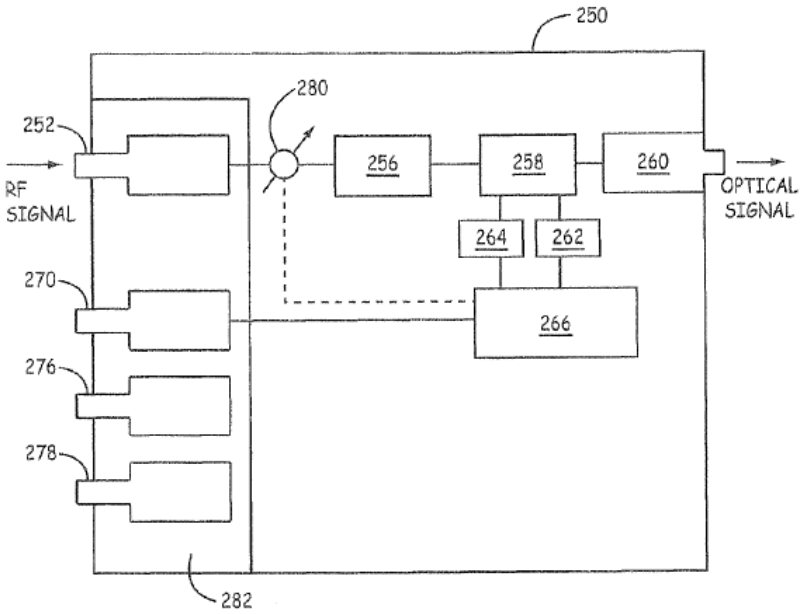
Claim 8	Disclosure
	 FIG. 1 Ishaug, Figure 1.

Claim 8	Disclosure
	A method includes reading operation parameters from a non-volatile memory located in a pluggable module that is coupled to a host module; processing the operational parameters with a processor located in the host module to control operation of a predistortion circuit located in the host module; adding predistortion to a signal with the predistortion circuit located in the host module and then sending the predistorted signal to the pluggable module. An apparatus includes a host module including a predistortion circuit and a processor coupled to the predistortion circuit; and a pluggable module coupled to the host module, wherein the pluggable module includes a non-volatile memory containing operational parameters for the predistortion circuit of the host module, wherein the operational parameters are processed by the processor of the host module to control the predistortion circuit of the host module. Ishaug, Abstract. As shown in FIG. 1, both the pluggable module and the host module have predistortion circuits 110, 120. However, the pluggable module does not need to contain a predistortion circuit. The pluggable module includes a processor 130 having non-volatile memory. In addition, the host module is shown as a standalone unit including a processor 140 having non-volatile memory. However, the host may be incorporated into a larger system that provides the analog RF input. There is also gain, attenuation and RF power monitoring shown in the pluggable and host modules that are not required, but may be needed to set the drive level to the laser modulation at an appropriate level. Finally, the optical transmitter is shown as a laser 150. The laser may be a directly modulated laser or an externally modulated laser in which the output is modulated using techniques such as electro-absorption modulation, Mach-Zender modulation or any appropriate technique for producing analog modulated optical signal. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion

793

Claim 8	Disclosure
	circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:34–5:4.  FIG. 3A Piehler, Figure 3A.

794

Claim 8	Disclosure
795	 FIG. 4 Piehler, Figure 4. Linearized optical transmitter units are described for a hybrid optical fiber coaxial cable network. The linearized optical transmitter unit can comprise a directly-modulated or externally-modulated laser optically coupled to an optical conduit directed to an optical fiber communications link and electrically coupled to an electrical RF source line that provides an RF source to drive the laser or an external modulator for a light beam from the laser. A linearization information electrical component comprising memory and/or a processor, and a data output configured to transmit linearization enabling data for input into a direct digital synthesis engine that enables the direct digital synthesis engine to generate an RF

Claim 8	Disclosure
	signal wherein nonlinear responses of the transmitter and/or the optical fiber communications link are pre-compensated, in which the data is specific for the optical transmitter and/or the optical fiber communications link. An electronic communication channel can connected to the data output of the electrical component to communicate the linearization enabling data external to the linearized optical transmitter unit suitable for communication to the direct digital synthesis engine. The linearized optical transmitter can be assembled in a plug-in module. Corresponding HFC systems and methods are also described. Piehler, Abstract. FIG. 3A shows edge QAM 202 and optical transmitter 204 connected by interface 206, e.g., a 25 ohm RF interface. While the division between edge-QAM 202 and optical transmitter 204 can be considered arbitrary, the division provides a convenient design framework since the edge-QAM 202 provides electrical RF signals based on user specific content and optical transmitter 204 provides laser module specific performance. The interface involves transfer of appropriate signals to provide the respective functionalities including, for example, pre-compensation for optical transmitter and fiber nonlinearities for efficient production of optical signals with low distortion. As described below, the division of the optical transmitter provides for the formation of separate modules, which can be plug-in components. A digital switch or other device passes a signal to edge QAM 202 to provide user specific digital content. Piehler, 9:38–53. In additional or alternative embodiments, electrical component 224 may comprise a microprocessor and further receive signal information that is used in creation of the pre-distortion parameters for dynamic correction of nonlinearities. As shown in FIG. 3A, electrical component 224 can receive optionally signals from sensor 228 and/or from an optical coupler or tap 230 and RF coupler or tap 232. Sensor 228 can measure, for example, laser temperature, laser bias voltage and/or other local conditions that can influence optical transmitter nonlinearities. Optical coupler 230 and RF coupler may be integrated into the system at a position downstream and upstream, respectively, of laser driver/laser assembly 222. One or both of the couplers/tap may be used to provide signal information, as indicated by dashed lines

796

Claim 8	Disclosure
	connecting the couplers and microprocessor. If both couplers 230, 232 are used, the microprocessor may compare the upstream and downstream signal information to isolate distortion created by laser driver/laser assembly 222 and RF amplifier 220, and use the comparison data to create data related to the pre-distortion correction that is provided to processing unit 208. The correction may include predetermined parameters specific to laser modulator 222 or be based only on the comparison data. Electrical component 224 may be an integrated circuit or other microprocessor programmed for active processing and/or supply of data or a passive device, e.g., random access memory, non-volatile memory, or other memory. Information in electrical component 224 may comprise parameters for the specific laser module in the device. For example, the laser may be tested at the point of manufacturing or at other times to measure the distortion that it creates. The distortion has consistent components that are specific to the laser; these components may be analyzed and appropriate compensatory corrections created that are placed in, or otherwise introduced by, the electrical component. The pre-distortion correction provided to the DDS may comprise initial values for pre-distortion based on the laser's intrinsic nonlinearity at normal operating conditions as well as other nonlinearities which may exist in the system. The pre-distortion corrections may thus compensate for the net nonlinearity of the optical transmitter which includes, the laser modulator, other optical components, and other RF components including RF amplifiers. Examples of corrections are compensations for one or more of second, third, and higher nth order distortions in the optical transmitter. Piehler, 10:35–11:12. **[0038]** With reference to the drawings, FIG. **2** is a flow chart of the operation of an example embodiment, wherein the LA/CDR unit **144** in conjunction with the microcontroller

797

Claim 8	Disclosure
	150 performs pre-equalization and optimizes power consumption of the transmission over host channel **103** via physical conductors **114** based on the feedback received from the SERDES module **110** via the low-speed digital interface **152**. In FIG. **2**, the operational flow of the LA/CDR unit **144** and the microcontroller **150** begins at step **202** when the microcontroller receives feedback via the low speed digital interface **152**. At step **204**, the microcontroller **150** instructs the LA/CDR unit **144** to equalize its transmitted data signal in response to the feedback, applying sufficient transmit pre-emphasis and other pre-equalization effects to the transmitted data signal to achieve predetermined target eye characteristics at the SERDES module **110**. Once the feedback from the SERDES module **110** indicates that the predetermined target eye characteristics have been achieved, the microcontroller **150** at step **206** instructs the LA/CDR unit **144** to begin reducing the launch amplitude of the transmitted data signal as much as it can without letting the received signal at the SERDES module **110** fall short of the predetermined target eye characteristics. The process is iterative, with feedback being monitored at step **202** and equalization being adjusted at step **204** while the launch amplitude is being reduced at step **206**. In alternative embodiments, the process may operate according to a different sequence of steps while achieving the same result, namely the optimization of signal quality and power consumption over the host channel **103**. Tremblay, paragraph 38

798

I. DEPENDENT CLAIM 9

Claim 9	Disclosure
9[a] The method of claim 8, further comprising: providing a first controller device that configures the first DRR device in response to a power-on event,	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.

799

Claim 9	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.

800

Claim 9	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.

801

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.

802

Claim 9	Disclosure
803	**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 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.

804

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.

805

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.

808

Claim 9	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.

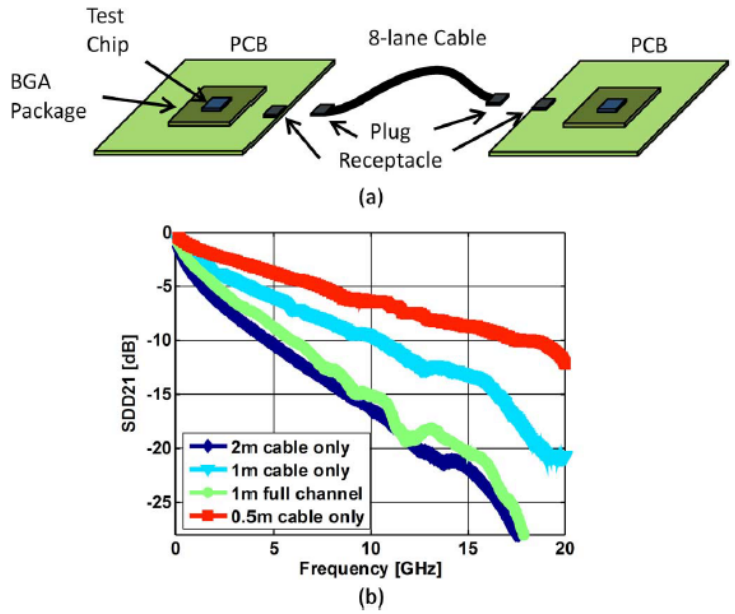
807

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

808

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.

809

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

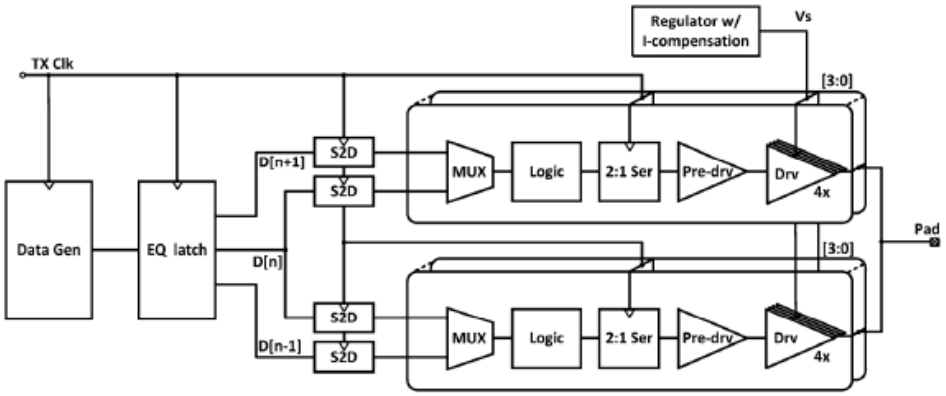
810

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.

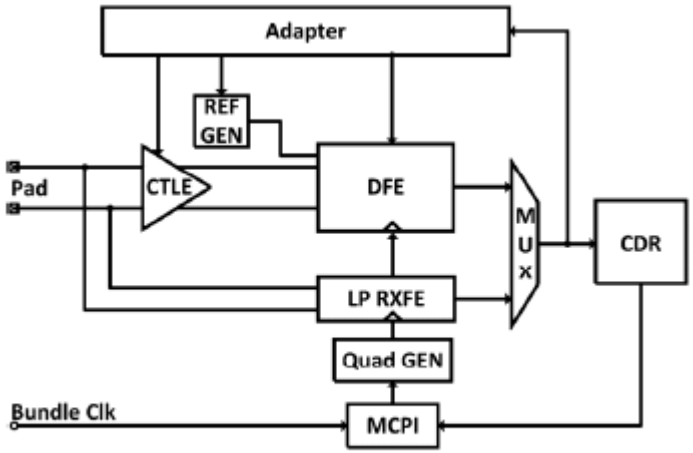
811

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.

812

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

813

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.

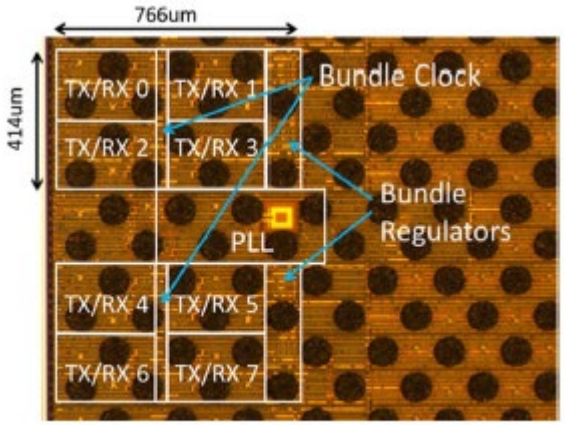
814

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.

815

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.

816

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.

817

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. 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

818

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.

819

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.

820

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.

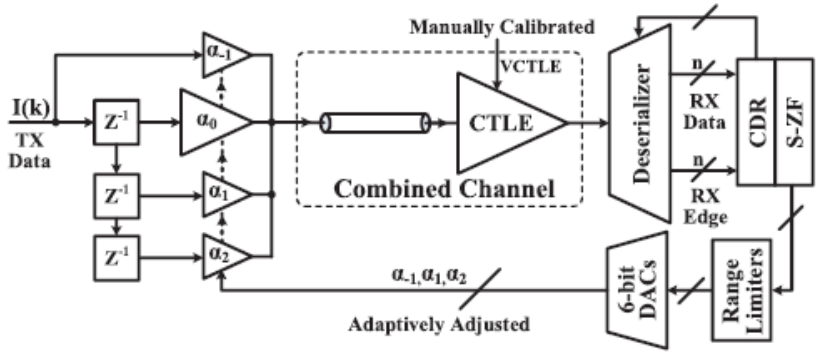


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

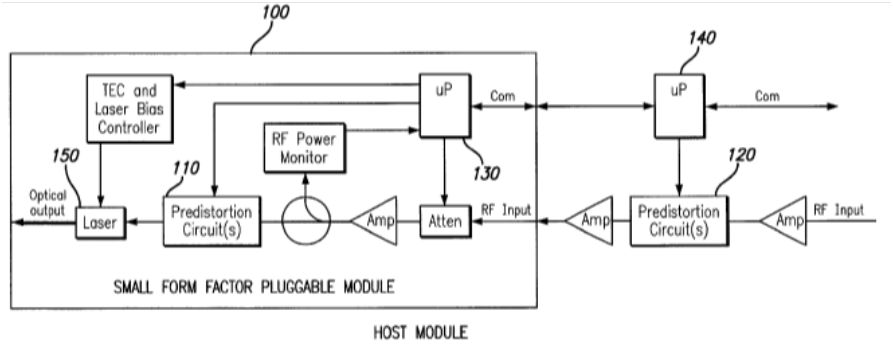
821

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.

822

Claim 9	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. 7,233,617 to Gorecki* Gorecki teaches "[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM . In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up," which is a power-on event. *See* Gorecki at 7:18-33, 16:47-53. In addition, Gorecki describes a "controller" that "distribute[s] to the transmitter" the filter coefficients thereby configuring a filter for pre-emphasis at this power-on event. *See* Gorecki at 7:4-17; 7:26-33.

823

Claim 9	Disclosure
824	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for

Claim 9	Disclosure
	each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4. **[0033]** In example embodiments a low speed digital connection or interface **152** is provided between host SERDES **110** and the interface module **100** to allow back channel communications between transmitting elements of the SERDES module **110** and corresponding receiving elements of the interface module **100** and between transmitting elements of the interface module **100** and receiving components of the SERDES Module **110**. In some implementations, the interface module **100** includes a uC (microcontroller) module **150** that facilitates communications over the low speed digital interface **152**. By way of non-limiting example, the low speed digital interface **152** could include an I2C ("Inter-Integrated Circuit") interface, or SPI ("Serial Peripheral Interface"), MDIO ("Management Data Input/Output") interface, or similar interface. In some example embodiments, the low speed data connection **152** is used for data rates under 10 MB/s. Tremblay, paragraph 33.
9[b] the first controller device retrieving the transmit filter coefficient	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:

825

Claim 9	Disclosure
values from an internal nonvolatile memory as part of said configuring.	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.

826

Claim 9	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.

827

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.

828

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.

829

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.

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.

831

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.

832

Claim 9	Disclosure
833	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.

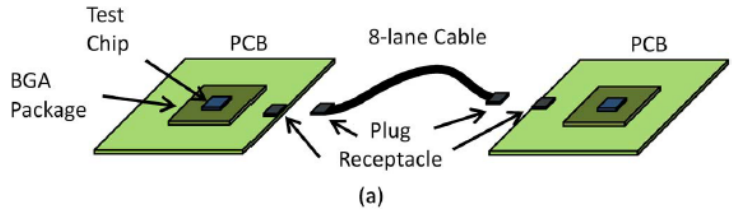
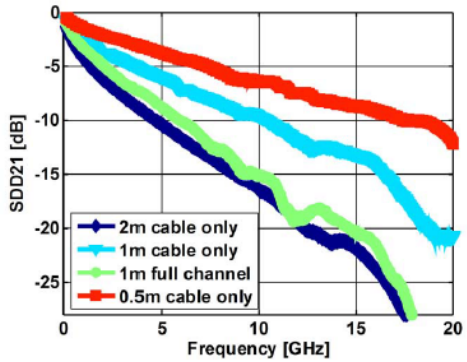
Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze 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 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

834

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.

835

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

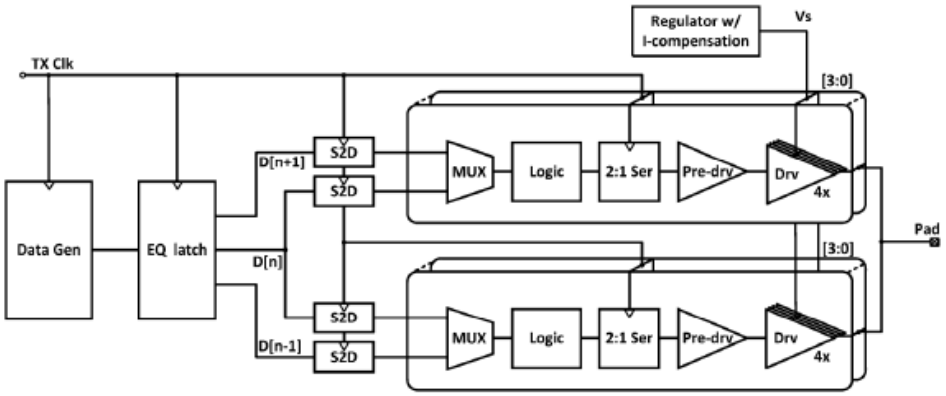
836

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.

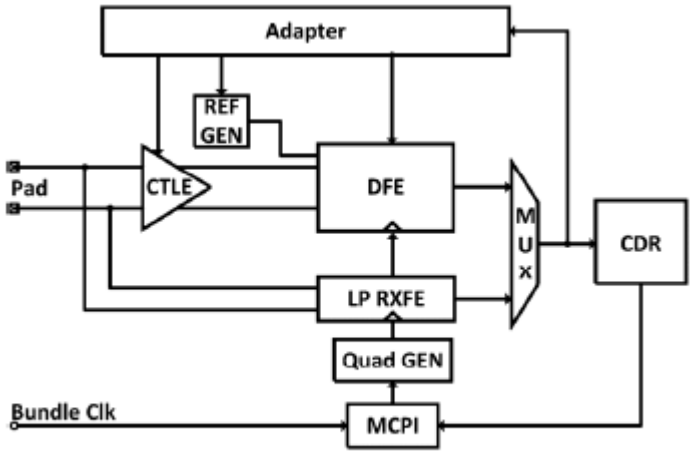
837

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.

838

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

839

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.

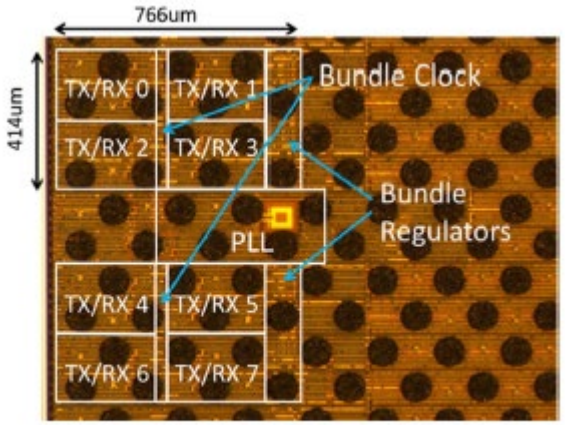
840

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.

841

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.

842

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.

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. 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

844

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.

845

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.

846

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.

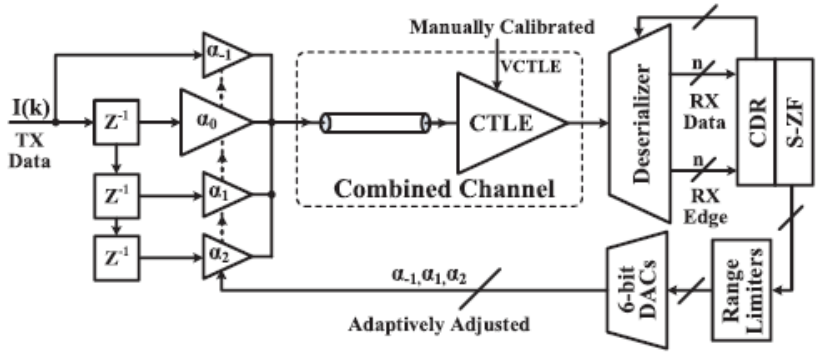


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

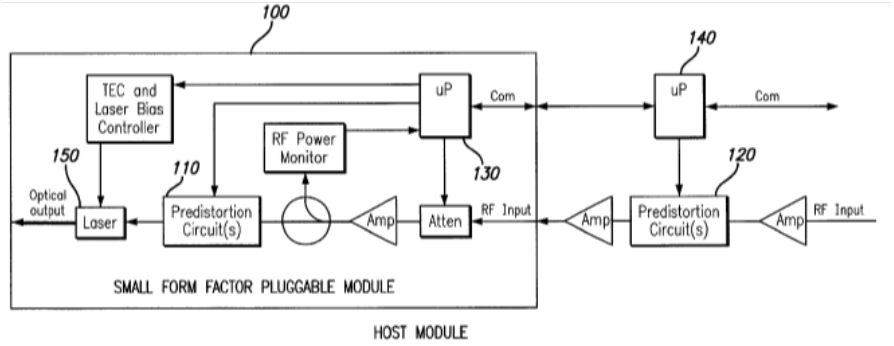
847

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.

848

Claim 9	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. 7,233,617 to Gorecki* Gorecki discloses a "controller" that "distribute[s] to the transmitter" the filter coefficients thereby configuring a filter for pre-emphasis at this power-on event. *See* Gorecki at 7:4-17; 7:26-33.

849

Claim 9	Disclosure
850	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for

Claim 9	Disclosure
	each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4. **[0033]** In example embodiments a low speed digital connection or interface **152** is provided between host SERDES **110** and the interface module **100** to allow back channel communications between transmitting elements of the SERDES module **110** and corresponding receiving elements of the interface module **100** and between transmitting elements of the interface module **100** and receiving components of the SERDES Module **110**. In some implementations, the interface module **100** includes a uC (microcontroller) module **150** that facilitates communications over the low speed digital interface **152**. By way of non-limiting example, the low speed digital interface **152** could include an I2C ("Inter-Integrated Circuit") interface, or SPI ("Serial Peripheral Interface"), MDIO ("Management Data Input/Output") interface, or similar interface. In some example embodiments, the low speed data connection **152** is used for data rates under 10 MB/s. Tremblay, paragraph 33.

J. DEPENDENT CLAIM 10

Claim 10	Disclosure
10. The method of claim 8, further comprising: programming the first and second DRR devices to use the transmit coefficient values each time power is supplied to the first and second end connector plugs.	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.

852

Claim 10	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.

853

Claim 10	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.

854

Claim 10	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.

855

Claim 10	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.

856

Claim 10	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.

857

Claim 10	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.

858

Claim 10	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.

859

Claim 10	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.

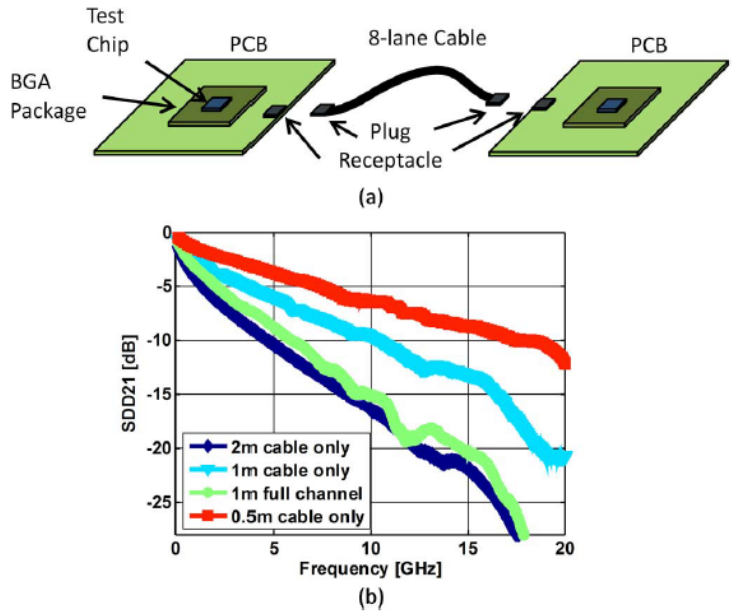
860

Claim 10	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

861

Claim 10	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.

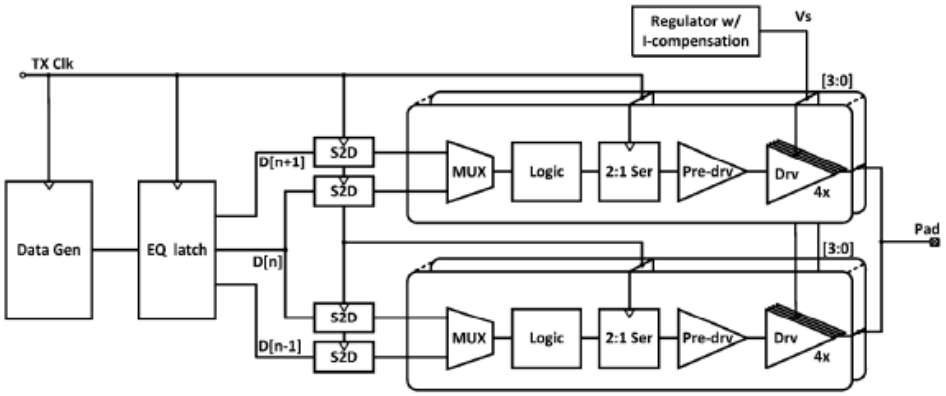
862

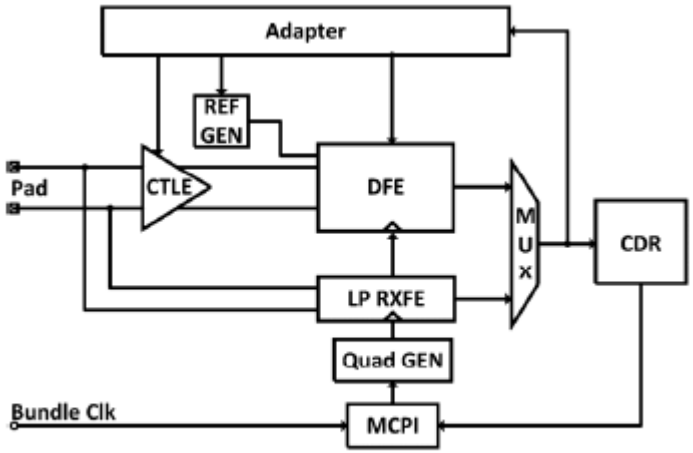
Claim 10	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 10	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 10	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.

865

Claim 10	Disclosure
868	 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 10	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.

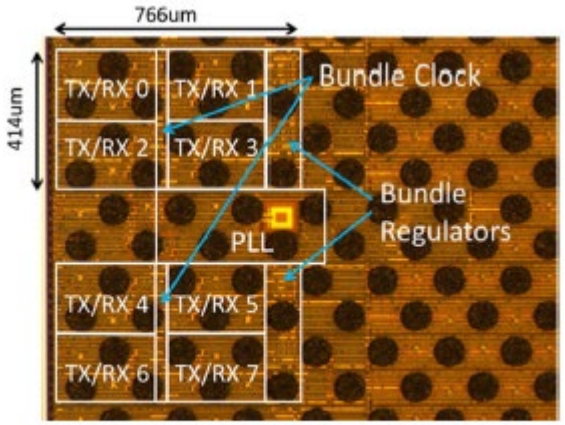
867

Claim 10	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.

868

Claim 10	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.

869

Claim 10	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.

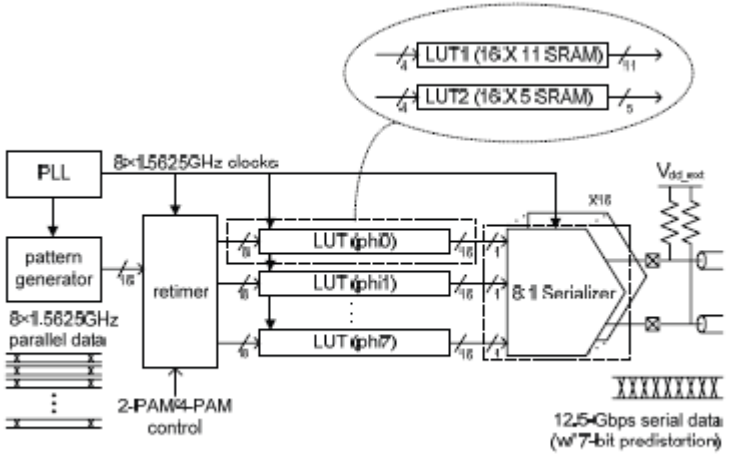
870

Claim 10	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.

871

Claim 10	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.

872

Claim 10	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

873

Claim 10	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

874

Claim 10	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.

875

Claim 10	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

876

Claim 10	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.

877

Claim 10	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.

878

Claim 10	Disclosure

879

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.

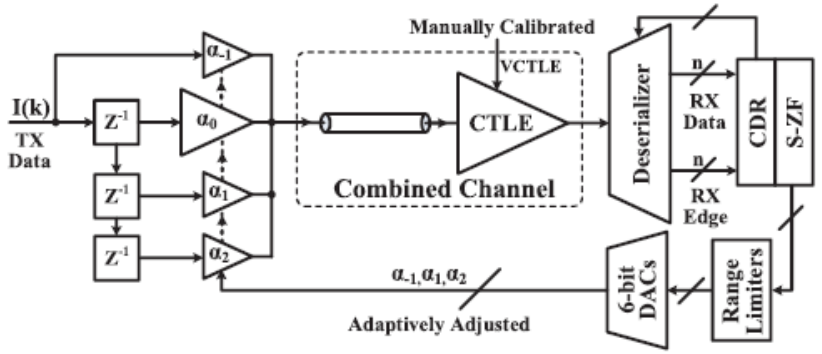


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

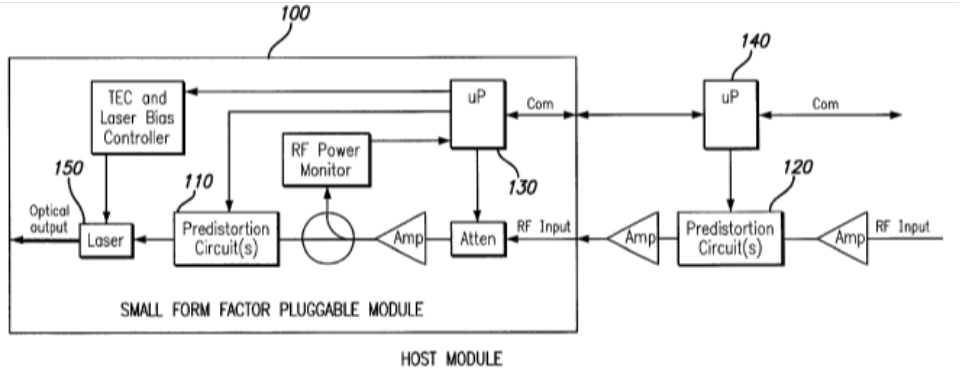
Claim 10	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

880

Claim 10	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. 7,233,617 to Gorecki* For the limited purpose of this analysis, the phrase "transmit coefficient values" is interpreted as referring to "transmit filter coefficient values," (*see, e.g.*, claim limitation 1[f]) as the term otherwise has no antecedent basis. Like the '233 Patent's specification, Gorecki discloses storing of coefficient values in memory as programming. *See* Gorecki at 7:1-33 ("the adjustment or control of the taps of the equalization circuitry may be by the user via programming… coefficients of the tap(s), and/or the pulse durations of the tap(s) may be pre-programmed"); *see also*'233 Patent at 11:3-6 ("the automated tester equipment 'burns' the flash memory or otherwise programs the initial default values of the filter coefficients[.]").

881

Claim 10	Disclosure
	Gorecki teaches that the programmed memory "may reside on the integrated circuit containing the transmitter, receiver or transceiver." *See* Gorecki at 7:43-49. This programming is done to use the transmit coefficient values when power is supplied to the first and second end connector plugs. Gorecki discloses that stored filter coefficients are retrieved from the non-volatile memory at start-up/power-up. *See*, *e.g.*, Gorecki at 7:18-33; *see also* claim 2. Power-up indicates that power is supplied to the first and second end connector plugs. An ordinarily skilled artisan would have understood that power is supplied from a host device when the enclosing connector's connector plug was electrically connected to the host interface port in a corresponding electronic device 101a/b as described claim limitation 1[a] because this was conventionally how standardized connectors like QSFP or SFP worked . *See, e.g.*, SFF-8436 at 21 ("A host board together with the QSFP+ module(s) forms an integrated power system. The host supplies stable power to the module."), 22 Fig. 4 ("Recommended Host Board Power Supply Filtering"). Regardless, an ordinarily skilled artisan would have found it conventional and obvious to implement the QSFP connectors powered by a connected host device. An ordinarily skilled artisan would have understood that Gorecki uses the transmit coefficient values stored in memory each time power is supplied to the first and second end connector plugs, and teaches that the programmed values can be stored "permanently, semi-permanently or temporarily[.]" *See* Gorecki at 7:21. When stored permanently, an ordinarily skilled artisan would have understood them to be used each time power is supplied to the first and second end connector plugs.

Claim 10	Disclosure
	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for

883

Claim 10	Disclosure
	each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 3:51–5:4.

884

K. DEPENDENT CLAIM 11

Claim 11	Disclosure
11[a] The method of claim 10, further comprising: after connecting the electrical conductors, characterizing channel characteristics of the electrical conductors to determine the transmit filter coefficient values; 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.

885

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.

886

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.

887

Claim 11	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.

888

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.

889

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.

890

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.

891

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.

892

Claim 11	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)

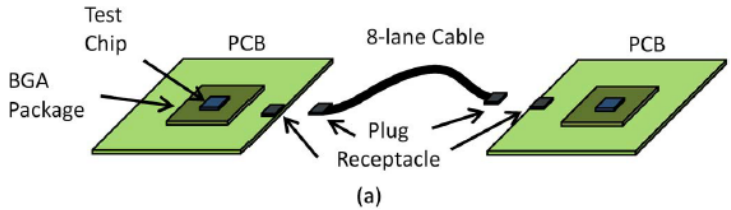
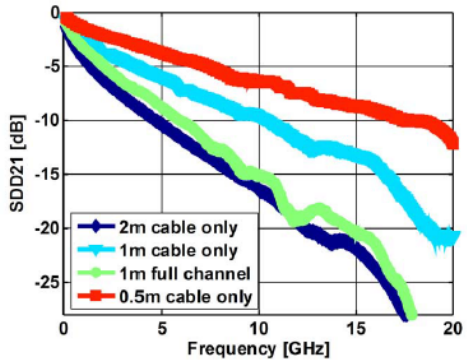
Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze 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 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

894

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.

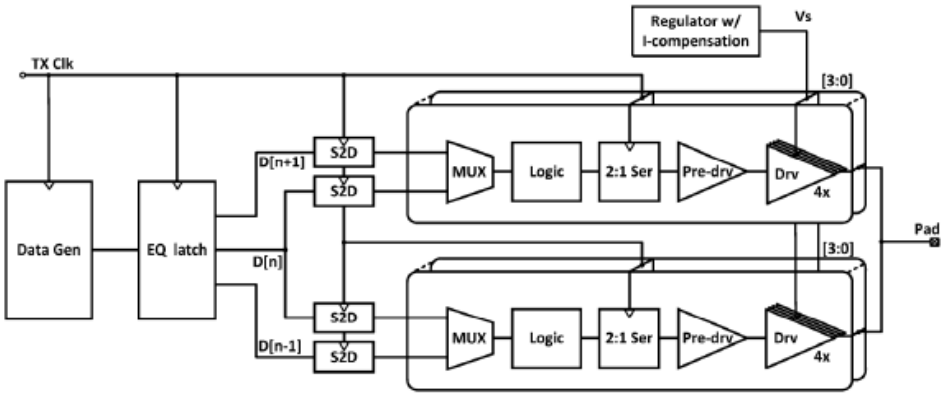
Claim 11	Disclosure
968	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

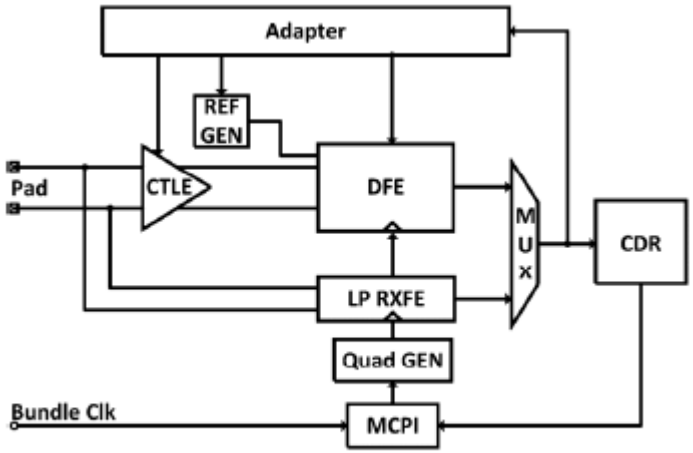
Claim 11	Disclosure
	The CDR block diagram in Fig. 5 shows the bundle timing recovery loop. Similar to the approach introduced in [7], the collaborative CDR combines phase error information from the 4 lanes in a bundle to synthesize the recovered clock. However, instead of using full phase information 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.

897

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.

898

Claim 11	Disclosure
668	 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.

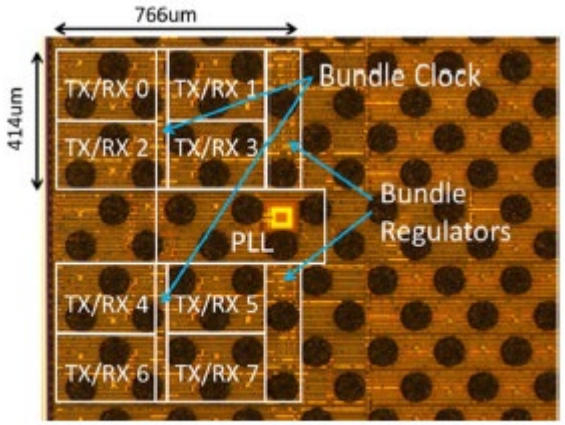
900

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.

901

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.

902

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.

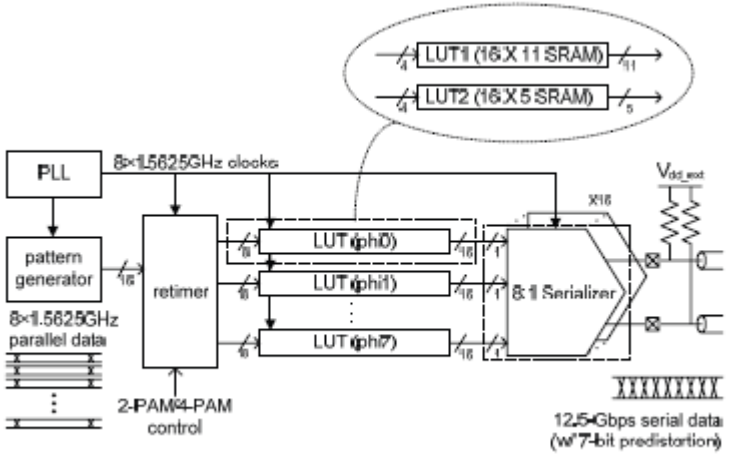
903

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.

904

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.

905

Claim 11	Disclosure
906	 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

907

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.

908

Claim 11	Disclosure
606	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 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.

910

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.

911

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.

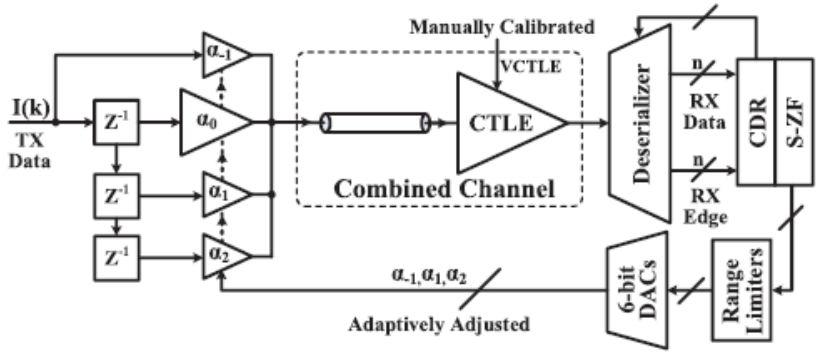


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

912

Claim 11	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.

913

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. *U.S. Patent No. 7,233,617 to Gorecki* An ordinarily skilled artisan would have had reasons to determine transmit filter coefficient values by training as part of the cable manufacture after connecting transceivers to provide the initial pre-set coefficient values. *See, e.g.*, Gorecki at 7:18-33 (describing storing pre-set coefficients). Gorecki discloses that "tap weights or coefficients (values) may be determined or controlled[.]" *See* Gorecki at 6:56-57. Gorecki discloses the storing of filter coefficient values in EEPROM, and use those values to set filter coefficients to provide pre-equalization. *See* Gorecki at 7:18-33 ("[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM. In this way, the transmitter may access the memory to retrieve the necessary information during start-up/power-up,

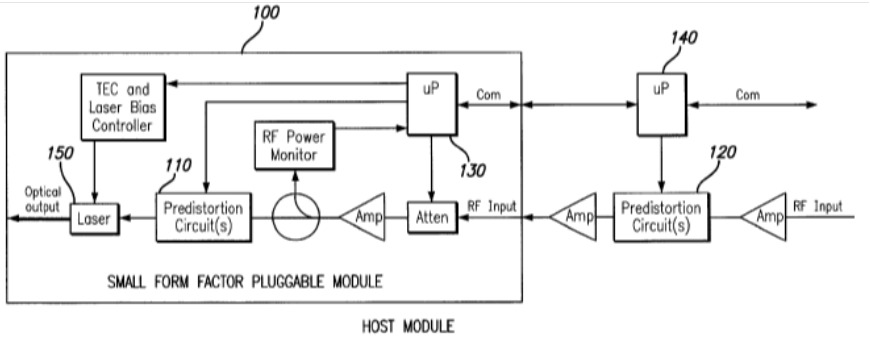
914

Claim 11	Disclosure
	initialization or re-initialization."). An EEPROM [electrically erasable programmable read-only memory] is non-volatile memory. *See, e.g.*, U.S. Patent App. Pub. No. 2015/0106536 at [0081] ("nonvolatile memory can include… electrically erasable PROM (EEPROM)"); U.S. Patent No. 8,000,176 at 2:53-54 ("EEPROM is a non-volatile memory"); U.S. Patent App. Pub. No. 2005/0013317 at [0012] ("non-volatile memory may include… an EEPROM"). The L-G combination teaches programming the coefficients of the FIR filter 17 using the values stored in non-volatile EEPROM to pre-equalize the TX signal in transceiver 107a and 107b. Gorecki discloses determining and then storing the transmit filter coefficient values in the non-volatile memories after assembly of the cable. The values should be selected to offset effects of a communication channel. *See* Gorecki at 9:8-10 (The equalizers are "configured to introduce compensation that is ideally the inverse of the effects caused by the communications channel."). An ordinarily skilled artisan would have found it obvious to determine and then store filter coefficients, such as by measuring the effects in that communication channel after assembly. Such an approach was known in the art . *See, e.g.*, U.S. Patent App. Pub. No. 2014/0119425 at [0033] ("The programming is preferably performed with the cable installed between the actual source and actual sink it will be connecting."). Gorecki describes "the coefficients of the tap(s)… are adjusted or controlled during operation of the transmitter, for example using an adaptive algorithm." *See* Gorecki at 10:27-51, 6:50-53, 17:41-51. Determining values "during operation" happens after cable assembly. The adaptive algorithm—which determines transmit filter coefficient values—is performed "after (or during) the performance of an initialization or re-initialization process." *See* Gorecki at 7:50-61. Gorecki discloses that stored coefficient values can be "fine-tuned to enhance the system performance" (Gorecki at 7:51-53), and that the values can be "re-programmed" and, for that reason, are stored in writable non-volatile memory. *See, e.g.*, Gorecki at 7:18-33, 10:37-51. An ordinarily skilled artisan would have had reasons to use NVM including writable NVM ("EEPROM or the like"). *See* Gorecki at 7:24-25, 10:46-47. EEPROM is electrically erasable programmable read only memory, e.g., rewritable NVM. An ordinarily skilled artisan would have known that flash memory was another well-known rewritable non-volatile memory (e.g., "EEPROM or the like") conventionally used with transmit FIR filters. *See, e.g.*, U.S. Patent No. 6,975,140 at 3:42-4:2 ("Any type of NVRAM can be used as the rewritable non-volatile storage, so long as [it] is of the type suitable for integration on an

915

Claim 11	Disclosure
	integrated circuit."). An ordinarily skilled artisan would have understood such memories would be used when the content of those memories is both re-programmed and preserved through "startup/power-up, initialization or re-initialization," as described in Gorecki at 7:18-33. Goreki discloses rewritable non-volatile memory that could be updated with new coefficient values that fine-tuned pre-equalization or equalization with new coefficient values learned through training at manufacture or during usage with a particular device in a particular environment . *See, e.g.*, Gorecki '716 at 7:50-61, 10:37-51; *see also* U.S. Patent No. 6,975,140 at 3:42-4:2, 9:21-28. Preserving such adapted pre-equalization filter coefficients in rewritable non-volatile memory like an EEPROM would have ensured that the fine-tuned coefficients remained available after an AEC transceiver—or the host powering it—was power cycled, initialized, or re-initialized. *See* Gorecki at 7:31-33 ("In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up, initialization or re-initialization."). Regardless, an ordinarily skilled artisan would have known that storing updated filter coefficients in the L-G combination's non-volatile memory was conventional and desirable so that the adapted coefficient values would be available after power cycling the active electric cable transceiver by unplugging it from a host device powering it, or power cycling the host device powering the active element cable transceiver. *See* U.S. Patent No. 6,975,140 at Abstract ("[D]ata transmitter includes a rewriteable non-volatile storage, operable to be rewritten with control information representing the values of coefficients updated during operation off the FIR driver."), 2:46-49 (same), 3:45-52 ("Updated values of the coefficients used in the taps of the FIR transmitter 30 are determined and stored to the flash memory 20 during the operation of the FIR transmitter . When the FIR transmitter is powered off and back on again, the coefficient values as last updated are available to be retrieved from the flash memory 20 and applied to the taps of the FIR transmitter again.").

916

Claim 11	Disclosure
	 FIG. 1 Ishaug, Figure 1. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or

917

Claim 11	Disclosure
	parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:63–5:4.
11[b] storing the transmit filter coefficient values in the nonvolatile memories.	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 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.

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.

920

Claim 11	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.

921

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.

922

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.

923

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.

924

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.

925

Claim 11	Disclosure

926

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	

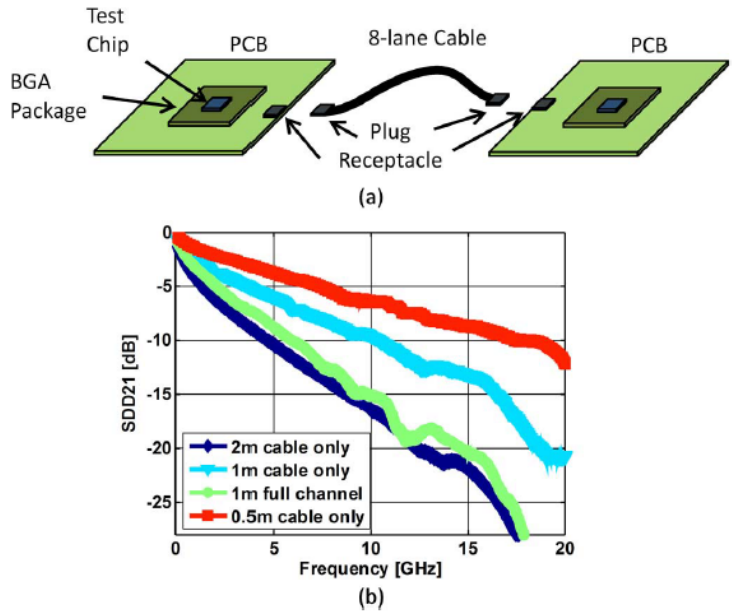
CMIS, 71.

Claim 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

927

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.

928

Claim 11	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

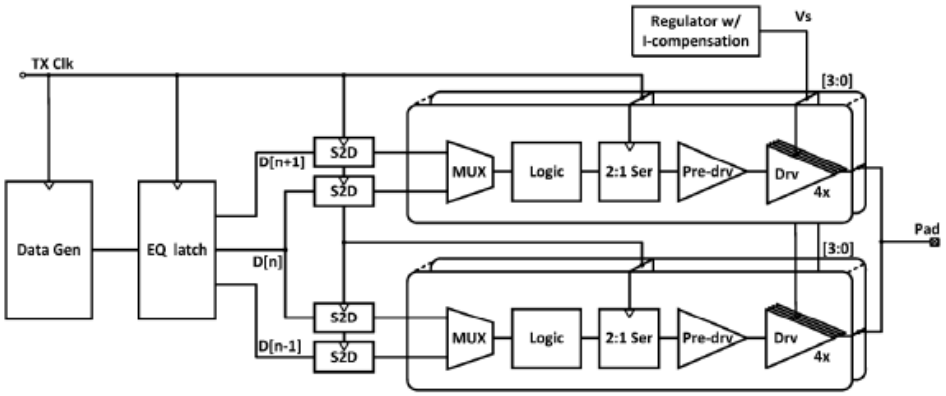
929

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.

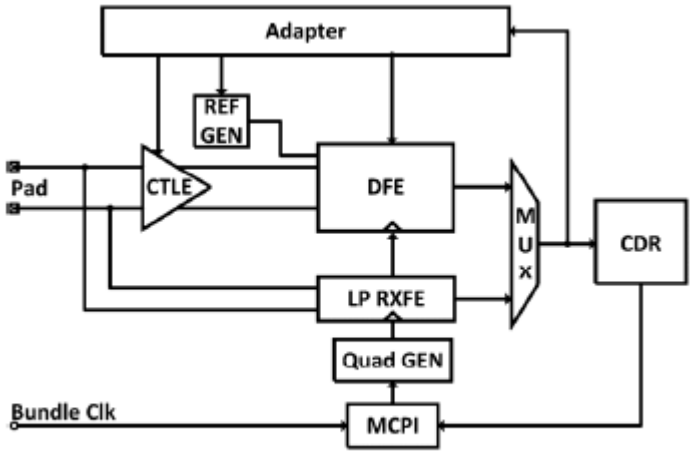
930

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.

931

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

932

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.

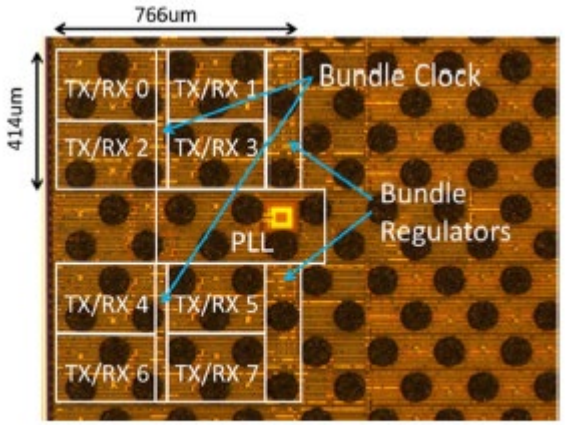
933

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.

934

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.

935

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.

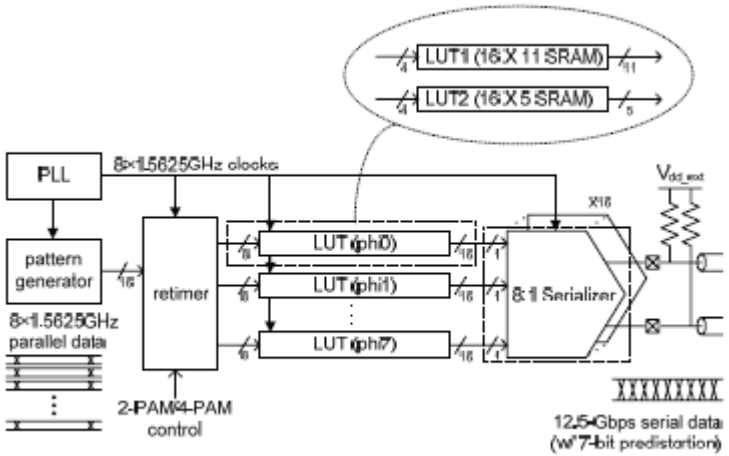
936

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.

937

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.

938

Claim 11	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

939

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

940

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.

941

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

942

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.

943

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.

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.

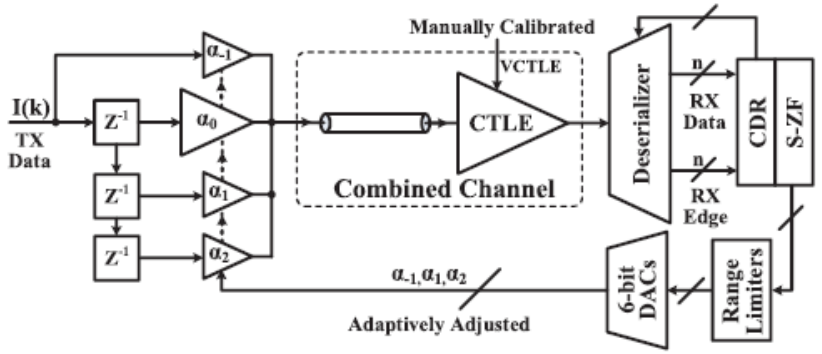


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

945

Claim 11	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.

946

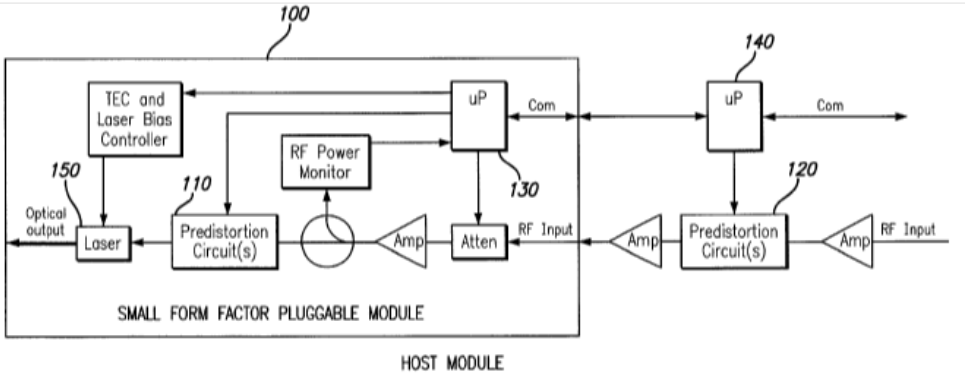
Claim 11	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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. *U.S. Patent No. 7,233,617 to Gorecki* An ordinarily skilled artisan would have had reasons to determine transmit filter coefficient values by training as part of the cable manufacture after connecting transceivers to provide the initial pre-set coefficient values. *See, e.g.*, Gorecki at 7:18-33 (describing storing pre-set coefficients). Gorecki discloses that "tap weights or coefficients (values) may be determined or controlled[.]" *See* Gorecki at 6:56-57. Gorecki discloses the storing of filter coefficient values in EEPROM, and use those values to set filter coefficients to provide pre-equalization. *See* Gorecki at 7:18-33 ("[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM. In this way, the transmitter may access the memory to retrieve the necessary information during start-up/power-up,

947

Claim 11	Disclosure
	initialization or re-initialization."). An EEPROM [electrically erasable programmable read-only memory] is non-volatile memory. *See, e.g.*, U.S. Patent App. Pub. No. 2015/0106536 at [0081] ("nonvolatile memory can include… electrically erasable PROM (EEPROM)"); U.S. Patent No. 8,000,176 at 2:53-54 ("EEPROM is a non-volatile memory"); U.S. Patent App. Pub. No. 2005/0013317 at [0012] ("non-volatile memory may include… an EEPROM"). The L-G combination teaches programming the coefficients of the FIR filter 17 using the values stored in non-volatile EEPROM to pre-equalize the TX signal in transceiver 107a and 107b. Gorecki discloses determining and then storing the transmit filter coefficient values in the non-volatile memories after assembly of the cable. The values should be selected to offset effects of a communication channel. *See* Gorecki at 9:8-10 (The equalizers are "configured to introduce compensation that is ideally the inverse of the effects caused by the communications channel."). An ordinarily skilled artisan would have found it obvious to determine and then store filter coefficients, such as by measuring the effects in that communication channel after assembly. Such an approach was known in the art . *See, e.g.*, U.S. Patent App. Pub. No. 2014/0119425 at [0033] ("The programming is preferably performed with the cable installed between the actual source and actual sink it will be connecting."). Gorecki describes "the coefficients of the tap(s)… are adjusted or controlled during operation of the transmitter, for example using an adaptive algorithm." *See* Gorecki at 10:27-51, 6:50-53, 17:41-51. Determining values "during operation" happens after cable assembly. The adaptive algorithm—which determines transmit filter coefficient values—is performed "after (or during) the performance of an initialization or re-initialization process." *See* Gorecki at 7:50-61. Gorecki discloses that stored coefficient values can be "fine-tuned to enhance the system performance" (Gorecki at 7:51-53), and that the values can be "re-programmed" and, for that reason, are stored in writable non-volatile memory. *See, e.g.*, Gorecki at 7:18-33, 10:37-51. An ordinarily skilled artisan would have had reasons to use NVM including writable NVM ("EEPROM or the like"). *See* Gorecki at 7:24-25, 10:46-47. EEPROM is electrically erasable programmable read only memory, e.g., rewritable NVM. An ordinarily skilled artisan would have known that flash memory was another well-known rewritable non-volatile memory (e.g., "EEPROM or the like") conventionally used with transmit FIR filters. *See, e.g.*, U.S. Patent No. 6,975,140 at 3:42-4:2 ("Any type of NVRAM can be used as the rewritable non-volatile storage, so long as [it] is of the type suitable for integration on an

Claim 11	Disclosure
	integrated circuit."). An ordinarily skilled artisan would have understood such memories would be used when the content of those memories is both re-programmed and preserved through "startup/power-up, initialization or re-initialization," as described in Gorecki at 7:18-33. Goreki discloses rewritable non-volatile memory that could be updated with new coefficient values that fine-tuned pre-equalization or equalization with new coefficient values learned through training at manufacture or during usage with a particular device in a particular environment . *See, e.g.*, Gorecki '716 at 7:50-61, 10:37-51; *see also* U.S. Patent No. 6,975,140 at 3:42-4:2, 9:21-28. Preserving such adapted pre-equalization filter coefficients in rewritable non-volatile memory like an EEPROM would have ensured that the fine-tuned coefficients remained available after an AEC transceiver—or the host powering it—was power cycled, initialized, or re-initialized. *See* Gorecki at 7:31-33 ("In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up, initialization or re-initialization."). Regardless, an ordinarily skilled artisan would have known that storing updated filter coefficients in the L-G combination's non-volatile memory was conventional and desirable so that the adapted coefficient values would be available after power cycling the active electric cable transceiver by unplugging it from a host device powering it, or power cycling the host device powering the active element cable transceiver. *See* U.S. Patent No. 6,975,140 at Abstract ("[D]ata transmitter includes a rewriteable non-volatile storage, operable to be rewritten with control information representing the values of coefficients updated during operation off the FIR driver."), 2:46-49 (same), 3:45-52 ("Updated values of the coefficients used in the taps of the FIR transmitter 30 are determined and stored to the flash memory 20 during the operation of the FIR transmitter . When the FIR transmitter is powered off and back on again, the coefficient values as last updated are available to be retrieved from the flash memory 20 and applied to the taps of the FIR transmitter again.").

949

Claim 11	Disclosure
	 FIG. 1 Ishaug, Figure 1.

950

Claim 11	Disclosure
	In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. Ishaug, 4:51–54. **[0043]** In some example embodiments, the interface module **100** includes a printed circuit board to which the CDR module **128**, LD module **130**, laser diode module **140**, photo diode module **148**, TIA module **146**. LA/CDR module **144** and uC module **150** are each mounted. At least some the modules mounted to the printed circuit board may be implemented as monolithic integrated circuits on respective silicon chips, including for example CDR module **126** and LA/CDR module **144**. Tremblay, paragraph 43.

951

L. DEPENDENT CLAIM 12

Claim 12	Disclosure
12. The method of claim 11, wherein the first and second DRR devices employ receiver-based equalization of the electrical transit signals using coefficient values stored in the nonvolatile memories.	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.

952

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.

953

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.

954

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.

955

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.

956

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.

957

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.

958

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.

959

Claim 12	Disclosure

960

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	

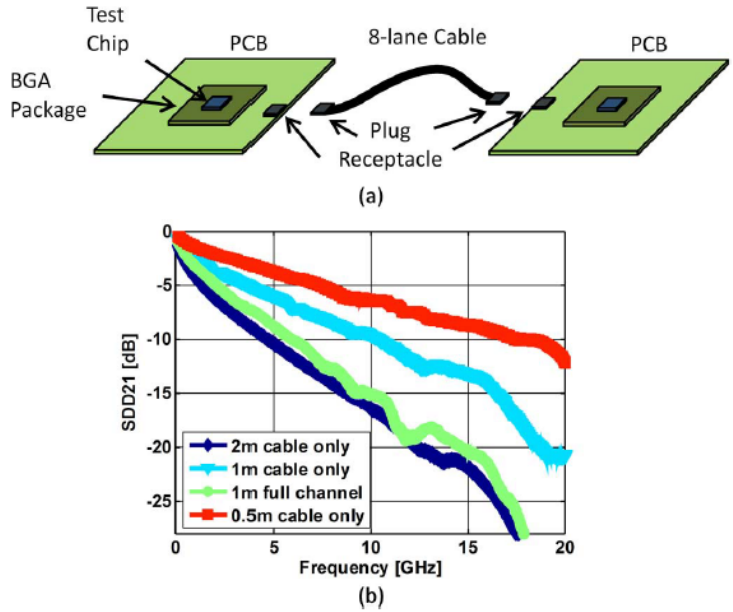
CMIS, 71.

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

961

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.

962

Claim 12	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

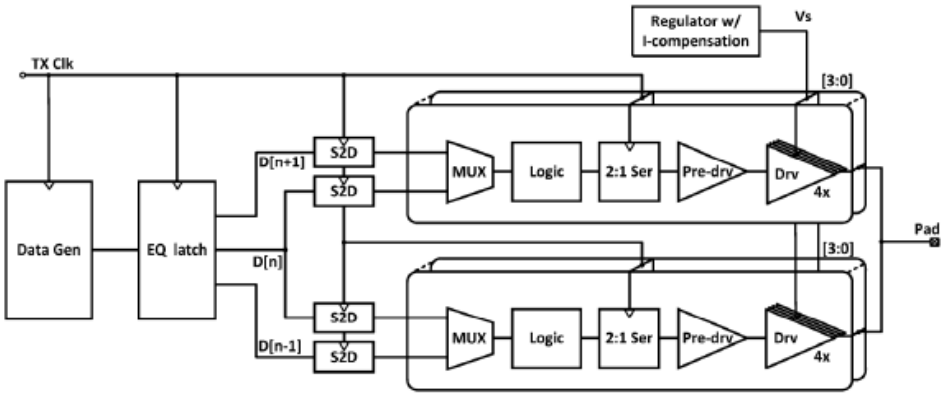
963

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.

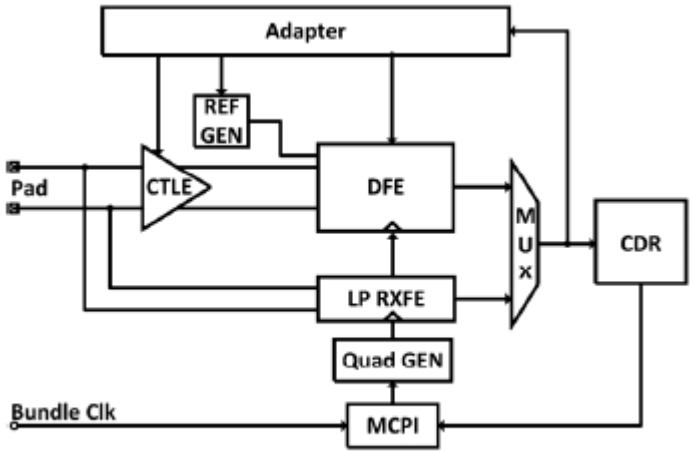
964

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.

965

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

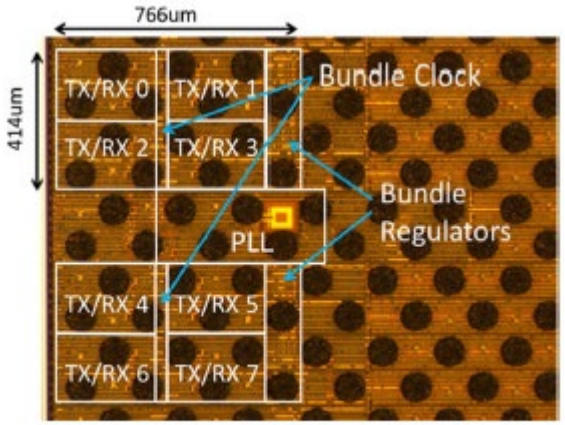
996

Claim 12	Disclosure
967	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 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.

968

Claim 12	Disclosure
696	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.

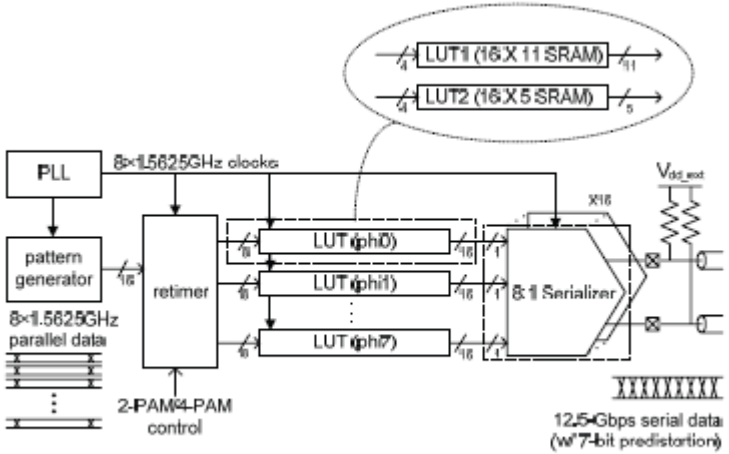
Claim 12	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

970

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.

971

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.

Claim 12	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

973

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

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.

975

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

976

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.

977

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.

978

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.

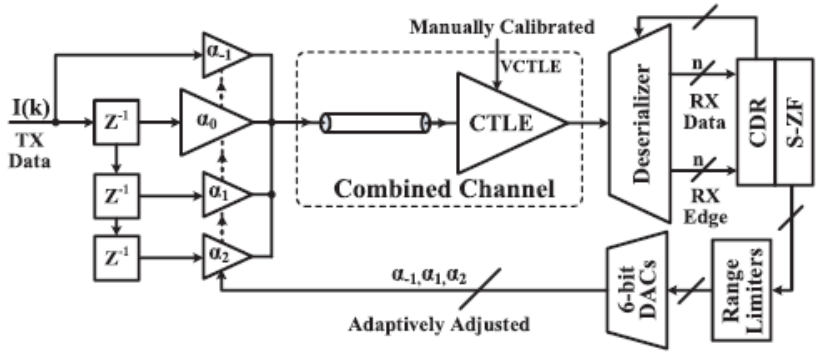


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

979

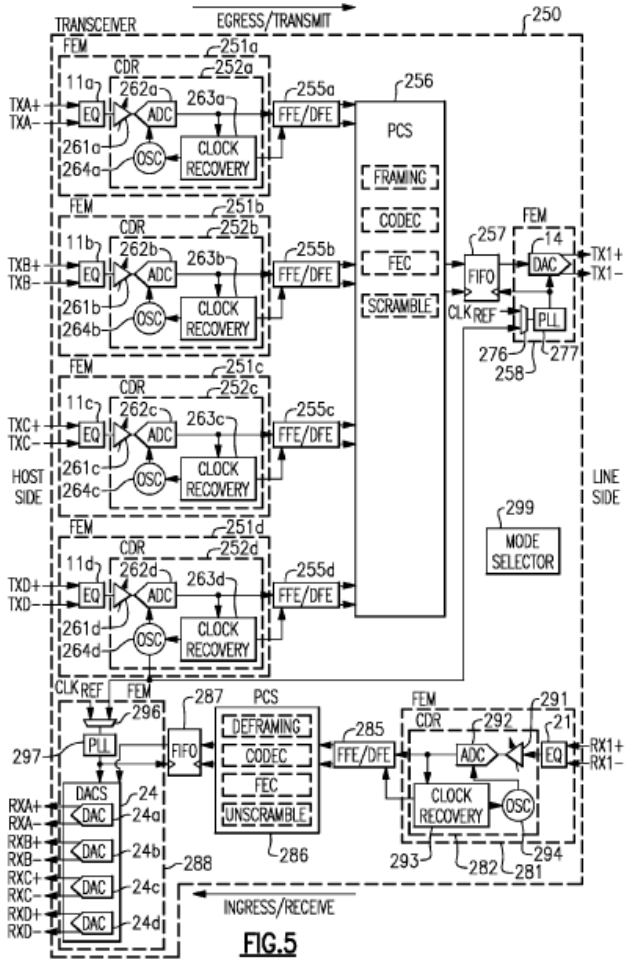
Claim 12	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

980

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. *U.S. Patent No. 7,233,617 to Gorecki* The '233 Patent specification repeats the "receiver-based equalization" language of dependent claim 5, but fails to provide any explanation of what the phrase means. *See, e.g.*, '233 Patent at 2:42-45, 2:61-63. Nevertheless, "receiver-based equalization" is met by equalizing a received signal to compensate for distortion incurred during transmission. *See, e.g.*, U.S. Patent No. 9,178,542 at 10:22-35; U.S. Patent App. Pub. No. 2005/0078758 at [0004]; U.S. Patent No. 7,570,708 at 2:17-21; Pavan Kumar Hanumolu, *et al.*, "Equalizers for High-Speed Serial Links," International Journal of High Speed Electronics and Systems, vol. 15, no. 2, 2005 ("Hanumolu") at p.185 (describing "receive-side equalization" suppressing channel loss).

981

Claim 12	Disclosure
	Gorecki suggest that equalizers are, like filters, configured to compensate for signal distortion over the cable. *See* Gorecki at 1:59-67 ("[C]onventional high speed digital baseband communications systems often employ circuitry, for example, a finite impulse response filter ('FIR filter'), in the receiver.… Such circuitry typically includes [sic] one or more taps having fixed or pre-programmed 'positions' and coefficients."); *see id.* at 2:17-21 ("Regardless of where the equalization circuitry is implemented, the duration of the equalization signal and the relative position or placement of the tap(s) of the equalization circuitry are selected or designed to avoid interference with the signal representative of the transmitted information.").
	The equalization functionality depends on filter coefficient settings. *See, e.g.,* Hanumolu at pp. 185-198 (describing different receive-side equalization architectures configured with coefficients [C-1, C0, C1]); U.S. Patent App. Pub. No. 2014/0281068 at [0069] (describing applying coefficients to a receiver's settings to minimize communication data loss).
	Though Gorecki describes implementing its equalization structure and techniques at a transmitter "may have an advantage," (*see, e.g.*, Gorecki at 18:47-53), an ordinarily skilled artisan would have recognized from this statement that implementation in the receiver was also an option, particularly as the combination explains implementation more generally in connection with equalization circuitry more generally. *See* Gorecki at 2:17-21, 6:14-23 ("The circuitry and techniques include leading and/or trailing taps to reduce, minimize, mitigate or effectively eliminate precursor and/or post-cursor intersymbol interference due to, for example, bandwidth limitations and reflections in high-speed digital communication systems. In this way, the equalization circuitry and techniques may reduce, minimize or eliminate non optimum (e.g., over-equalization) at the boundaries of the data signal (i.e., the symbol)."). An ordinarily skilled artisan would have considered such an approach would have been applicable to equalizers 21 as they provide equalization in the same way as transmit filter 17.
	Additionally, an ordinarily skilled artisan would have understood that the equalization performed at both a transmit filter and a receive path equalizer should collectively compensate for the effects introduced the communication channel and would have had a reason to set equalization parameters for both in the same way.

Claim 12	Disclosure
	 FIG.5 Lugthart, Figure 5. The illustrated transmit path line-side FEM 258 includes a transmit path DAC 14, a multiplexer 276 and a phase-locked loop (PLL) 277. Additionally, the receive path line-side FEM 281 of FIG. 5 includes a

983

Claim 12	Disclosure
	receive path equalizer 21 and a receive path CDR circuit 282, which includes a thresholder 291, an ADC 292, a clock recovery circuit 293, and a controllable oscillator 294. Furthermore, the illustrated receive path host-side FEM 288 includes a multiplexor 296, a PLL 297, and a plurality of receive path DACs 24 including a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. Lugthart, 5:24–34. 10. A transceiver comprising: an equalizer configured to receive a data signal and to generate a clock and data recovery (CDR) input signal based on equalizing the data signal; and a CDR circuit comprising: a data recovery analog-to-digital converter (ADC) configured to receive the CDR input signal and to generate a first digital output signal; a clock recovery ADC configured to receive the CDR input signal and to generate a second digital output signal, wherein the data recovery ADC has a higher resolution than the clock recovery ADC; a timing control circuit configured to receive the second digital output signal and to generate a clock control signal based on the second digital output signal; and a clock generation unit configured to control a timing of the data recovery ADC and a timing of the clock recovery ADC based on the clock control signal. Lugthart, 58:64–59:13.

984

Claim 12	Disclosure
	[0029] Example embodiments will now be described with reference to FIG. **1** which illustrates a host **102** that is connected through a host channel **103** to a pluggable network transceiver interface module **100**, which in turn is connected to a transmission medium **104**. In the illustrated embodiment, the transmission medium **104** includes fibre optic cable **124** for transmitting high speed data from the interface module **100** and fibre optic cable **126** for transmitting high speed data to the interface module **100**, however embodiments described herein can also be applied to non-fibre transmission mediums such as copper for example. In the illustrated embodiment, the host **102** includes a network interface card or board **106** that allows data to be exchanged between the host **102** and the network interface module **100**. In at least one example embodiment, the host interface board **106** includes data processing circuitry, which could for example be implemented an ASIC **108** ("Application Specific Integrated Circuit"), and a SERDES ("SERializer-DESerializer") chip or module **110**. SERDES module **110** receives parallel data from the host ASIC **108** over a parallel data channel **170** and converts it into high speed serial data for transmission over the host channel **103** to the network interface module **100**, and receives high speed serial data from the network interface module **100** through the host channel **103** which it then converts into parallel data and passes on to the ASIC **108** over the parallel data channel **170**. Tremblay, paragraph 29.

985

Claim 12	Disclosure
	[0038] With reference to the drawings, FIG. **2** is a flow chart of the operation of an example embodiment, wherein the LA/CDR unit **144** in conjunction with the microcontroller **150** performs pre-equalization and optimizes power consumption of the transmission over host channel **103** via physical conductors **114** based on the feedback received from the SERDES module **110** via the low-speed digital interface **152**. In FIG. **2**, the operational flow of the LA/CDR unit **144** and the microcontroller **150** begins at step **202** when the microcontroller receives feedback via the low speed digital interface **152**. At step **204**, the microcontroller **150** instructs the LA/CDR unit **144** to equalize its transmitted data signal in response to the feedback, applying sufficient transmit pre-emphasis and other pre-equalization effects to the transmitted data signal to achieve predetermined target eye characteristics at the SERDES module **110**. Once the feedback from the SERDES module **110** indicates that the predetermined target eye characteristics have been achieved, the microcontroller **150** at step **206** instructs the LA/CDR unit **144** to begin reducing the launch amplitude of the transmitted data signal as much as it can without letting the received signal at the SERDES module **110** fall short of the predetermined target eye characteristics. The process is iterative, with feedback being monitored at step **202** and equalization being adjusted at step **204** while the launch amplitude is being reduced at step **206**. In alternative embodiments, the process may operate according to a different sequence of steps while achieving the same result, namely the optimization of signal quality and power consumption over the host channel **103**. Tremblay, paragraph 38.

986

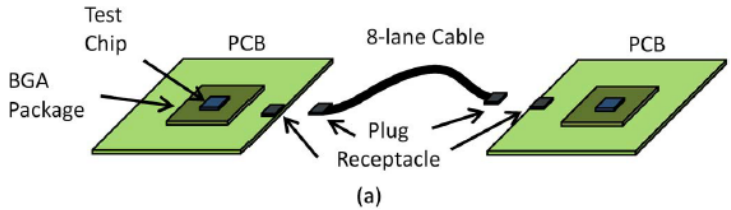
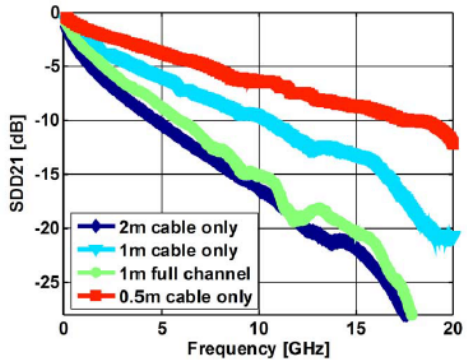
M. DEPENDENT CLAIM 13

Claim 13	Disclosure
13. The method of claim 11, wherein the electrical conductors comprise twin-axial conductors that carry each of the electrical transit signals in differential form.	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

987

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.

988

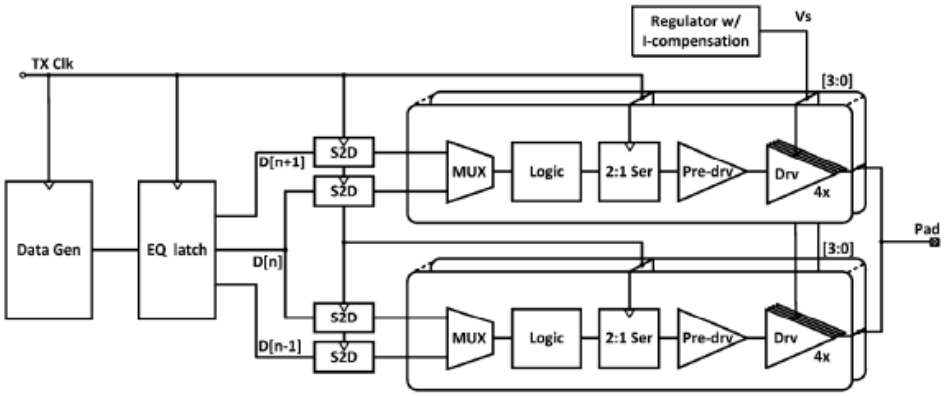
Claim 13	Disclosure
686	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

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.

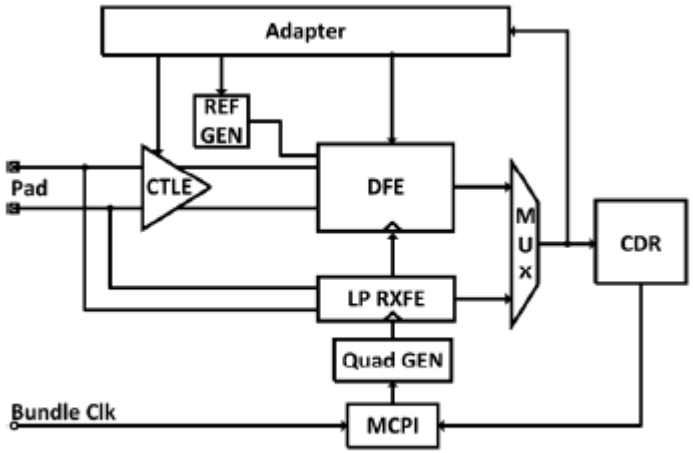
066

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.

991

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

992

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.

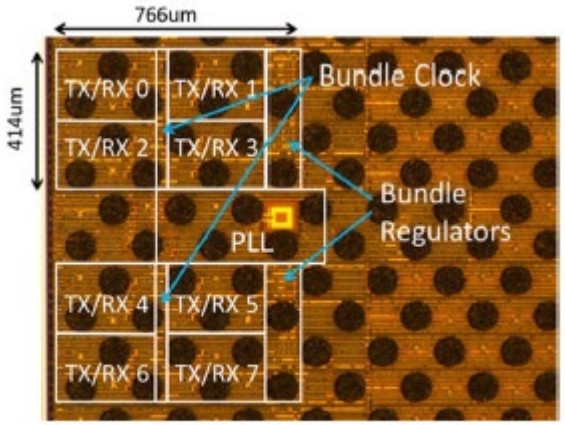
993

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.

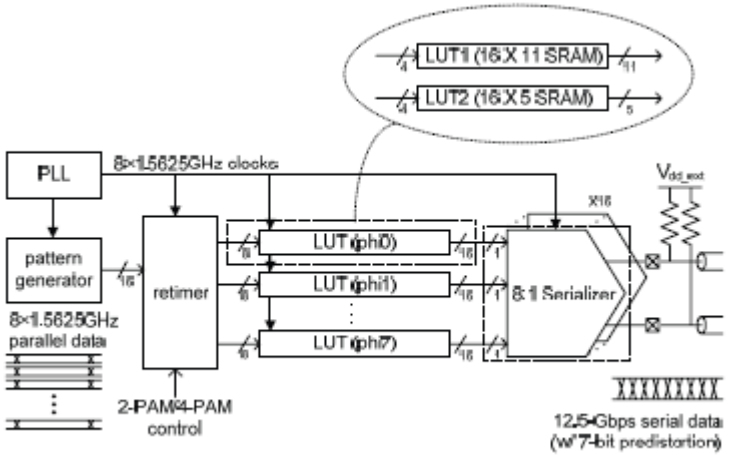
994

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.

995

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.

966

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.  Fig. 1. Overall transmitter block diagram Chung, Figure 1.

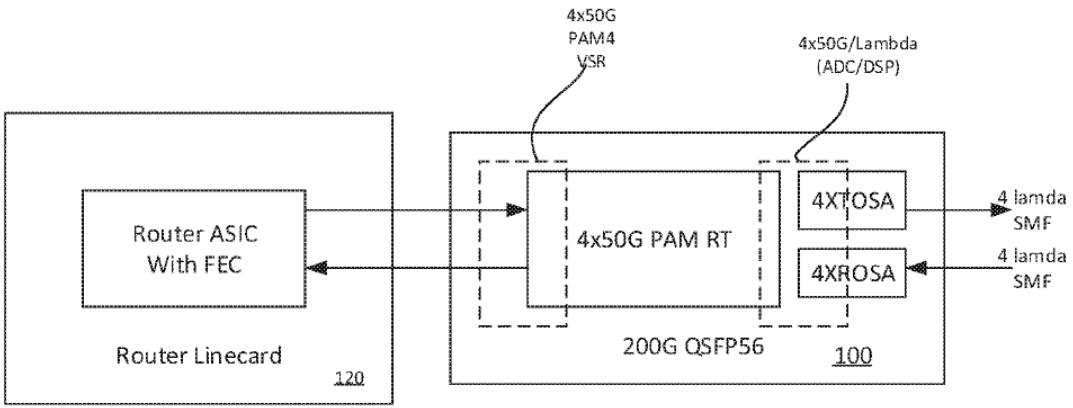
997

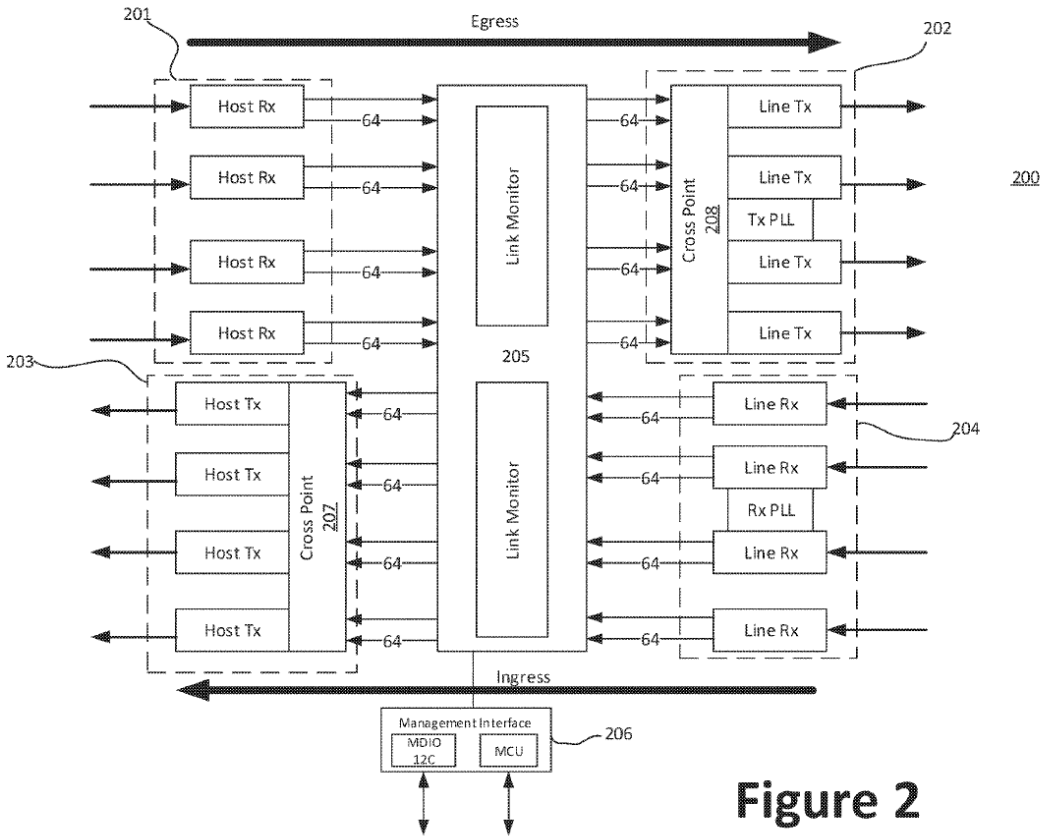
Claim 13	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

998

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.

666

Claim 13	Disclosure
1000	 **Figure 1** Riani, Figure 1.

Claim 13	Disclosure
1001	 Figure 2 Riani, Figure 2.

Claim 13	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.

1002

Claim 13	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 13	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.

1004

Claim 13	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.

1005

Claim 13	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.

1006

Claim 13	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.

1007

Claim 13	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.

1008

Claim 13	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.

1009

Claim 13	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.

1010

	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.	

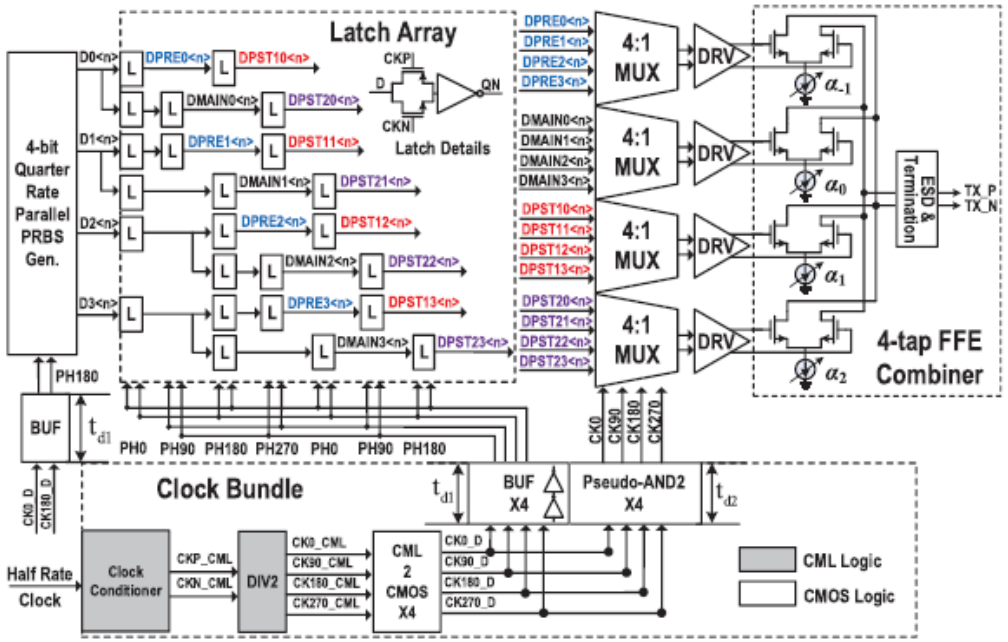
1011

Claim 13	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.

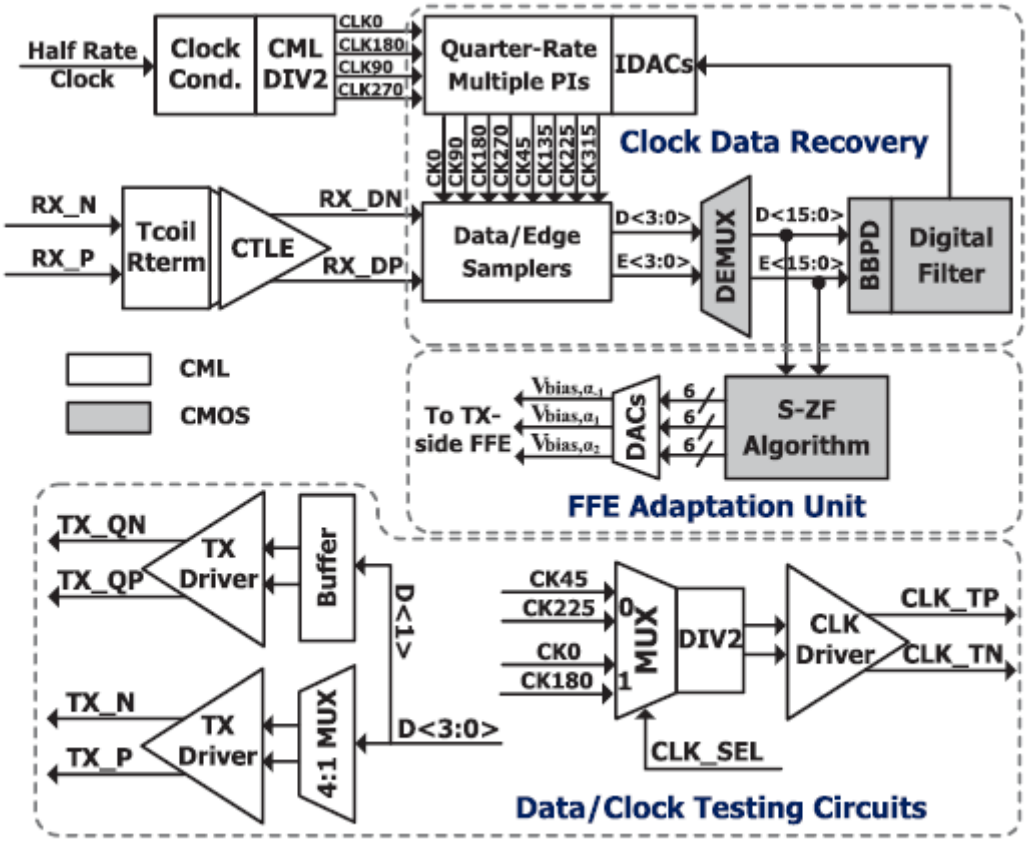
1012

Claim 13	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.

1013

Claim 13	Disclosure
	 Fig. 1. Block diagram of the TX chip. Zheng, 2964.

1014

Claim 13	Disclosure
	 Fig. 8. Block diagram of the RX chip. Zheng, 2966.

1015

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

1016

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.

1017

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.

1018

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.

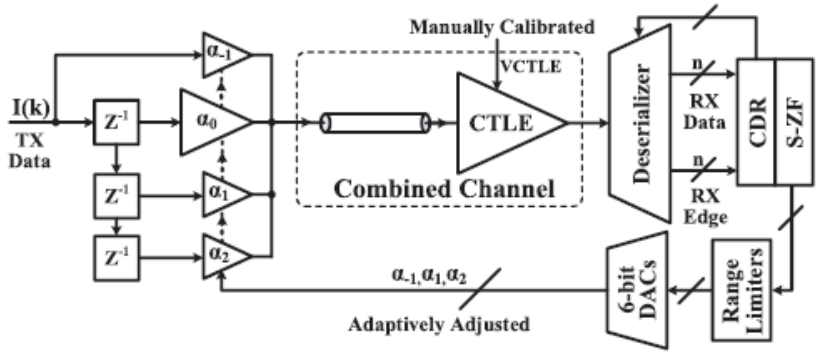
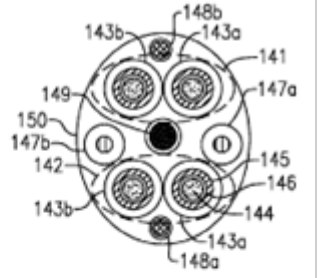
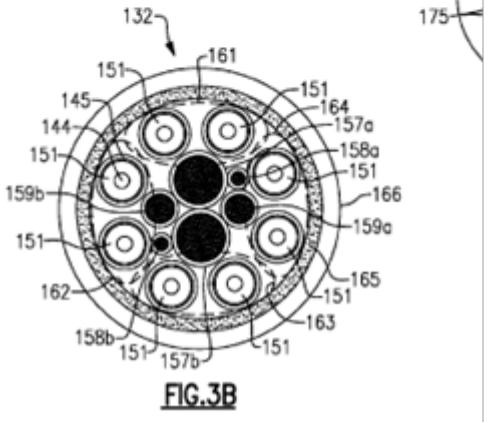


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1019

Claim 13	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.

1020

Claim 13	Disclosure
	 FIG.3A  FIG.3B Lugthart, Figures 3A and 3B. FIG. 3A illustrates a cross-sectional view of one embodiment of a cable 131.

1021

Claim 13	Disclosure
	The illustrated cable 131 includes a first pair of conducting lines 141 and a second pair of conducting lines 142. Each of the pairs of conducting lines 141-142 can be configured for differential signal transmission. For instance, the individual conducting lines in the pair can be implemented coaxially in certain embodiments. In the illustrated embodiment, each pair of conducting lines 141-142 includes a first micro coaxial cable 143 a and a second micro coaxial cable 143 b. In the depicted embodiment, each of the micro coaxial cables includes an inner conductor 144 and an outer conductor 145. An insulating material or a dielectric 146 is disposed between the inner conductor 144 and the outer conductor 145. In various implementations, the insulating material or dielectric 146 can include a polymer or a fluoroplastic material. Lugthart, 17:9–24. For example, FIG. 3B illustrates a cross-sectional view of another embodiment of a cable 132. The cable 132 includes eight individual conducting lines 151, which can be, for example, micro coaxial cables. As depicted by the dashed lines, according to some embodiments, each of the conducting lines 151 can be paired with another of the conducting lines 151 to form a differential pair, and a single corresponding logical channel. For example, the cable 132 can include a first pair of conducting lines 161, a second pair of conducting lines 162, a third pair of conducting lines 163, and a fourth pair of conducting lines 164. Lugthart, 17:63–18:6.

1022

N. DEPENDENT CLAIM 14

Claim 14	Disclosure
14. The method of claim 8, wherein the first and second DRR devices do not perform pre-equalization of the multi-lane data streams provided to the first and second host interface ports.	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.

1023

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.

1024

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.

1025

Claim 14	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.

1026

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.

1027

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.

1028

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.

1029

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.

1030

Claim 14	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) 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	

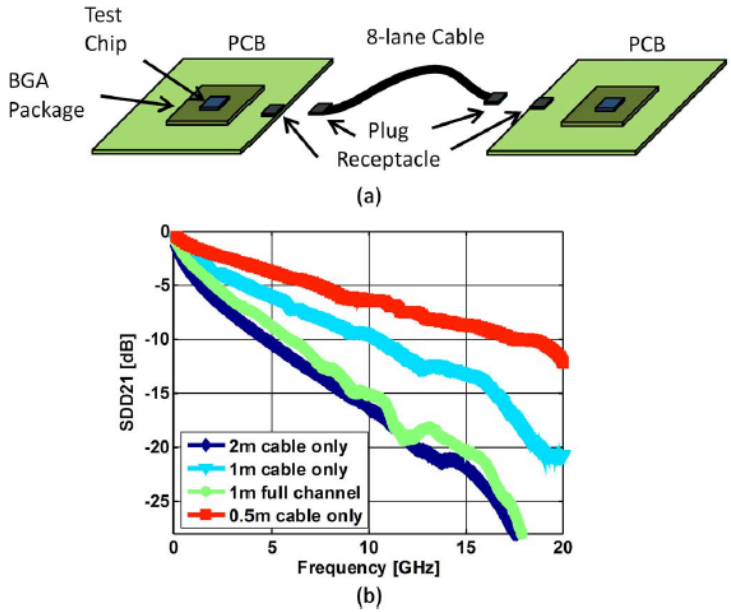
1031

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

1032

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.

1033

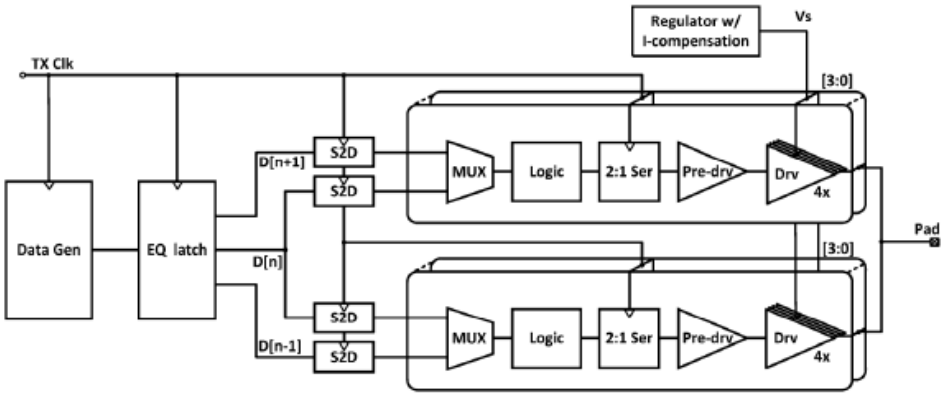
Claim 14	Disclosure
1034	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

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.

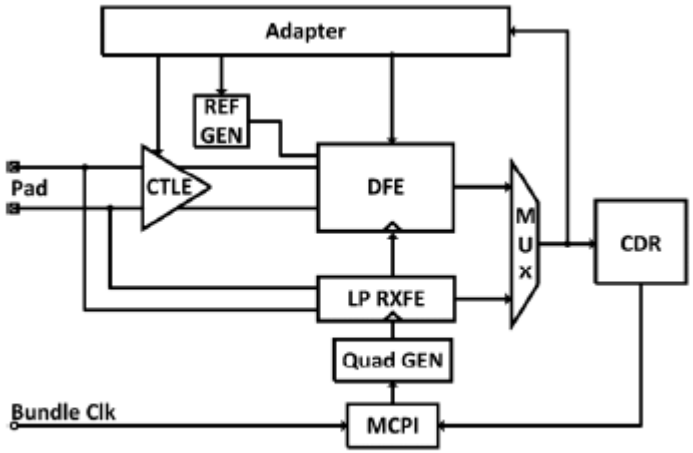
1035

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.

1036

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

1037

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.

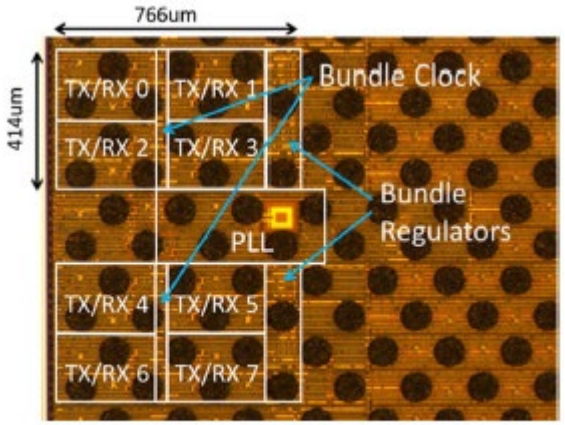
1038

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.

1039

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.

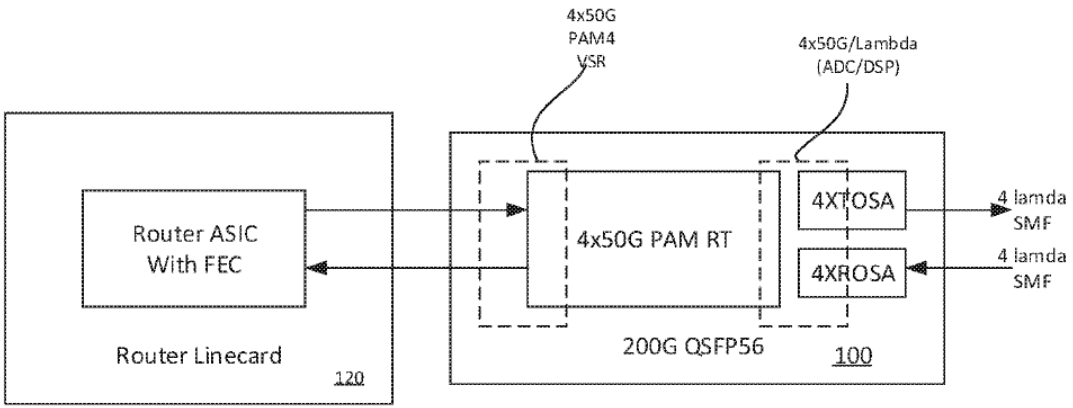
1040

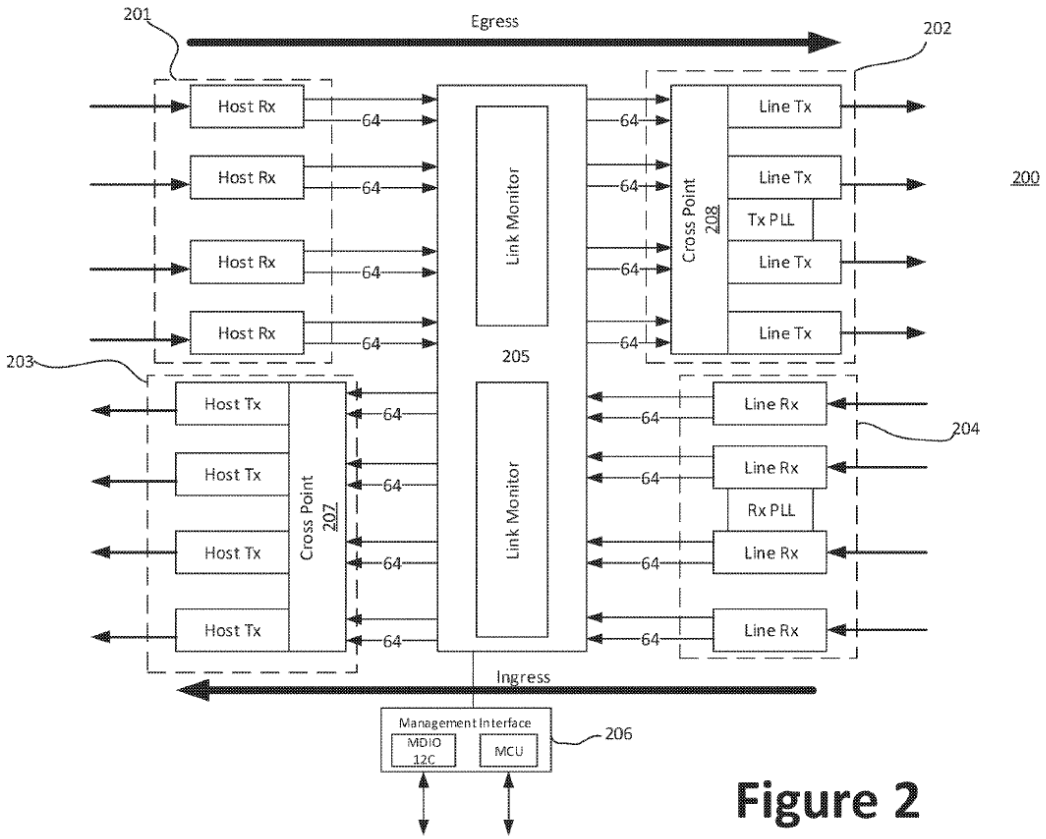
Claim 14	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

1041

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.

1042

Claim 14	Disclosure
1043	 **Figure 1** Riani, Figure 1.

Claim 14	Disclosure
1044	 Figure 2 Riani, Figure 2.

Claim 14	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.

1045

Claim 14	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 14	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.

1047

Claim 14	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.

1048

Claim 14	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.

1049

Claim 14	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.

1050

Claim 14	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.

1051

Claim 14	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.

1052

Claim 14	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.

1053

	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.	

1054

Claim 14	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.

1055

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

1056

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.

1057

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.

1058

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.

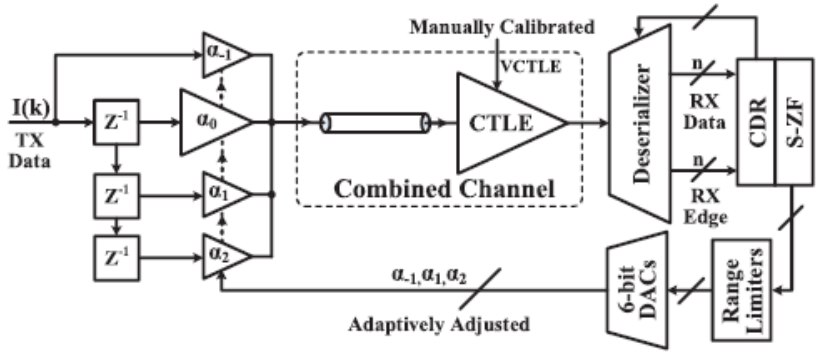


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1059

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.

1060

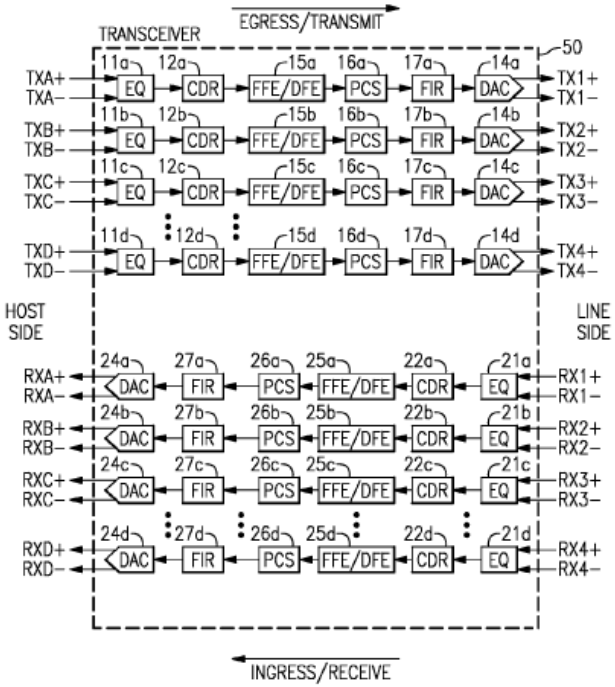
Claim 14	Disclosure
1061	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. 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. 7,233,617 to Gorecki* Gorecki discloses that "tap weights or coefficients (values) may be determined or controlled[.]" *See* Gorecki at 6:56-57. Gorecki discloses the storing of filter coefficient values in EEPROM, and use those values to set filter coefficients to provide pre-equalization. *See* Gorecki at 7:18-33 ("[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM. In this way, the transmitter may access the memory to retrieve the necessary information during start-up/power-up, initialization or re-initialization."). An EEPROM [electrically erasable programmable read-only memory] is non-volatile memory. *See, e.g.*, U.S. Patent App. Pub. No. 2015/0106536 at [0081] ("nonvolatile memory can include… electrically erasable PROM (EEPROM)"); U.S. Patent No. 8,000,176 at 2:53-54 ("EEPROM is a non-volatile memory"); U.S. Patent App. Pub. No. 2005/0013317 at [0012] ("non-

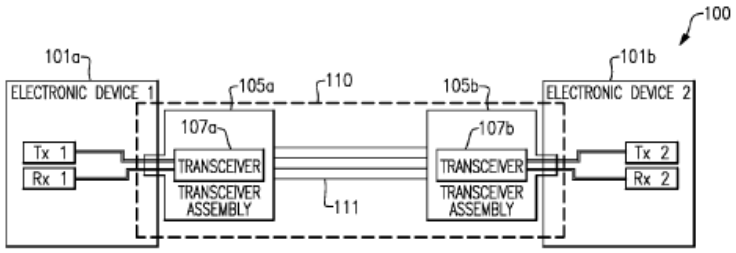
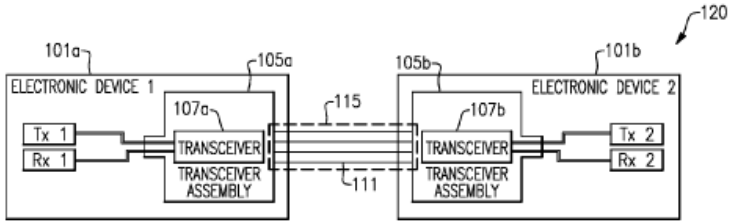
Claim 14	Disclosure
	volatile memory may include… an EEPROM"). The L-G combination teaches programming the coefficients of the FIR filter 17 using the values stored in non-volatile EEPROM to pre-equalize the TX signal in transceiver 107a and 107b. *IEEE 802.3-2015* The IEEE 802.3-2015 standard discloses setting transmit filter coefficients to disable pre-equalization. As was known, an ordinarily skilled artisan would have disabled pre-equalization "if not needed to consume minimum power while satisfying the performance or margin objective" for a system. *See* U.S. Patent No. 9,137,063 to Zerbe at 27:24-37. Thus, for instances where signal transmission between host devices using the active cable meets performance objectives without pre-equalization of the multi-lane data streams provided to the first and second host interface ports, a user would disable those pre-equalization based on the IEEE 802.3-2015 standard.

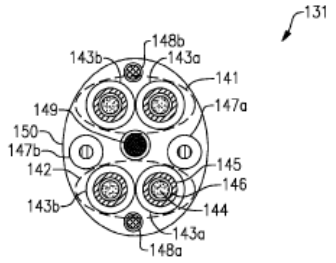
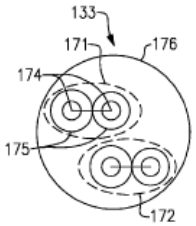
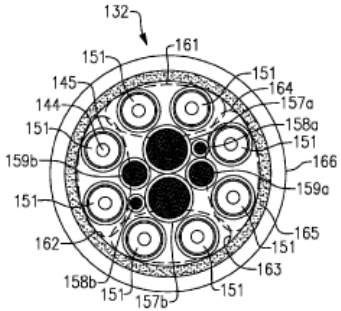
1062

Claim 14	Disclosure
	72.6.10.3.4 Functions COEFF_UPDATE(coefficient, preset, initialize, inc, dec) Returns an updated coefficient based on the contents of the coefficient update field in the training frame. Sets a fixed coefficient value, or adds, or subtracts from the current coefficient value to create the updated coefficient. If multiple actions are requested in the coefficient update field, then the priority is: 1) preset 2) initialize 3) inc/dec Values: preset; If preset is TRUE then the function returns the coefficient value equivalent to no equalization [$c(-1)$ and $c(1)$ coefficients are set to zero, $c(0)$ set to maximum]. initialize; If initialize is TRUE, then the function returns the coefficient value such that the transmit output meets the conditions defined in 72.6.10.4.2. inc; If inc is TRUE then the function returns (coefficient + step). dec; If dec is TRUE then the function returns (coefficient – step). The requirements for the value of step are defined in 72.7.1.10 and Table 72–7. IEEE 802.3-2015, Section 5 at 480.

1063

Claim 14	Disclosure
1064	 FIG.1D Lugthart, Figure 1D.

Claim 14	Disclosure
1065	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

Claim 14	Disclosure
	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3C. According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver

1066

Claim 14	Disclosure
	including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:66–4:5. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive

1067

Claim 14	Disclosure
	path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart 12:1–53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable.

1068

Claim 14	Disclosure
	The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:3–38. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

1069

Claim 14	Disclosure
	[0037] Accordingly, the system of FIG. **1** provides a back-channel using a low speed interface to set the operating parameters of the serializer section **160** of SERDES module **110** and the operating parameters of the LA/CDR module **144** of the transceiver interface module **100**. In some applications, this allows the operating parameters to be change over time which may in turn facilitate adaptive optimization of high data speed transmission characteristics and power consumption used for the communications link. For example, some embodiments may optimize power consumption used for the communications link by reducing the launch amplitude of the signals transmitted across the host channel **103**. By monitoring the eye characteristics of the signal at the SERDES eye monitor **164** and the CDR eye monitor **142** and modifying the transmit equalization parameters used by the LA/CDR module **144** and the serializer section **160** of the SERDES module **110**, respectively, the predetermined target eye characteristics can be maintained through transmit equalization even while launch amplitude is decreased. Some embodiments may also optimize power consumption by deactivating the circuits that perform transmit equalization when the monitored eye characteristics indicate that pre-equalization is unnecessary for maintaining the predetermined target eye characteristics. Tremblay, paragraph 37.

1070

O. INDEPENDENT CLAIM 15

Claim 15	Disclosure
15[pre] A communications 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.

1071

Claim 15	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.

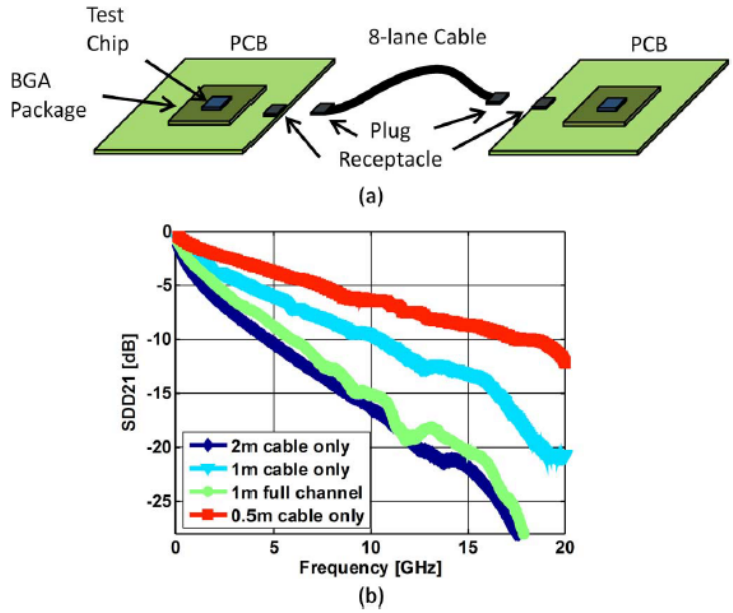
1072

Claim 15	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

1073

Claim 15	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.

1074

Claim 15	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

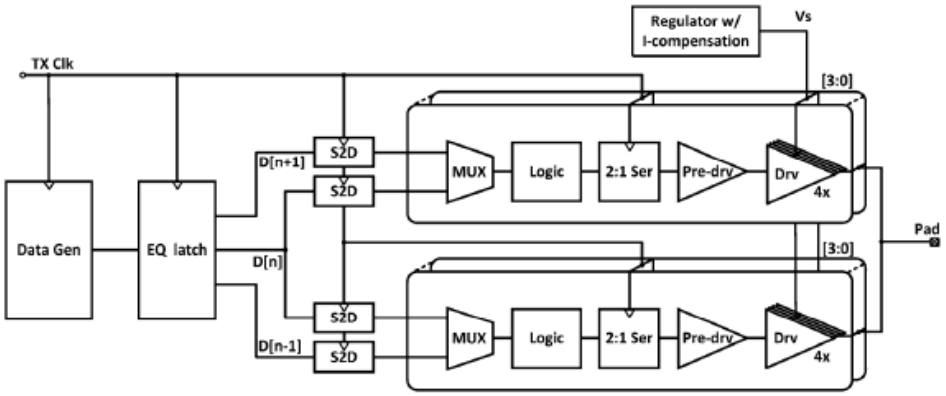
1075

Claim 15	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.

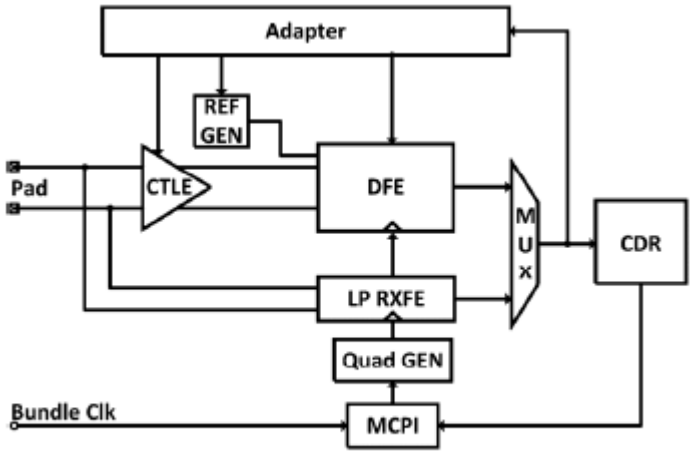
1076

Claim 15	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.

1077

Claim 15	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

1078

Claim 15	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.

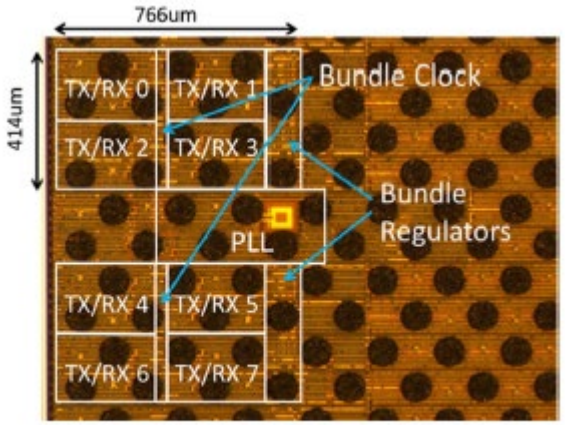
1079

Claim 15	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.

1080

Claim 15	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.

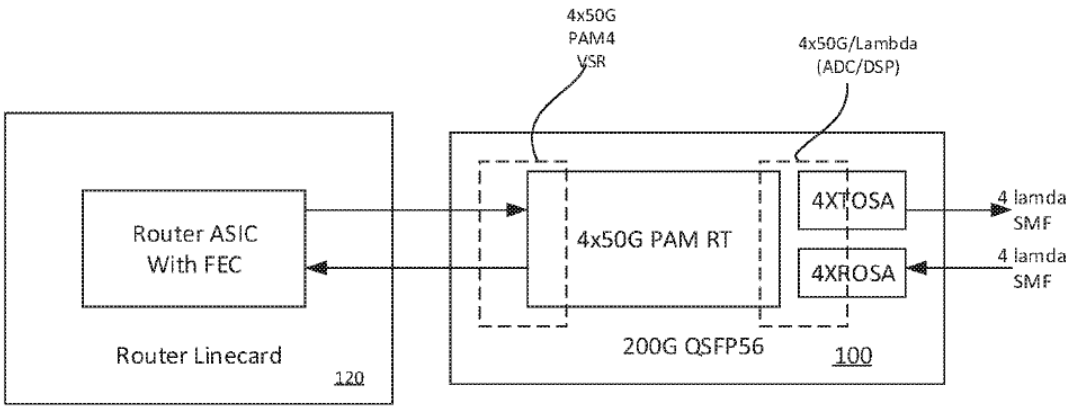
1081

Claim 15	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.

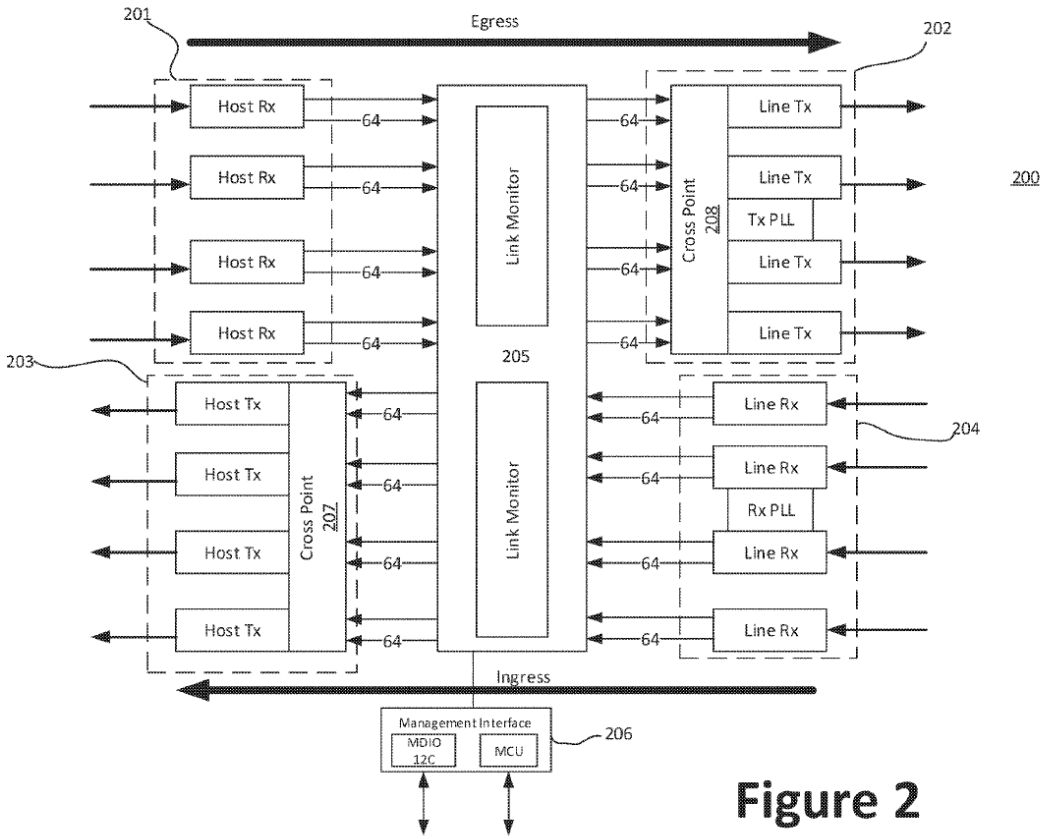
1082

Claim 15	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.

1083

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

1084

Claim 15	Disclosure
1085	 Figure 2 Riani, Figure 2.

Claim 15	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.

1086

Claim 15	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.

1087

Claim 15	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.

1088

Claim 15	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.

1089

Claim 15	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.

1090

Claim 15	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.

1091

Claim 15	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.

1092

Claim 15	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.

1093

Claim 15	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.

1094

	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.	

1095

Claim 15	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.

1096

Claim 15	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.

1097

Claim 15	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.

1098

Claim 15	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

1099

Claim 15	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.

1100

Claim 15	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.

1101

Claim 15	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.

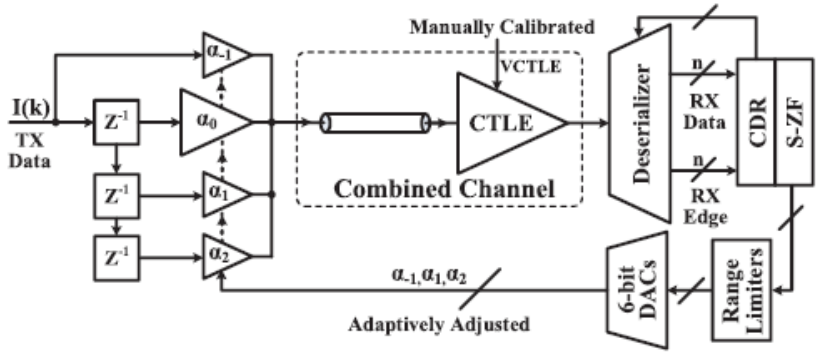
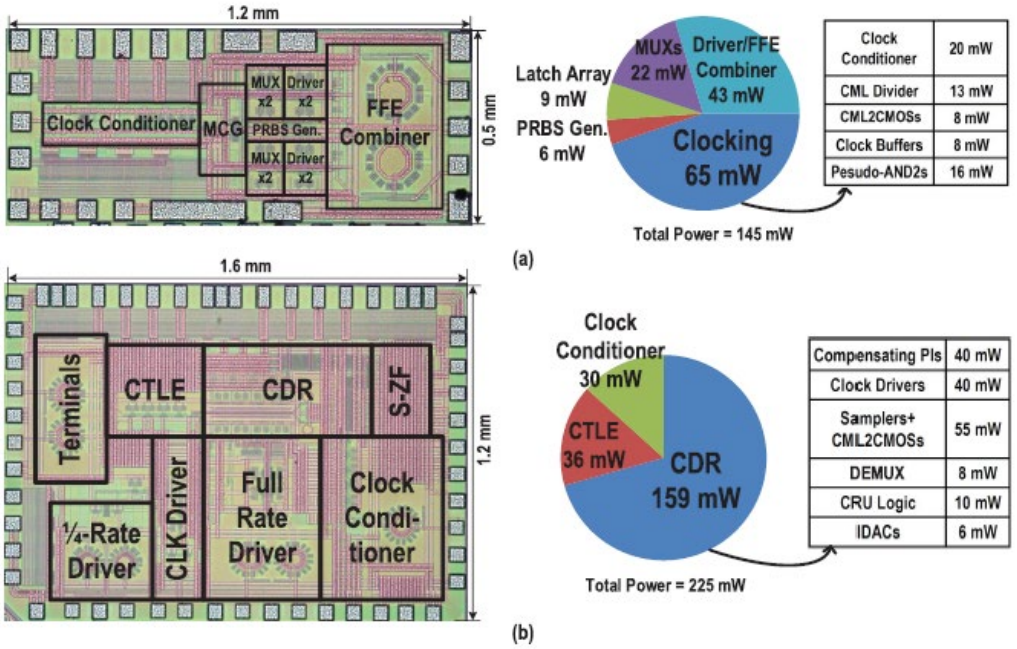


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1102

Claim 15	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.

1103

Claim 15	Disclosure
1104	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

Claim 15	Disclosure
1105	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. Systems and methods for high speed communications are described herein. In certain aspects, the systems and methods include innovative transceiver architectures and techniques for re-timing, multiplexing, de-multiplexing and transmitting data. The systems and methods can be used to achieve reliable high-speed point-to-point communication between different electronic devices, computing devices, storage devices and peripheral devices. Lugthart, Abstract. Circuits, methods, and apparatus that provide for the powering of active components in connector inserts at each end of a cable may in various ways. For example, where a host is coupled to a device that is not self-powered, the host may provide power for circuitry at each end of the cable. In various embodiments

Claim 15	Disclosure
	of the present invention, the device may request higher voltage from the host, such that more power can be delivered. In these cases, the device may regulate the voltage received from the host to a lower voltage, and then provide the lower voltage to circuitry at one or both ends of the cable. Where the host is connected to a device that is self-powered, the host and the self-powered device may power their respective connector insert circuits. Baker, Abstract. 0027. In high speed data networks a host network element Such as a Switch or router (also referred to as a host) can be connected to a transmission medium (for example fibre cable, copper cable or other physical medium) through an interface module that processes data coming in from the transmission medium to place it into a condition suitable for the host and that processes data from the host to place it into a condition suitable for the transmission medium. Tremblay, paragraph 27. 0029. Example embodiments will now be described with reference to FIG. 1 which illustrates a host 102 that is con nected through a host channel 103 to a pluggable network transceiver interface module 100, which in turn is connected to a transmission medium 104. In the illustrated embodiment, the transmission medium 104 includes fibre optic cable 124 for transmitting high speed data from the interface module 100 and fibre optic cable 126 for transmitting high speed data to the interface module 100, however embodiments described herein can also be applied to non-fibre transmission mediums Such as copper for example. Tremblay, paragraph 29.
15[a] inserting a first end connector plug of a cable into a first host interface	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:

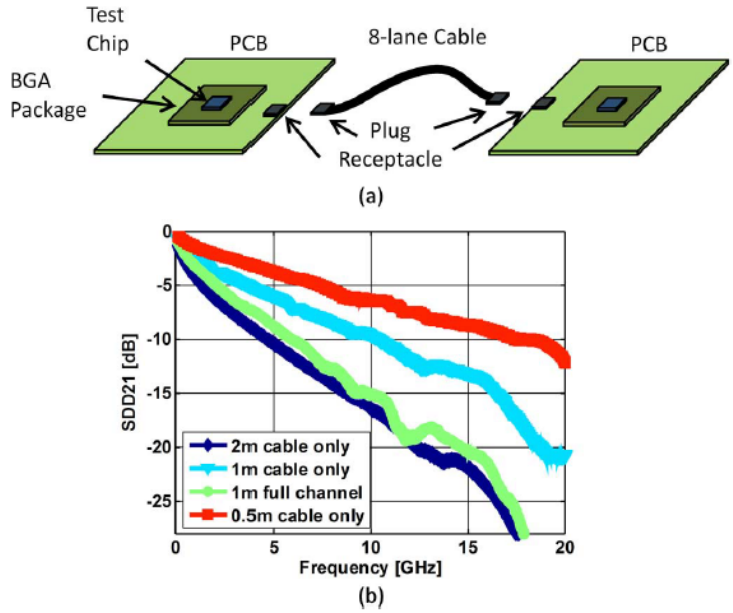
1106

Claim 15	Disclosure
port, the first end connector plug being connected to a first data recovery and re-modulation (DRR) device that converts a multi-lane data stream from the first host interface port into electrical transit signals conveyed by electrical conductors to a second DRR device connected to a second end connector plug of the cable; 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

1107

Claim 15	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.

1108

Claim 15	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

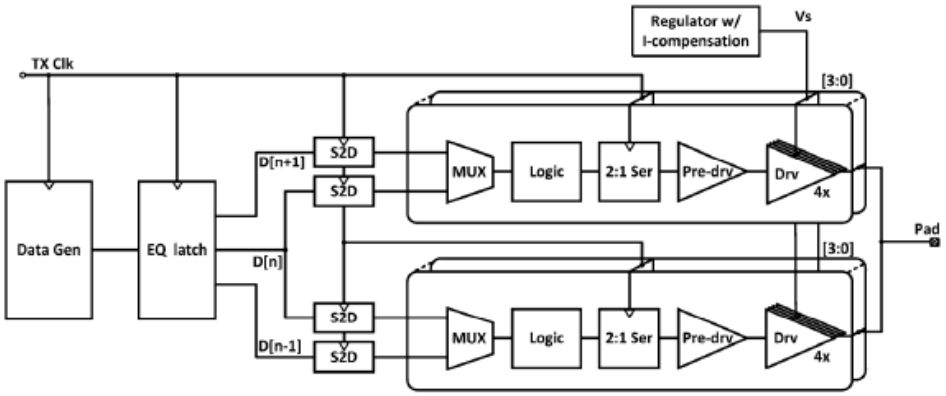
1109

Claim 15	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.

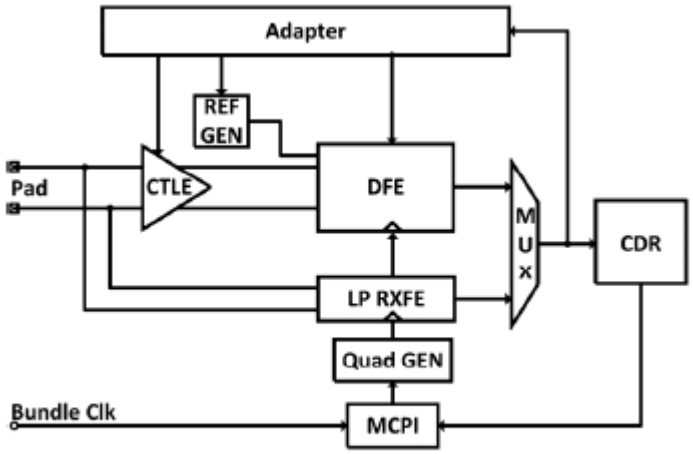
1110

Claim 15	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.

1111

Claim 15	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

1112

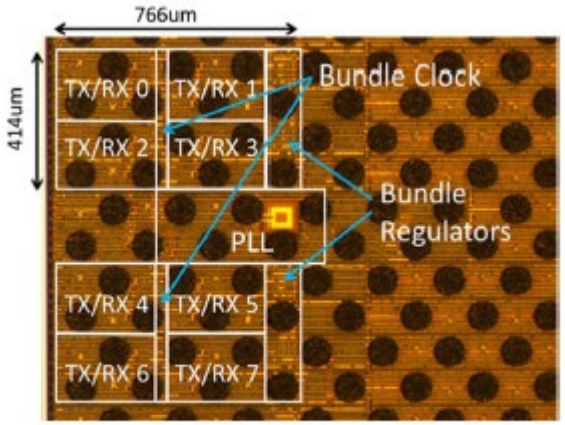
Claim 15	Disclosure
1113	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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 15	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.

1114

Claim 15	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.

1115

Claim 15	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.

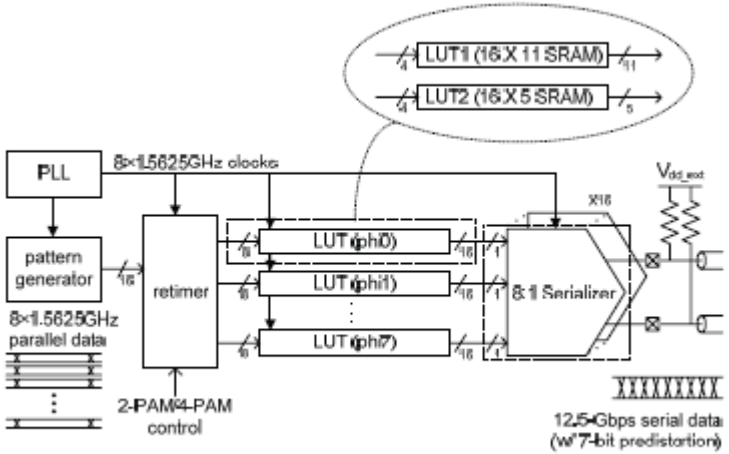
1116

Claim 15	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.

1117

Claim 15	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.

1118

Claim 15	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

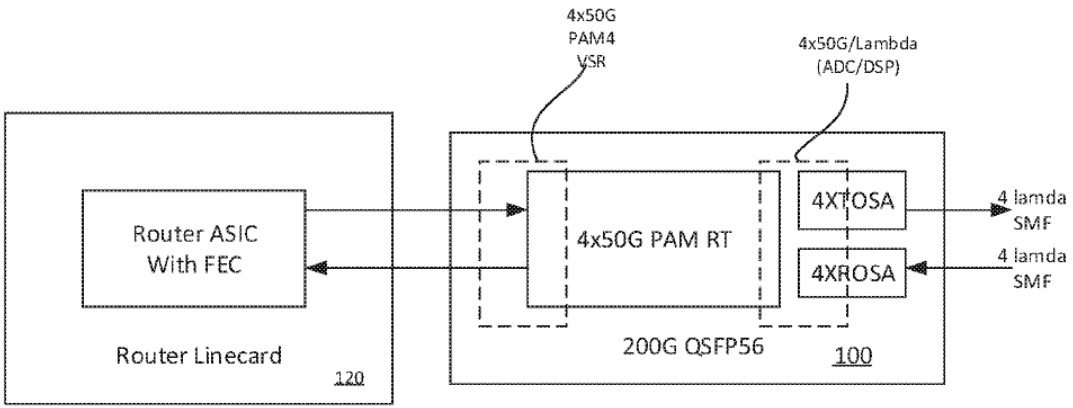
1119

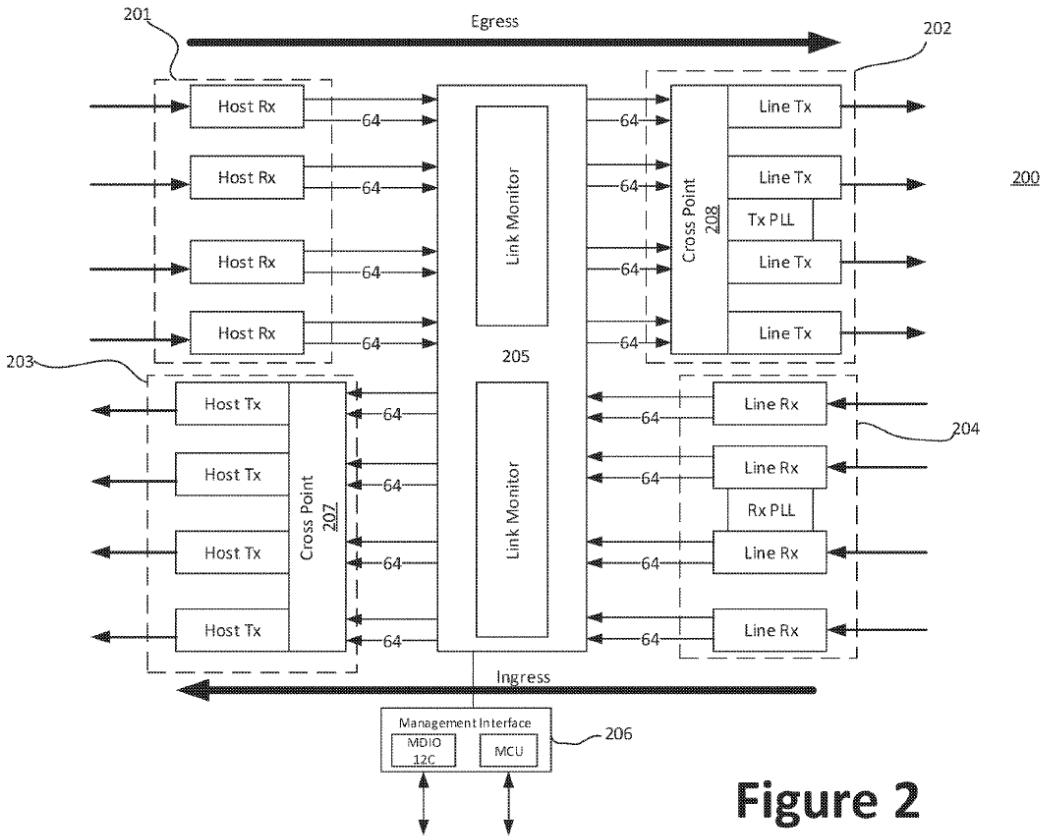
Claim 15	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

1120

Claim 15	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.

1121

Claim 15	Disclosure
1122	 **Figure 1** Riani, Figure 1.

Claim 15	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 15	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.

1124

Claim 15	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.

1125

Claim 15	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.

1126

Claim 15	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.

1127

Claim 15	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.

1128

Claim 15	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.

1129

Claim 15	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.

1130

Claim 15	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.

1131

Claim 15	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.

	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.	

1133

Claim 15	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.

1134

Claim 15	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.

1135

Claim 15	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.

1136

Claim 15	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

1137

Claim 15	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.

1138

Claim 15	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.

1139

Claim 15	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.

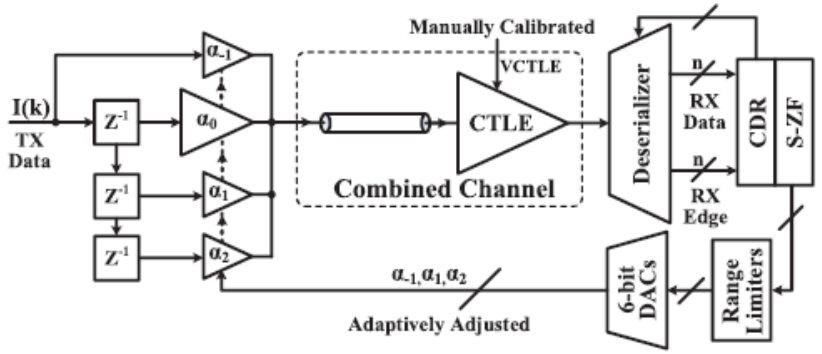
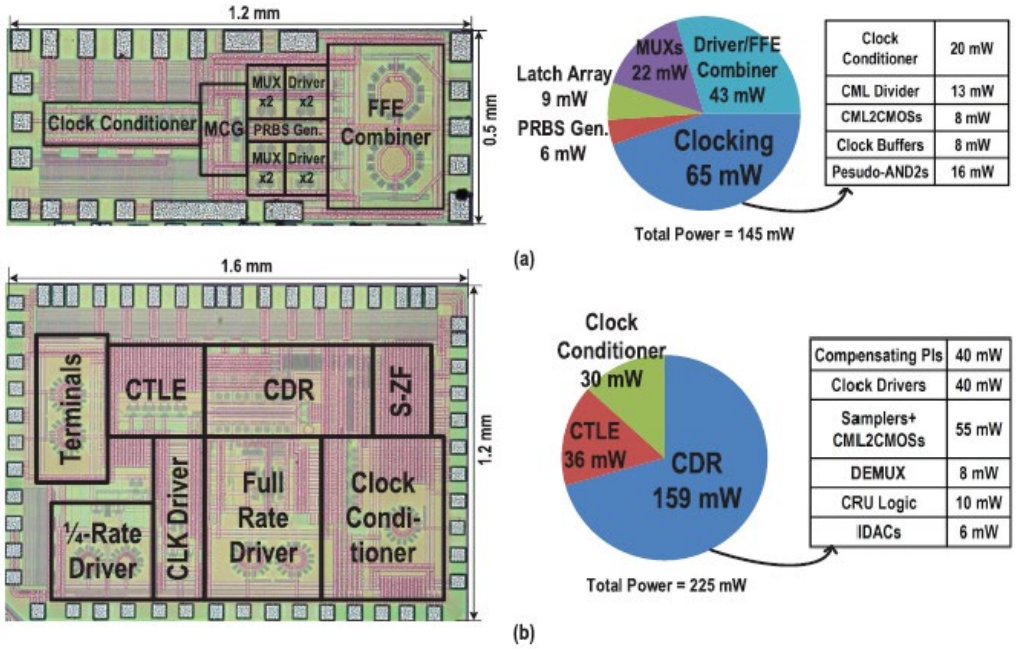


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1140

Claim 15	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.

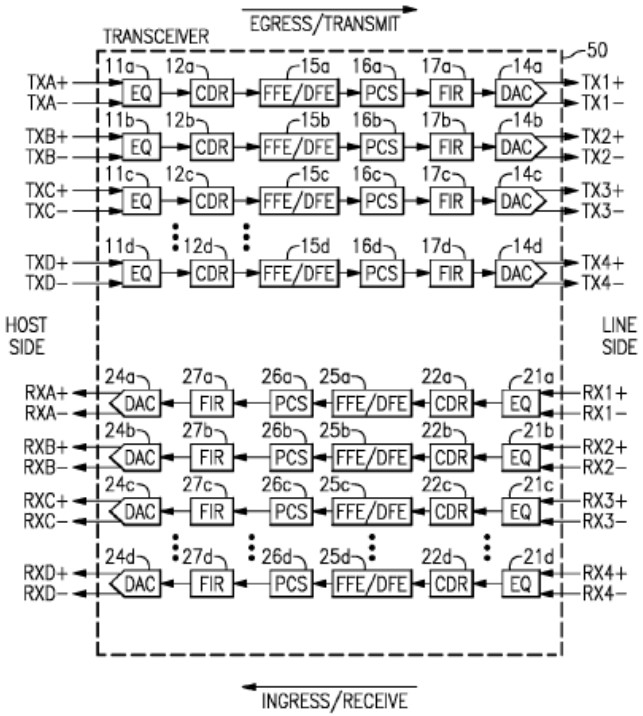
1141

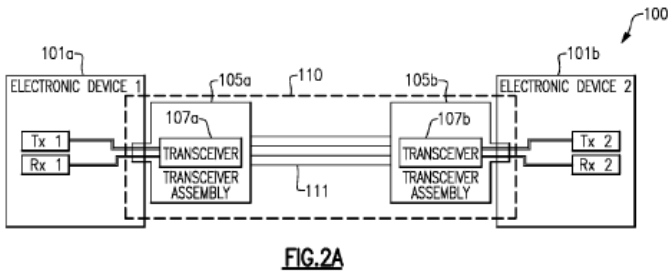
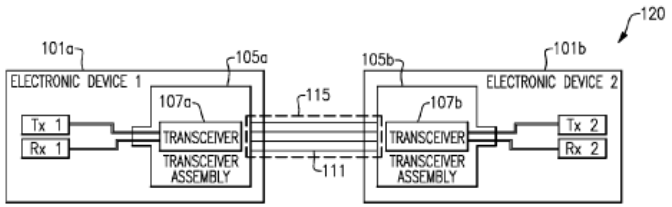
Claim 15	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

1142

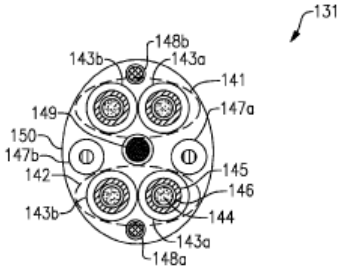
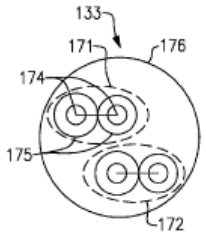
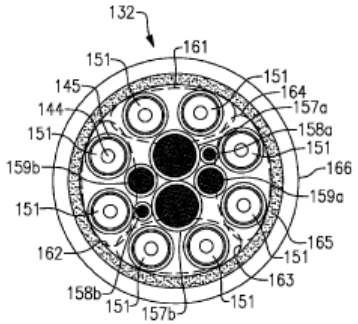
Claim 15	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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.

1143

Claim 15	Disclosure
	 FIG.1D Lugthart, Figure 1D.

Claim 15	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

1145

Claim 15	Disclosure
	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3B.

1146

Claim 15	Disclosure
	According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:66–4:5. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first

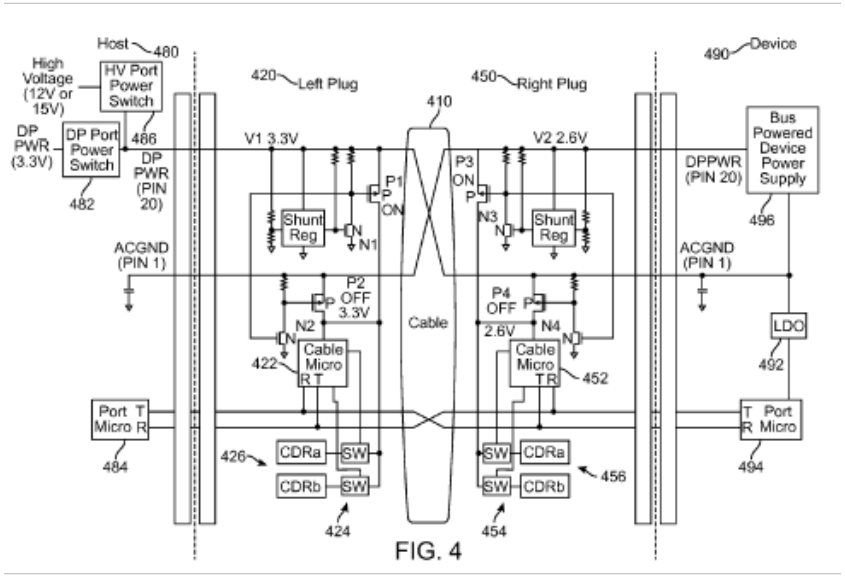
1147

Claim 15	Disclosure
	receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1–53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111.

1148

Claim 15	Disclosure
	Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:3–38. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

1149

Claim 15	Disclosure
1150	 FIG. 4 Baker, Figure 4. FIG. 4 illustrates an electronic system according to an embodiment of the present invention. This figure includes host 480 coupled to device 490 via cable 410. Cable 410 includes left plug 424 connected to host 480, and right plug 454 connected to device 490. Host 480 may be capable of providing one or more voltages to both device 490 and cable circuitry in cable 410. Host 480 may provide a lower voltage of 3.3

Claim 15	Disclosure
	V, or a higher voltage of 12 or 15 volts. In other embodiments of the present invention, host 480 may provide various voltage levels to host 490 and cable circuitry in cable 410. Baker, 5:26–36. Voltage V1 is received by low-drop-out regulator 492 in device 490, which provides power to port microcontroller 494. The host port microcontroller 484 may then communicate with cable microcontrollers s 422 and 452 and port microcontroller 494 to determine proper configuration for the cable. In a specific embodiment of the present invention, the host port microcontroller 484 may check with device port microcontroller 494 to determine whether device 490 requires a higher level of power. If it does, host port microcontroller 484 may check with cable microcontroller 422 to determine whether the cable can support the delivery of this higher level of power. If device 490 requires higher power, and the cable can deliver it, then host 480 may provide the higher level of power. In another embodiment of the present invention, host port microcontroller 484 may determine how much power will be needed by the cable and device 490. In some circumstances, one link including one pair of clock and data recovery circuits, or other circuitry, may need to be powered off. Baker, 5:49—67. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker, 6:1–12.

1151

Claim 15	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. Tremblay, paragraph 31.
15[b] inserting the second end connector plug of the cable into a second host interface port, the second DRR device converting a multi-lane data stream from the second host interface port into electrical transit signals conveyed by electrical conductors to the first DRR 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:

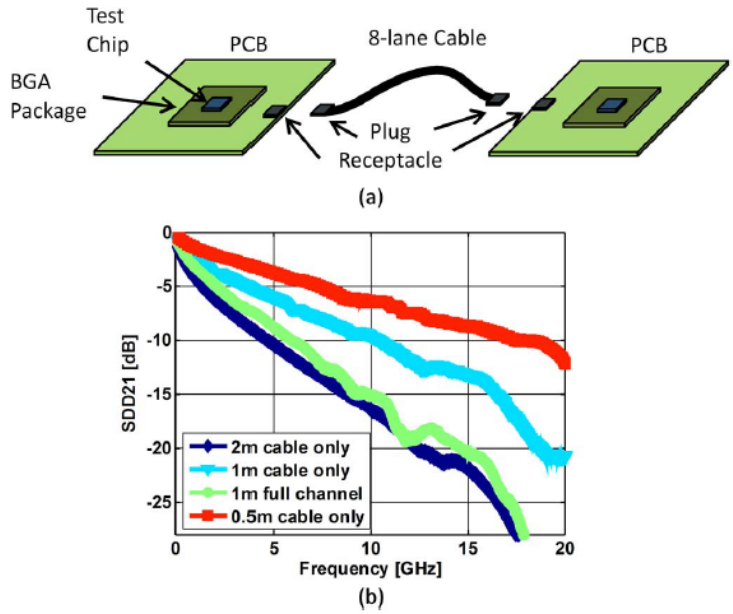
1152

Claim 15	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

1153

Claim 15	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.

1154

Claim 15	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

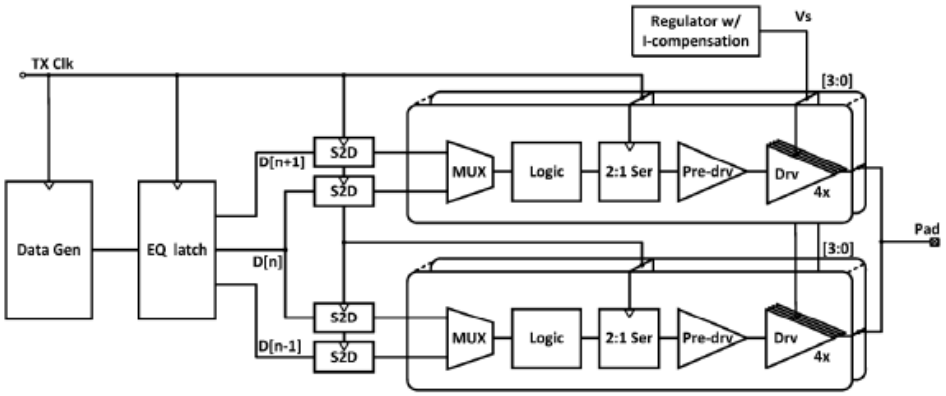
1155

Claim 15	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.

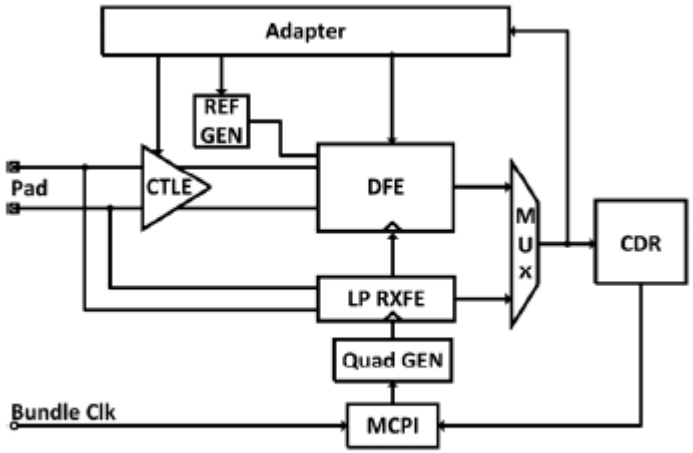
1156

Claim 15	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.

1157

Claim 15	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

1158

Claim 15	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.

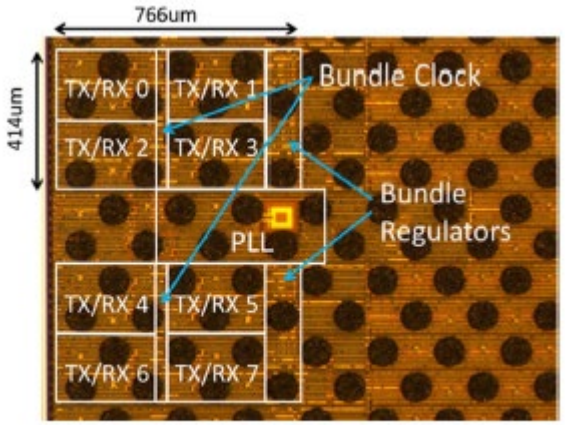
1159

Claim 15	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.

1160

Claim 15	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.

1161

Claim 15	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.

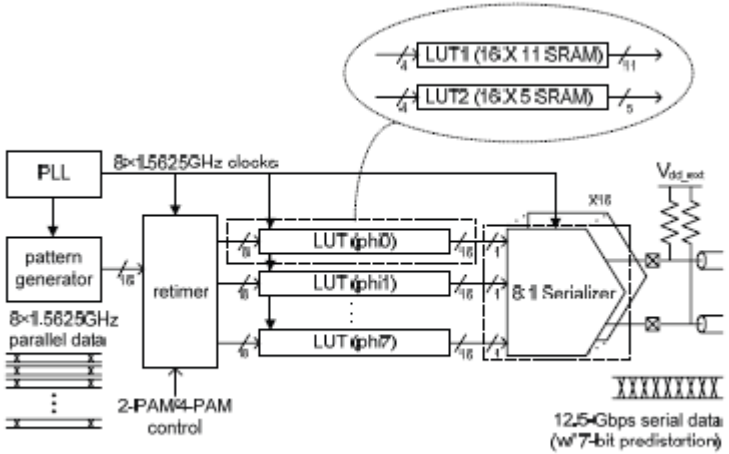
1162

Claim 15	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.

1163

Claim 15	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.

1164

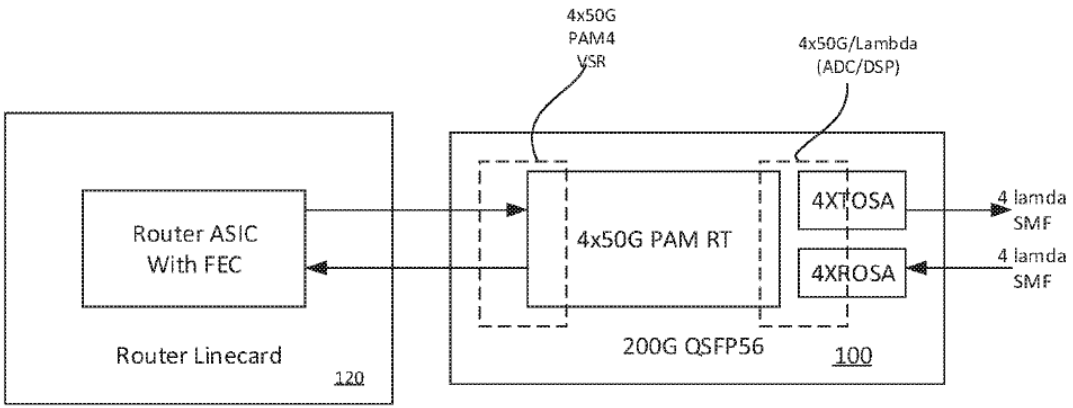
Claim 15	Disclosure
1165	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

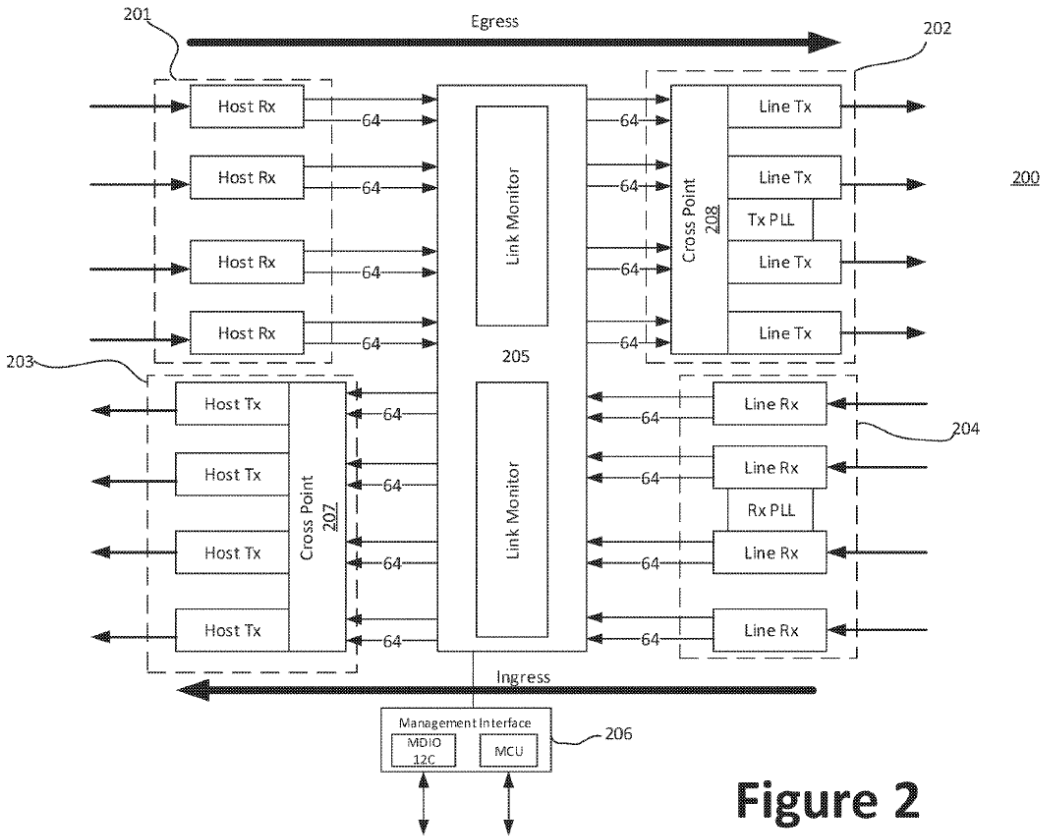
Claim 15	Disclosure
	Fig. 1 presents the block diagram of an 8-way interleaved transmitter architecture. It uses a ring-oscillator-based phase-locked loop (PLL) to generate 8 evenly-spaced 1.5625-GHz clock phases for 8:1 serialization. The retimer block samples and rearranges parallel data from a pseudo-random 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

1166

Claim 15	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.

1167

Claim 15	Disclosure
1168	 **Figure 1** Riani, Figure 1.

Claim 15	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 15	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.

1170

Claim 15	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 15	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.

1172

Claim 15	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.

1173

Claim 15	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.

1174

Claim 15	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.

1175

Claim 15	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.

1176

Claim 15	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.

1177

Claim 15	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.

1178

	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 15	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.

1180

Claim 15	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.

1181

Claim 15	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.

1182

Claim 15	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

1183

Claim 15	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.

1184

Claim 15	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.

1185

Claim 15	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.

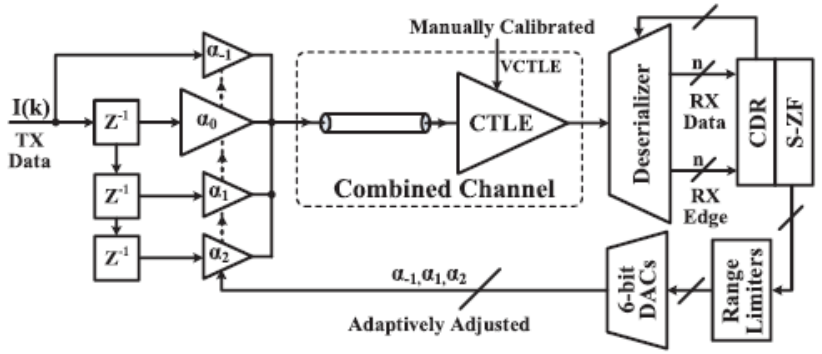
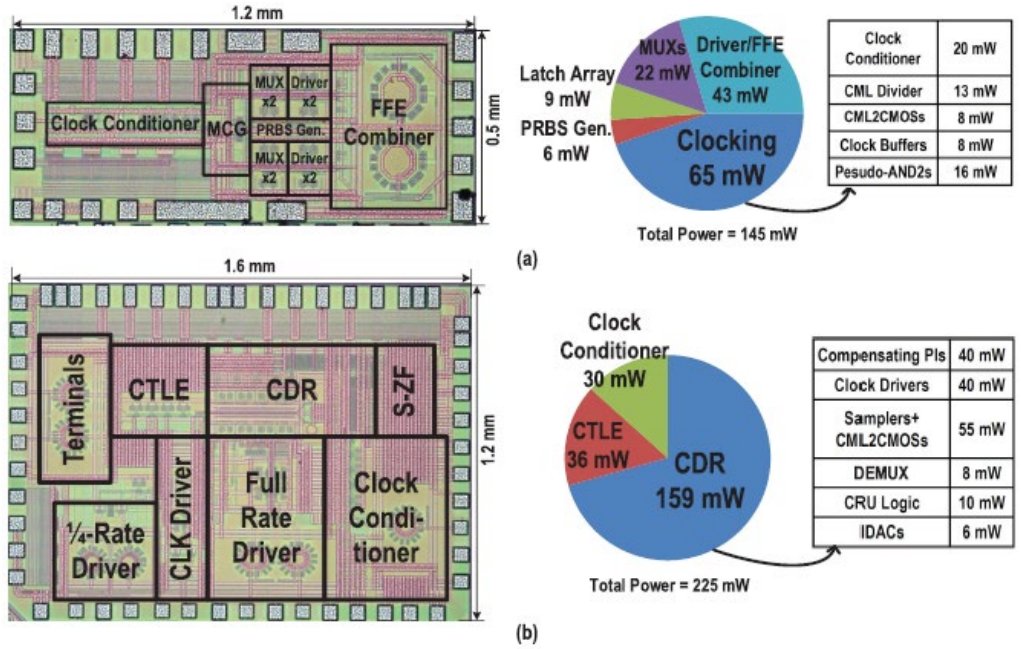


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1186

Claim 15	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.

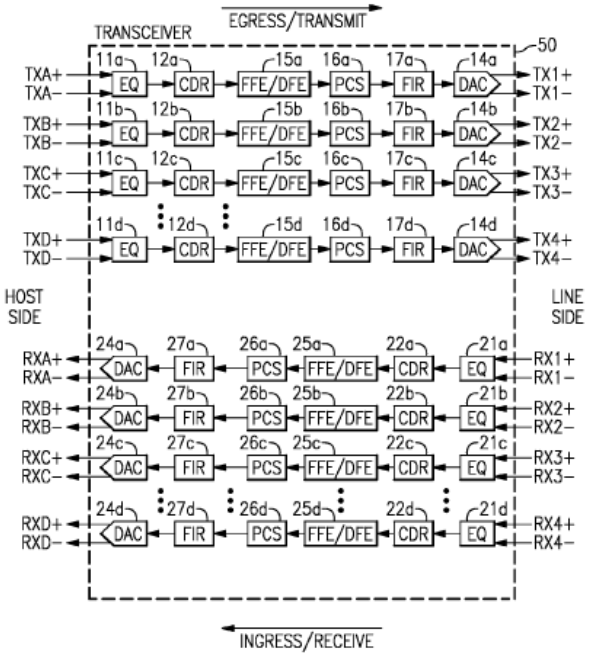
1187

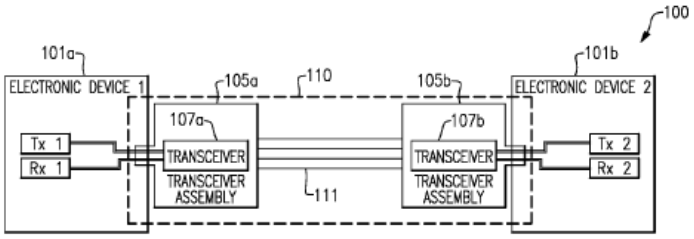
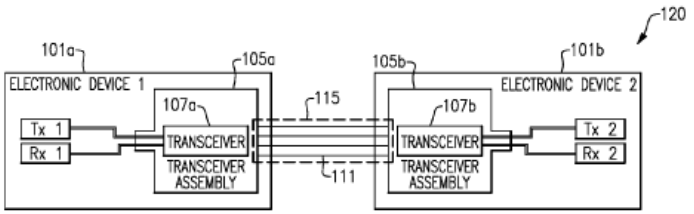
Claim 15	Disclosure
	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

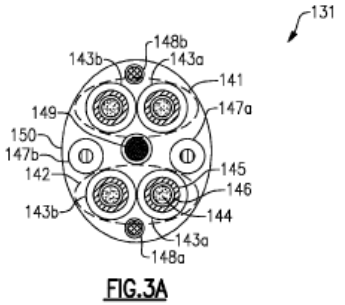
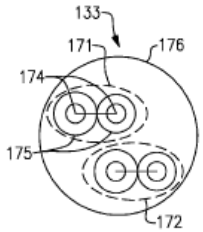
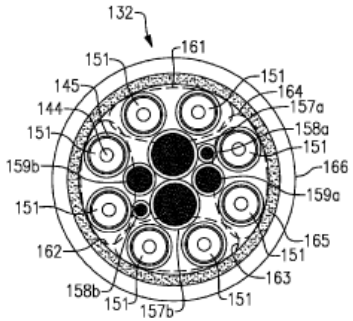
1188

Claim 15	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.

1189

Claim 15	Disclosure
1190	 FIG.1D Lugthart, Figure 1D.

Claim 15	Disclosure
1191	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

Claim 15	Disclosure
	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A–3C. According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals.

1192

Claim 15	Disclosure
	Lugthart, 3:66–4:5. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS

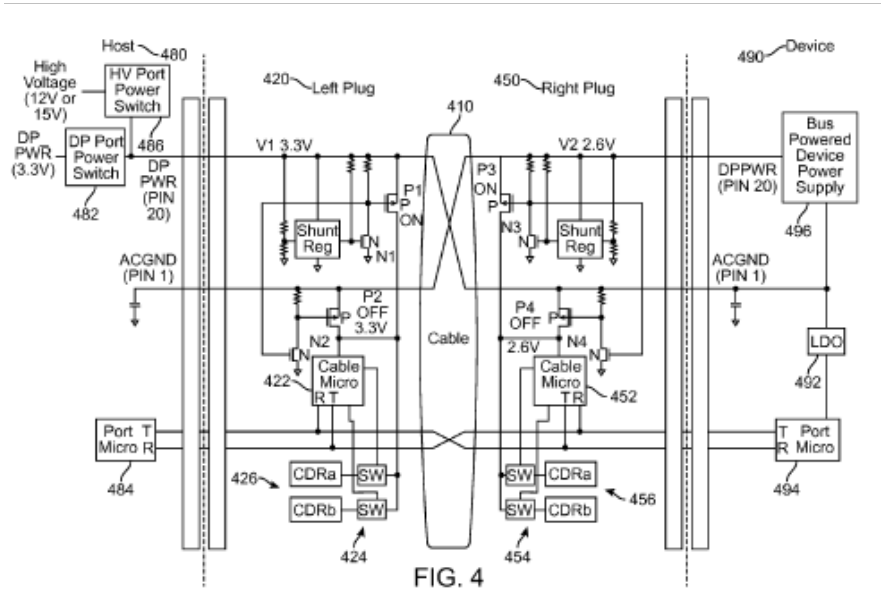
1193

Claim 15	Disclosure
	block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1–53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable. The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b,

1194

Claim 15	Disclosure
	respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:3–38. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

1195

Claim 15	Disclosure
1196	 FIG. 4 Baker, Figure 4. FIG. 4 illustrates an electronic system according to an embodiment of the present invention. This figure includes host 480 coupled to device 490 via cable 410. Cable 410 includes left plug 424 connected to host 480, and right plug 454 connected to device 490. Host 480 may be capable of providing one or more

Claim 15	Disclosure
	voltages to both device 490 and cable circuitry in cable 410. Host 480 may provide a lower voltage of 3.3 V, or a higher voltage of 12 or 15 volts. In other embodiments of the present invention, host 480 may provide various voltage levels to host 490 and cable circuitry in cable 410. Baker, 5:26–36. Voltage V1 is received by low-drop-out regulator 492 in device 490, which provides power to port microcontroller 494. The host port microcontroller 484 may then communicate with cable microcontrollers s 422 and 452 and port microcontroller 494 to determine proper configuration for the cable. In a specific embodiment of the present invention, the host port microcontroller 484 may check with device port microcontroller 494 to determine whether device 490 requires a higher level of power. If it does, host port microcontroller 484 may check with cable microcontroller 422 to determine whether the cable can support the delivery of this higher level of power. If device 490 requires higher power, and the cable can deliver it, then host 480 may provide the higher level of power. In another embodiment of the present invention, host port microcontroller 484 may determine how much power will be needed by the cable and device 490. In some circumstances, one link including one pair of clock and data recovery circuits, or other circuitry, may need to be powered off. Baker 5:49–67. In this example, power supply 496 receives only 3.3 V on line V1 from host 480. At this voltage, power supply 496 may be in an under-voltage lock-out state, and may therefore be powered off. In this state, power supply 496 does not provide power to the cable circuitry. In this example, the cable plug circuitry includes clock and data recovery circuits 426 and 456. These clock and data recovery circuits may receive and retime data received from host 480, device 490, and from each other. Examples of this can be found in co-pending U.S. patent application Ser. No. 13/173,739, titled Circuitry for Active Cable, which is incorporated by reference. Baker 6:1–12.

1197

Claim 15	Disclosure
	[0031] In the illustrated embodiment, the pluggable transceiver module **100** includes a transmit section **120** and a receive section **122** for processing outgoing and incoming high speed data respectively. The transmit section **120** includes a CDR ("Clock and Data Recovery Module") module **128**, an LD ("Laser Driver") module **130** and a laser diode module **140**. The CDR module **128** performs clock and data recovery on the high speed data stream received through connector **118** from the host channel **103**, and LD module **130** and laser diode module **140** modulate the resulting high speed data stream onto fibre optic cable **124**. In embodiments having multiple channels in the form of multiple receiving conductors **114**, the CDR module **128** performs these operations for each channel of the conductors **114**. Tremblay, paragraph 31.
15[c] the first and second DRR devices providing pre-equalization of the electrical transit signals using transmit filter coefficient values stored in nonvolatile memories.	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.

1198

Claim 15	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.

1199

Claim 15	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.

1200

Claim 15	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.

1201

Claim 15	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.

1202

Claim 15	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.

1203

Claim 15	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.

1204

Claim 15	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.

1205

Claim 15	Disclosure

1206

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	

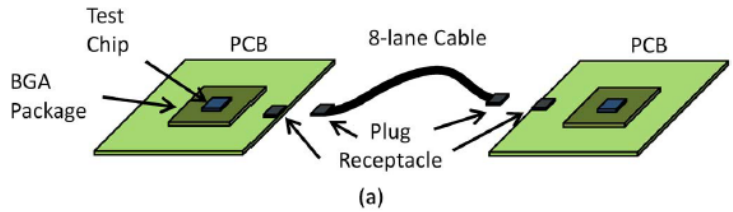
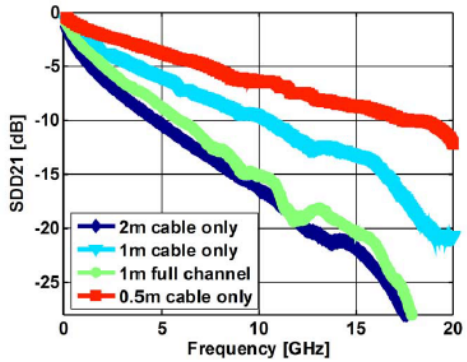
CMIS, 71.

Claim 15	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

1207

Claim 15	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.

1208

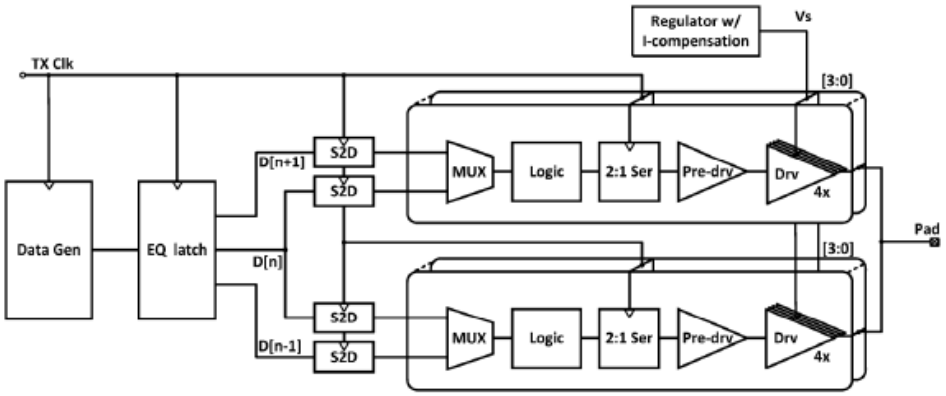
Claim 15	Disclosure
1209	 (a)  (b) Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 15	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.

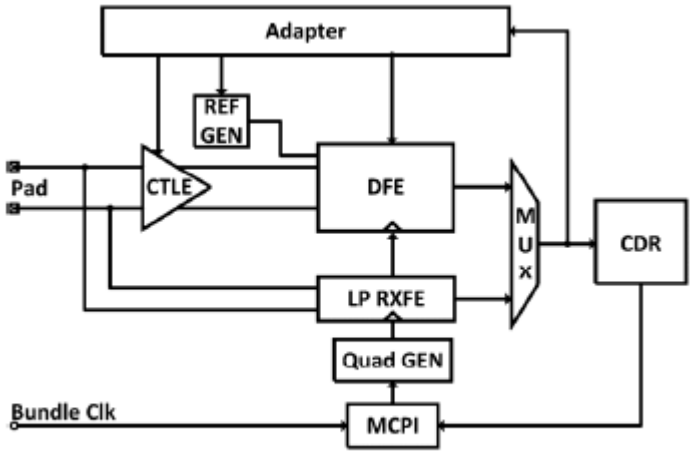
1210

Claim 15	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.

1211

Claim 15	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

1212

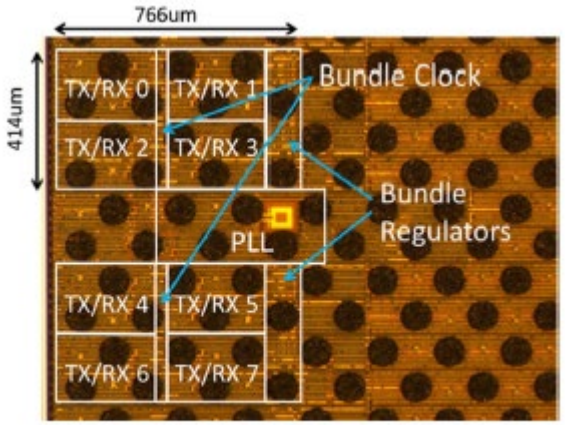
Claim 15	Disclosure
1213	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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 15	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.

1214

Claim 15	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.

1215

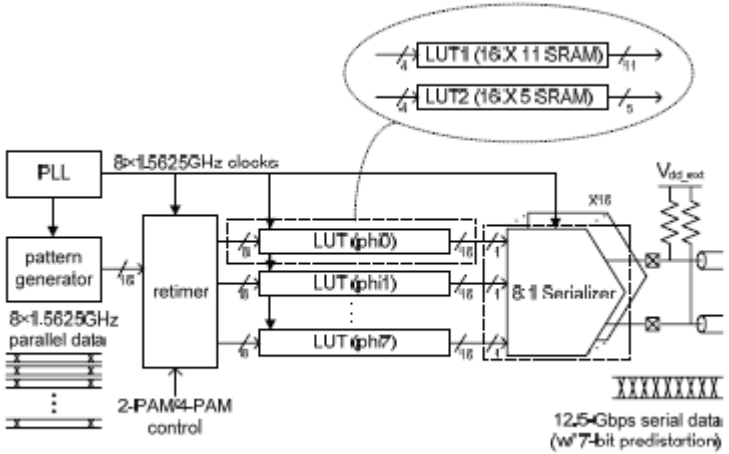
Claim 15	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses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

1216

Claim 15	Disclosure
1217	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 15	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.

1218

Claim 15	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1219

Claim 15	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

1220

Claim 15	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.

1221

Claim 15	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.

1222

Claim 15	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.

1223

Claim 15	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

1224

Claim 15	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.

1225

Claim 15	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.

1226

Claim 15	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.

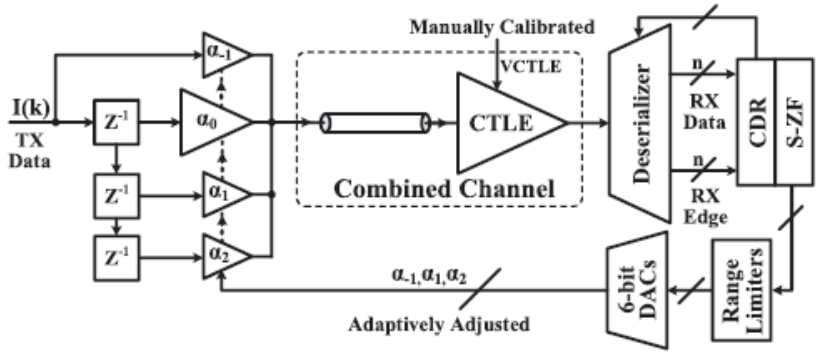
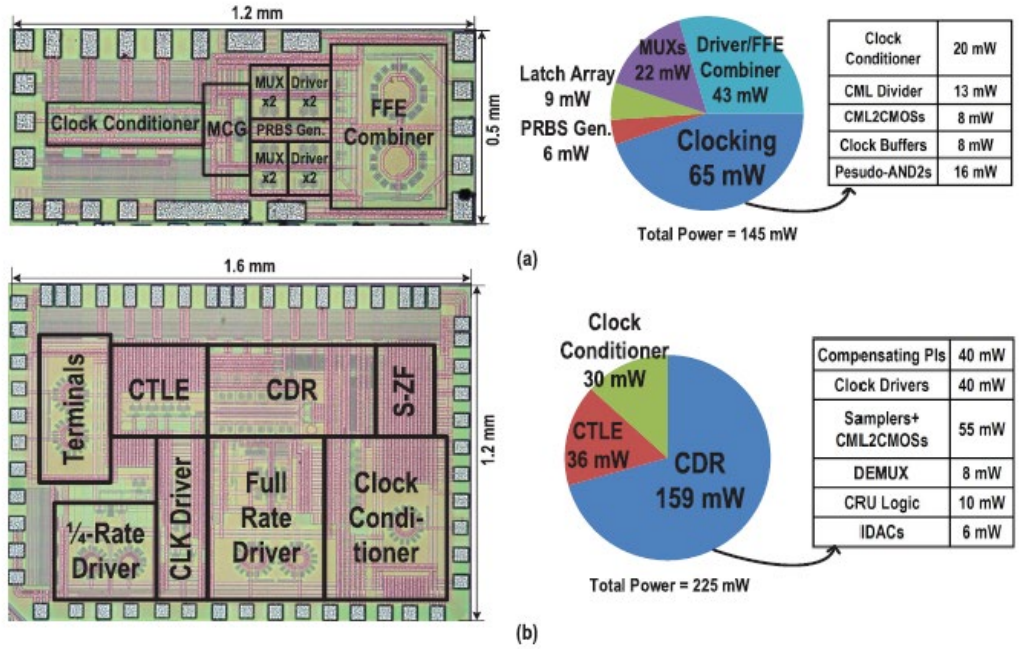


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1227

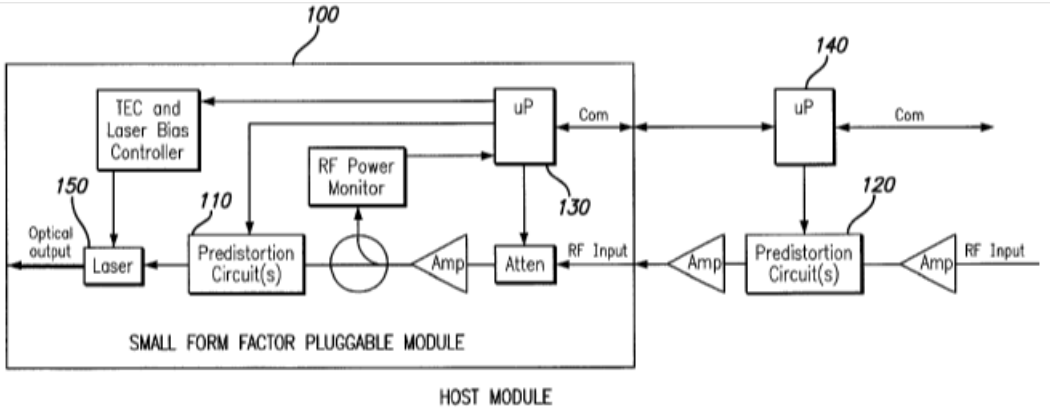
Claim 15	Disclosure
	Zheng, 2970. The TX and RX chips are fabricated in a 65-nm CMOS process. The chips are mounted on PCBs through wire bonding and they are connected to the testing instruments via SMA connectors and connection cables. Fig. 22 shows the micrographs and power breakdown when applying 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.

1228

Claim 15	Disclosure
1229	 Fig. 22. Micrographs and power breakdown of (a) TX chip and (b) RX chip. Zheng, 2973.

Claim 15	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.

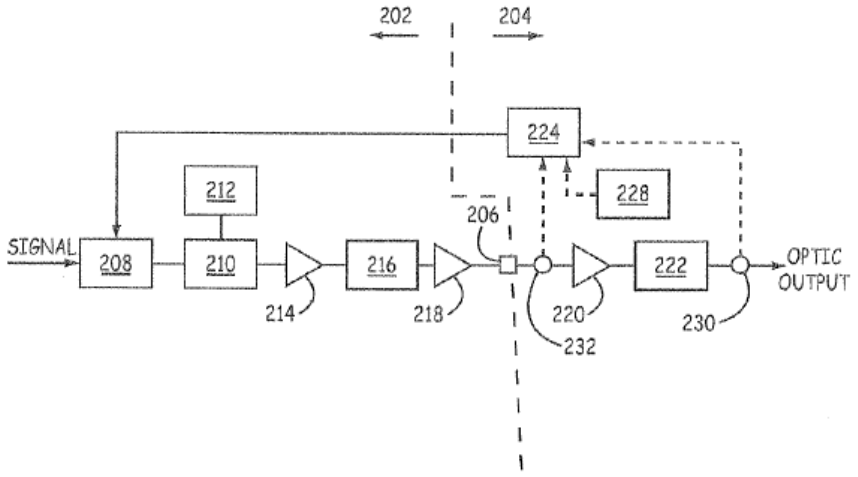
1230

Claim 15	Disclosure
	 FIG. 1 Ishaug, Figure 1.

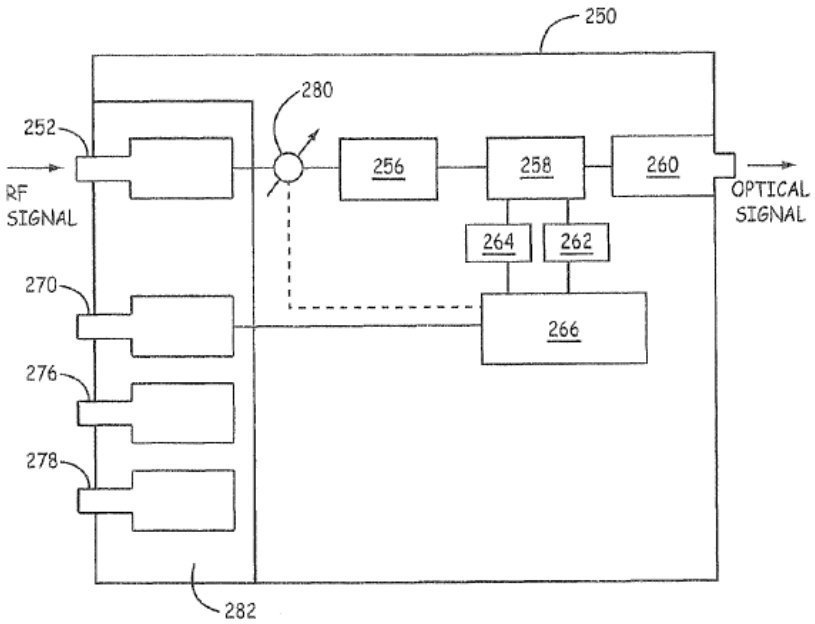
1231

Claim 15	Disclosure
	A method includes reading operation parameters from a non-volatile memory located in a pluggable module that is coupled to a host module; processing the operational parameters with a processor located in the host module to control operation of a predistortion circuit located in the host module; adding predistortion to a signal with the predistortion circuit located in the host module and then sending the predistorted signal to the pluggable module. An apparatus includes a host module including a predistortion circuit and a processor coupled to the predistortion circuit; and a pluggable module coupled to the host module, wherein the pluggable module includes a non-volatile memory containing operational parameters for the predistortion circuit of the host module, wherein the operational parameters are processed by the processor of the host module to control the predistortion circuit of the host module. Ishaug, Abstract. As shown in FIG. 1, both the pluggable module and the host module have predistortion circuits 110, 120. However, the pluggable module does not need to contain a predistortion circuit. The pluggable module includes a processor 130 having non-volatile memory. In addition, the host module is shown as a standalone unit including a processor 140 having non-volatile memory. However, the host may be incorporated into a larger system that provides the analog RF input. There is also gain, attenuation and RF power monitoring shown in the pluggable and host modules that are not required, but may be needed to set the drive level to the laser modulation at an appropriate level. Finally, the optical transmitter is shown as a laser 150. The laser may be a directly modulated laser or an externally modulated laser in which the output is modulated using techniques such as electro-absorption modulation, Mach-Zender modulation or any appropriate technique for producing analog modulated optical signal. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other “smart” equipment to

1232

Claim 15	Disclosure
	determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:34–5:4.  FIG. 3A Piehler, Figure 3A.

1233

Claim 15	Disclosure
	 FIG. 4 Piehler, Figure 4. Linearized optical transmitter units are described for a hybrid optical fiber coaxial cable network. The linearized optical transmitter unit can comprise a directly-modulated or externally-modulated laser optically coupled to an optical conduit directed to an optical fiber communications link and electrically coupled to an electrical RF source line that provides an RF source to drive the laser or an external modulator for a light beam from the laser. A linearization information electrical component comprising memory and/or a processor, and a data output configured to transmit linearization enabling data for input into a direct digital synthesis engine that enables the direct digital synthesis engine to generate an RF

1234

Claim 15	Disclosure
	signal wherein nonlinear responses of the transmitter and/or the optical fiber communications link are pre-compensated, in which the data is specific for the optical transmitter and/or the optical fiber communications link. An electronic communication channel can connected to the data output of the electrical component to communicate the linearization enabling data external to the linearized optical transmitter unit suitable for communication to the direct digital synthesis engine. The linearized optical transmitter can be assembled in a plug-in module. Corresponding HFC systems and methods are also described. Piehler, Abstract. FIG. 3A shows edge QAM 202 and optical transmitter 204 connected by interface 206, e.g., a 25 ohm RF interface. While the division between edge-QAM 202 and optical transmitter 204 can be considered arbitrary, the division provides a convenient design framework since the edge-QAM 202 provides electrical RF signals based on user specific content and optical transmitter 204 provides laser module specific performance. The interface involves transfer of appropriate signals to provide the respective functionalities including, for example, pre-compensation for optical transmitter and fiber nonlinearities for efficient production of optical signals with low distortion. As described below, the division of the optical transmitter provides for the formation of separate modules, which can be plug-in components. A digital switch or other device passes a signal to edge QAM 202 to provide user specific digital content. Piehler, 9:38–53. In additional or alternative embodiments, electrical component 224 may comprise a microprocessor and further receive signal information that is used in creation of the pre-distortion parameters for dynamic correction of nonlinearities. As shown in FIG. 3A, electrical component 224 can receive optionally signals from sensor 228 and/or from an optical coupler or tap 230 and RF coupler or tap 232. Sensor 228 can measure, for example, laser temperature, laser bias voltage and/or other local conditions that can influence optical transmitter nonlinearities. Optical coupler 230 and RF coupler may be integrated into the system at a position downstream and upstream, respectively, of laser driver/laser assembly 222. One or both of the couplers/tap may be used to provide signal information, as indicated by dashed lines

1235

Claim 15	Disclosure
	connecting the couplers and microprocessor. If both couplers 230, 232 are used, the microprocessor may compare the upstream and downstream signal information to isolate distortion created by laser driver/laser assembly 222 and RF amplifier 220, and use the comparison data to create data related to the pre-distortion correction that is provided to processing unit 208. The correction may include predetermined parameters specific to laser modulator 222 or be based only on the comparison data. Electrical component 224 may be an integrated circuit or other microprocessor programmed for active processing and/or supply of data or a passive device, e.g., random access memory, non-volatile memory, or other memory. Information in electrical component 224 may comprise parameters for the specific laser module in the device. For example, the laser may be tested at the point of manufacturing or at other times to measure the distortion that it creates. The distortion has consistent components that are specific to the laser; these components may be analyzed and appropriate compensatory corrections created that are placed in, or otherwise introduced by, the electrical component. The pre-distortion correction provided to the DDS may comprise initial values for pre-distortion based on the laser's intrinsic nonlinearity at normal operating conditions as well as other nonlinearities which may exist in the system. The pre-distortion corrections may thus compensate for the net nonlinearity of the optical transmitter which includes, the laser modulator, other optical components, and other RF components including RF amplifiers. Examples of corrections are compensations for one or more of second, third, and higher nth order distortions in the optical transmitter. Piehler, 10:35–11:12. **[0038]** With reference to the drawings, FIG. **2** is a flow chart of the operation of an example embodiment, wherein the LA/CDR unit **144** in conjunction with the microcontroller

1236

Claim 15	Disclosure
	150 performs pre-equalization and optimizes power consumption of the transmission over host channel **103** via physical conductors **114** based on the feedback received from the SERDES module **110** via the low-speed digital interface **152**. In FIG. **2**, the operational flow of the LA/CDR unit **144** and the microcontroller **150** begins at step **202** when the microcontroller receives feedback via the low speed digital interface **152**. At step **204**, the microcontroller **150** instructs the LA/CDR unit **144** to equalize its transmitted data signal in response to the feedback, applying sufficient transmit pre-emphasis and other pre-equalization effects to the transmitted data signal to achieve predetermined target eye characteristics at the SERDES module **110**. Once the feedback from the SERDES module **110** indicates that the predetermined target eye characteristics have been achieved, the microcontroller **150** at step **206** instructs the LA/CDR unit **144** to begin reducing the launch amplitude of the transmitted data signal as much as it can without letting the received signal at the SERDES module **110** fall short of the predetermined target eye characteristics. The process is iterative, with feedback being monitored at step **202** and equalization being adjusted at step **204** while the launch amplitude is being reduced at step **206**. In alternative embodiments, the process may operate according to a different sequence of steps while achieving the same result, namely the optimization of signal quality and power consumption over the host channel **103**. Tremblay, paragraph 38.

1237

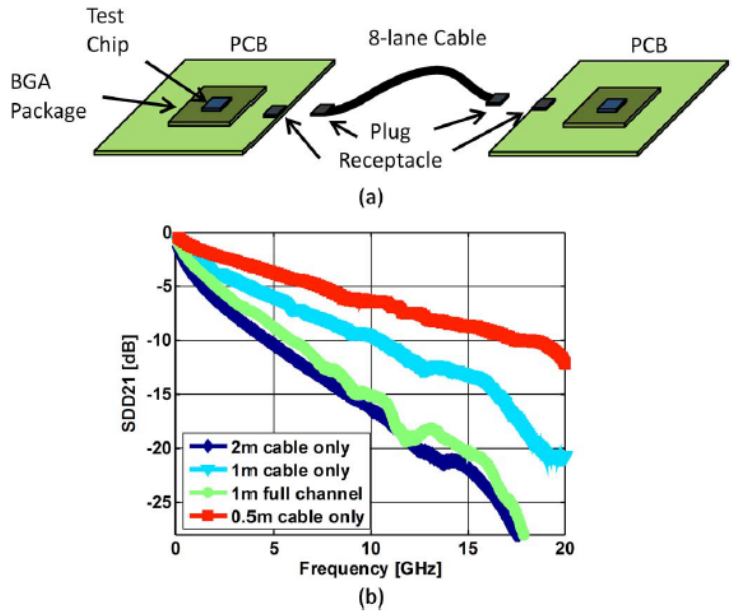
P. DEPENDENT CLAIM 16

Claim 16	Disclosure
16[a] The method of claim 15, further comprising: supplying power to the first and second end connector plugs,	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

1238

Claim 16	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.

1239

Claim 16	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

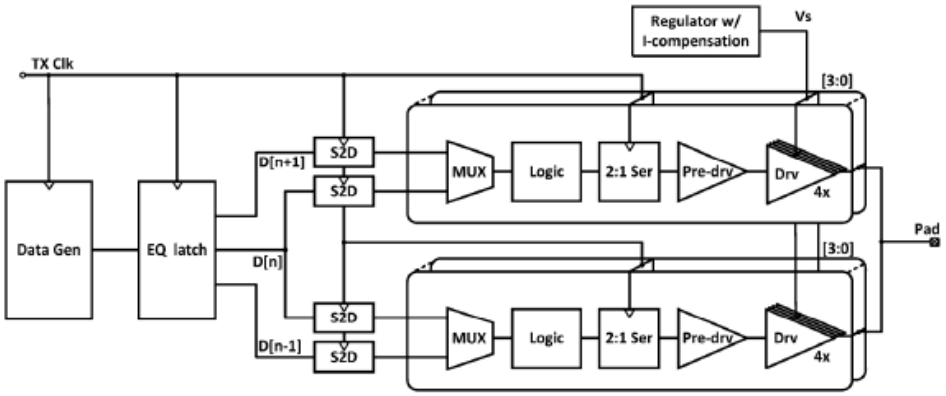
1240

Claim 16	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.

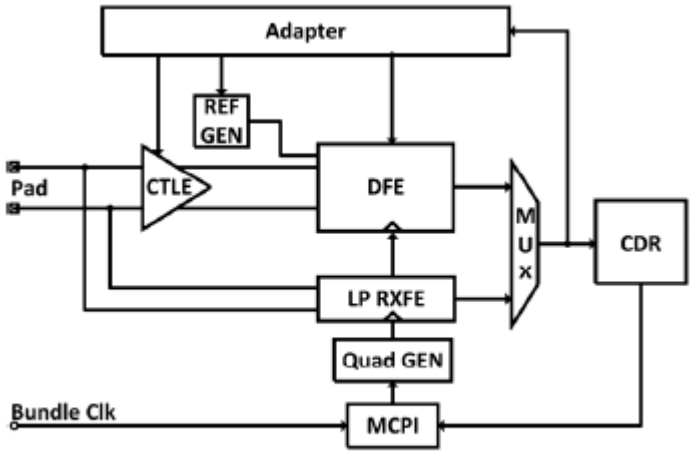
1241

Claim 16	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.

1242

Claim 16	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

1243

Claim 16	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.

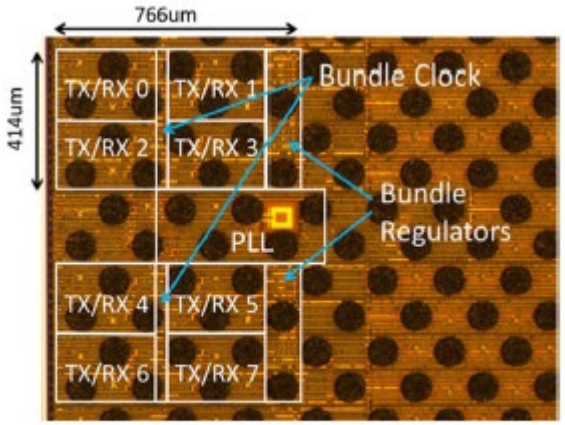
1244

Claim 16	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.

1245

Claim 16	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.

1246

Claim 16	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.

1247

Claim 16	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. 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

1248

Claim 16	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.

1249

Claim 16	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.

1250

Claim 16	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.

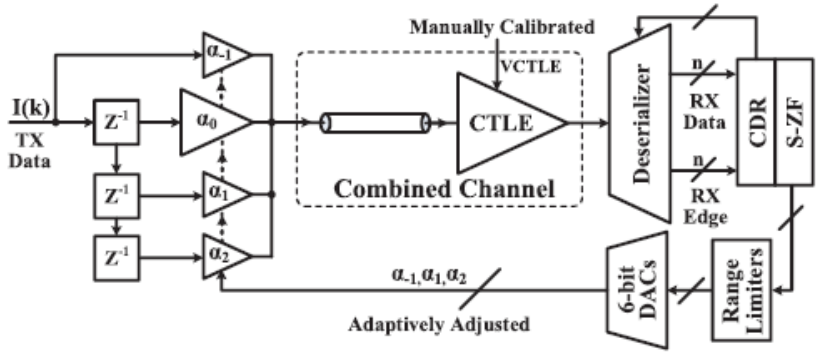


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1251

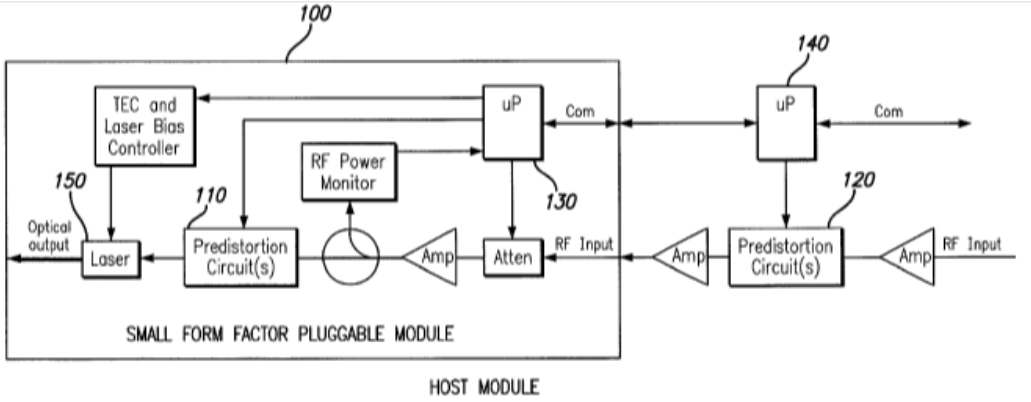
Claim 16	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

1252

Claim 16	Disclosure
1253	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. 7,233,617 to Gorecki* Like the '233 Patent's specification, Gorecki discloses storing of coefficient values in memory as programming. *See* Gorecki at 7:1-33 ("the adjustment or control of the taps of the equalization circuitry may be by the user via programming… coefficients of the tap(s), and/or the pulse durations of the tap(s) may be pre-programmed"); *see also*'233 Patent at 11:3-6 ("the automated tester equipment 'burns' the flash memory or otherwise programs the initial default values of the filter coefficients[.]"). Gorecki teaches that the programmed memory "may reside on the integrated circuit containing the transmitter, receiver or transceiver." *See* Gorecki at 7:43-49.

Claim 16	Disclosure
	This programming is done to use the transmit coefficient values when power is supplied to the first and second end connector plugs. Gorecki discloses that stored filter coefficients are retrieved from the non-volatile memory at start-up/power-up. *See*, *e.g.*, Gorecki at 7:18-33; *see also* claim 2. Power-up indicates that power is supplied to the first and second end connector plugs. An ordinarily skilled artisan would have understood that power is supplied from a host device when the enclosing connector's connector plug was electrically connected to the host interface port in a corresponding electronic device 101a/b as described claim limitation 1[a] because this was conventionally how standardized connectors like QSFP or SFP worked . *See, e.g.*, SFF-8436 at 21 ("A host board together with the QSFP+ module(s) forms an integrated power system. The host supplies stable power to the module."), 22 Fig. 4 ("Recommended Host Board Power Supply Filtering"). Regardless, an ordinarily skilled artisan would have found it conventional and obvious to implement the QSFP connectors powered by a connected host device. An ordinarily skilled artisan would have understood that Gorecki uses the transmit coefficient values stored in memory each time power is supplied to the first and second end connector plugs, and teaches that the programmed values can be stored "permanently, semi-permanently or temporarily[.]" *See* Gorecki at 7:21. When stored permanently, an ordinarily skilled artisan would have understood them to be used each time power is supplied to the first and second end connector plugs.

1254

Claim 16	Disclosure
	 FIG. 1 Ishaug, Figure 1.

1255

Claim 16	Disclosure
	In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4.
16[b] the first and second end connector plugs being connected to first and second controller devices, respectively,	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:

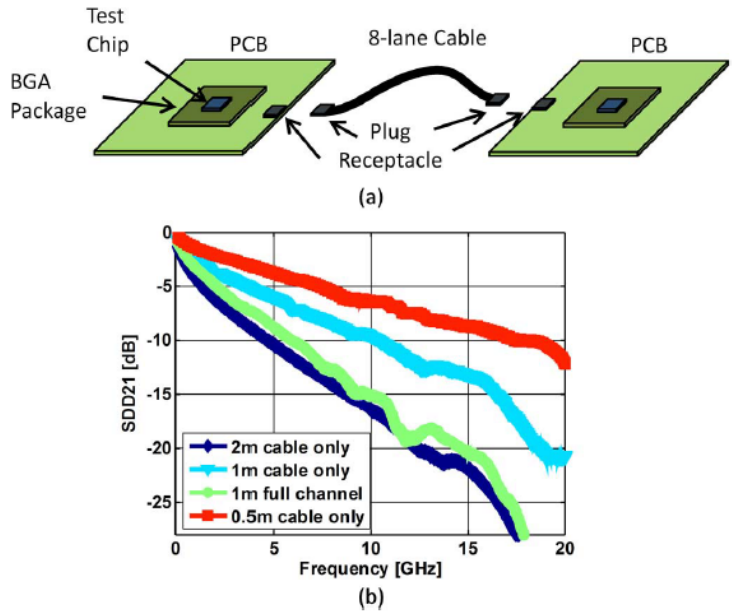
1256

Claim 16	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

1257

Claim 16	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.

1258

Claim 16	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

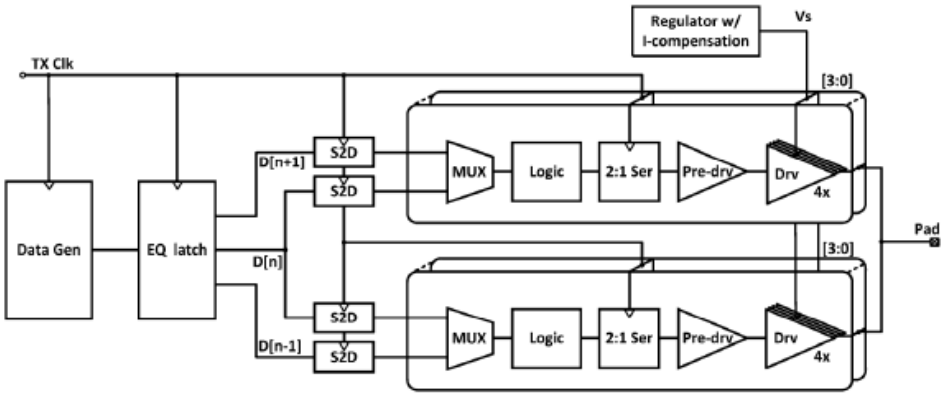
1259

Claim 16	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.

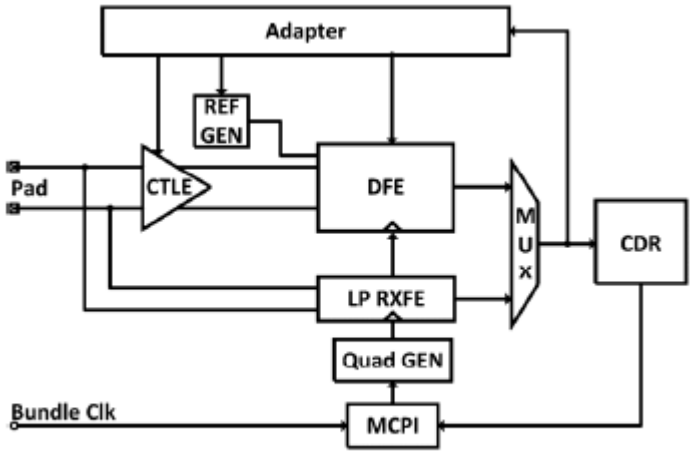
1260

Claim 16	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.

1261

Claim 16	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

1262

Claim 16	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.

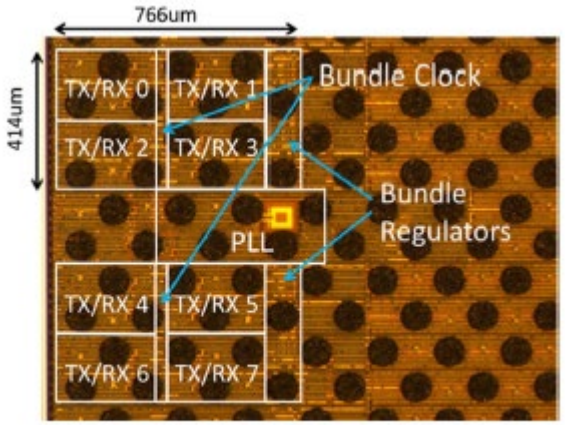
1263

Claim 16	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.

1264

Claim 16	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.

1265

Claim 16	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.

1266

Claim 16	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. 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

1267

Claim 16	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.

1268

Claim 16	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.

1269

Claim 16	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.

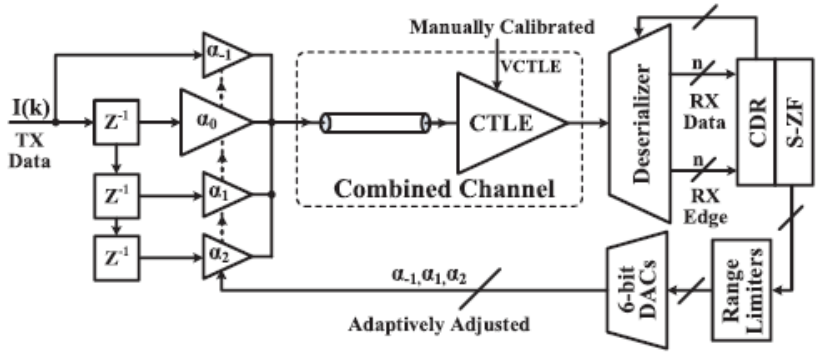


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

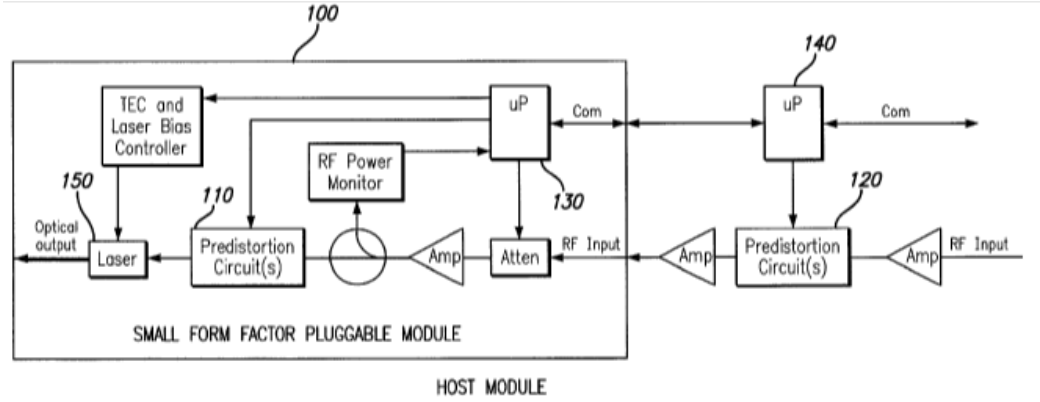
1270

Claim 16	Disclosure
1271	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 16	Disclosure
1272	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. 7,233,617 to Gorecki* Like the '233 Patent's specification, Gorecki discloses storing of coefficient values in memory as programming. *See* Gorecki at 7:1-33 ("the adjustment or control of the taps of the equalization circuitry may be by the user via programming… coefficients of the tap(s), and/or the pulse durations of the tap(s) may be pre-programmed"); *see also*'233 Patent at 11:3-6 ("the automated tester equipment 'burns' the flash memory or otherwise programs the initial default values of the filter coefficients[.]"). Gorecki teaches that the programmed memory "may reside on the integrated circuit containing the transmitter, receiver or transceiver." *See* Gorecki at 7:43-49.

Claim 16	Disclosure
	This programming is done to use the transmit coefficient values when power is supplied to the first and second end connector plugs. Gorecki discloses that stored filter coefficients are retrieved from the non-volatile memory at start-up/power-up. *See*, *e.g.*, Gorecki at 7:18-33; *see also* claim 2. Power-up indicates that power is supplied to the first and second end connector plugs. An ordinarily skilled artisan would have understood that power is supplied from a host device when the enclosing connector's connector plug was electrically connected to the host interface port in a corresponding electronic device 101a/b as described claim limitation 1[a] because this was conventionally how standardized connectors like QSFP or SFP worked . *See, e.g.*, SFF-8436 at 21 ("A host board together with the QSFP+ module(s) forms an integrated power system. The host supplies stable power to the module."), 22 Fig. 4 ("Recommended Host Board Power Supply Filtering"). Regardless, an ordinarily skilled artisan would have found it conventional and obvious to implement the QSFP connectors powered by a connected host device. An ordinarily skilled artisan would have understood that Gorecki uses the transmit coefficient values stored in memory each time power is supplied to the first and second end connector plugs, and teaches that the programmed values can be stored "permanently, semi-permanently or temporarily[.]" *See* Gorecki at 7:21. When stored permanently, an ordinarily skilled artisan would have understood them to be used each time power is supplied to the first and second end connector plugs.

1273

Claim 16	Disclosure
	 FIG. 1 Ishaug, Figure 1.

1274

Claim 16	Disclosure
	In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4.
16[c] each of the first and second controller devices operating to configure the first and second DRR devices in response to a power-on event,	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.

1275

Claim 16	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.

1276

Claim 16	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.

1277

Claim 16	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.

1278

Claim 16	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.

1279

Claim 16	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.

1280

Claim 16	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.

1281

Claim 16	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.

1282

Claim 16	Disclosure

1283

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	

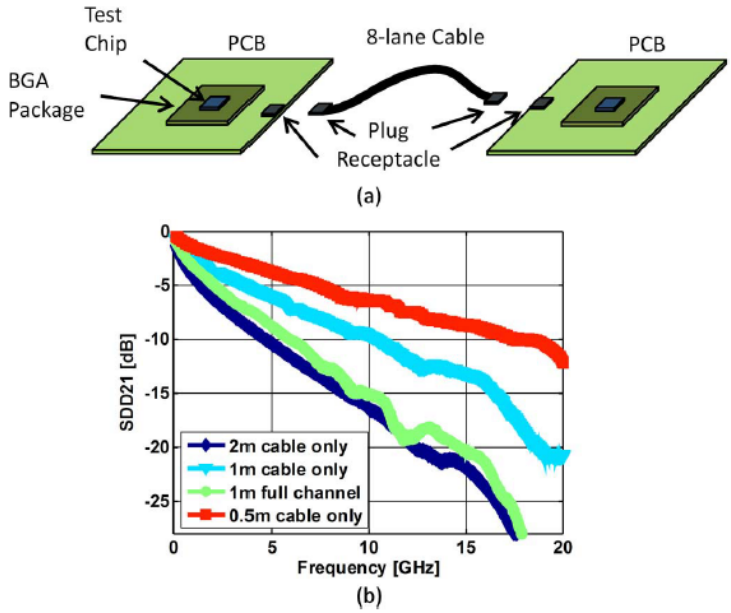
CMIS, 71.

Claim 16	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

1284

Claim 16	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.

1285

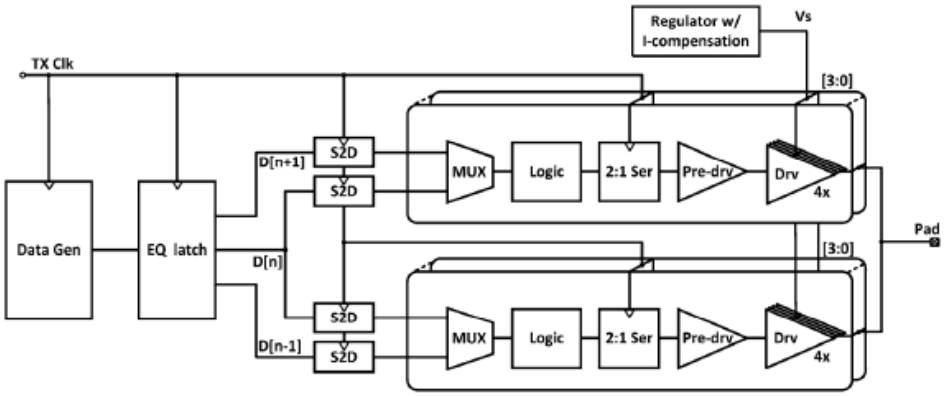
Claim 16	Disclosure
1286	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

Claim 16	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.

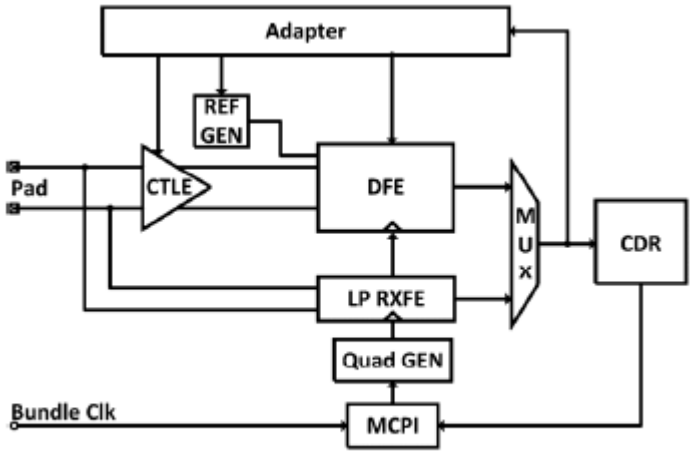
1287

Claim 16	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.

1288

Claim 16	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

1289

Claim 16	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.

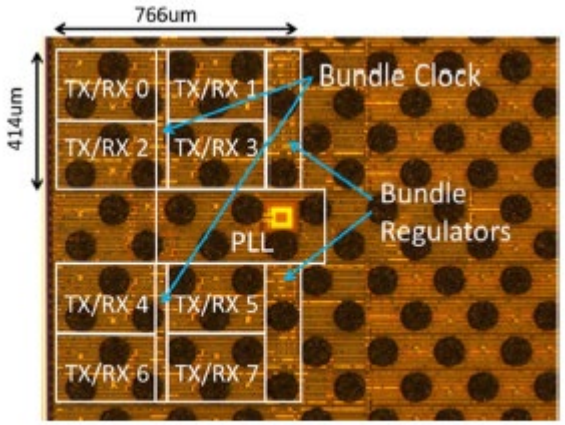
1290

Claim 16	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.

1291

Claim 16	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.

1292

Claim 16	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.

1293

Claim 16	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. 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

1294

Claim 16	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.

1295

Claim 16	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.

1296

Claim 16	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.

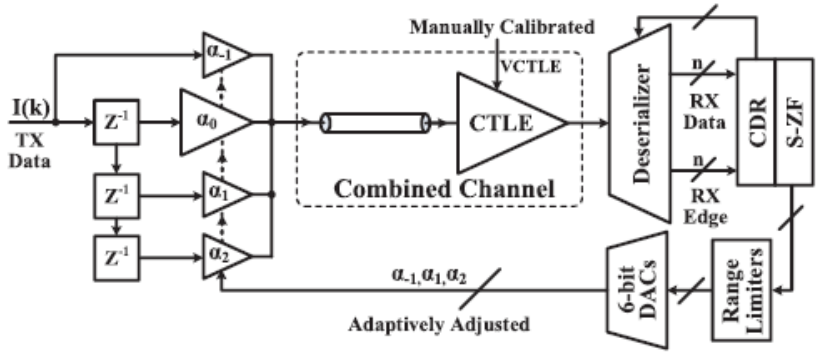


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1297

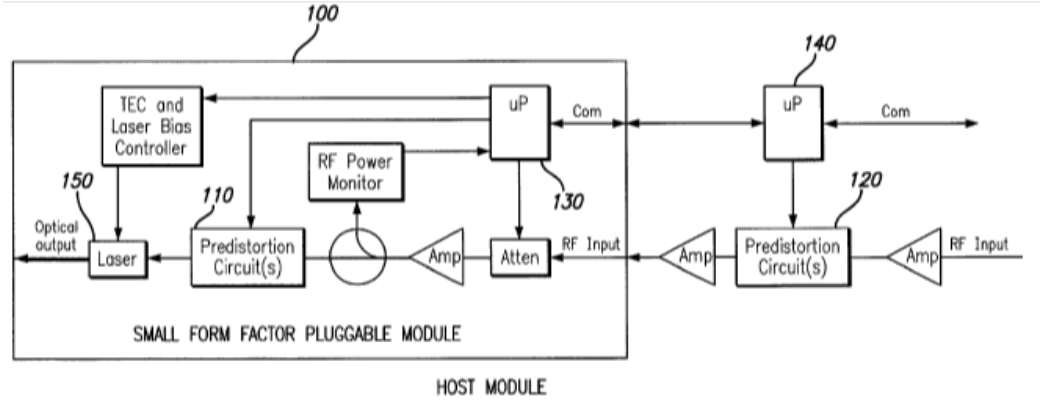
Claim 16	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

1298

Claim 16	Disclosure
1299	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. 7,233,617 to Gorecki* Like the '233 Patent's specification, Gorecki discloses storing of coefficient values in memory as programming. *See* Gorecki at 7:1-33 ("the adjustment or control of the taps of the equalization circuitry may be by the user via programming… coefficients of the tap(s), and/or the pulse durations of the tap(s) may be pre-programmed"); *see also*'233 Patent at 11:3-6 ("the automated tester equipment 'burns' the flash memory or otherwise programs the initial default values of the filter coefficients[.]"). Gorecki teaches that the programmed memory "may reside on the integrated circuit containing the transmitter, receiver or transceiver." *See* Gorecki at 7:43-49.

Claim 16	Disclosure
	This programming is done to use the transmit coefficient values when power is supplied to the first and second end connector plugs. Gorecki discloses that stored filter coefficients are retrieved from the non-volatile memory at start-up/power-up. *See*, *e.g.*, Gorecki at 7:18-33; *see also* claim 2. Power-up indicates that power is supplied to the first and second end connector plugs. An ordinarily skilled artisan would have understood that power is supplied from a host device when the enclosing connector's connector plug was electrically connected to the host interface port in a corresponding electronic device 101a/b as described claim limitation 1[a] because this was conventionally how standardized connectors like QSFP or SFP worked . *See, e.g.*, SFF-8436 at 21 ("A host board together with the QSFP+ module(s) forms an integrated power system. The host supplies stable power to the module."), 22 Fig. 4 ("Recommended Host Board Power Supply Filtering"). Regardless, an ordinarily skilled artisan would have found it conventional and obvious to implement the QSFP connectors powered by a connected host device. An ordinarily skilled artisan would have understood that Gorecki uses the transmit coefficient values stored in memory each time power is supplied to the first and second end connector plugs, and teaches that the programmed values can be stored "permanently, semi-permanently or temporarily[.]" *See* Gorecki at 7:21. When stored permanently, an ordinarily skilled artisan would have understood them to be used each time power is supplied to the first and second end connector plugs.

1300

Claim 16	Disclosure
1301	 FIG. 1 Ishaug, Figure 1.

Claim 16	Disclosure
	In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4.
16[d] the configuring including retrieving the transmit filter coefficient values from internal nonvolatile memory.	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.

1302

Claim 16	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.

1303

Claim 16	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.

1304

Claim 16	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.

1305

Claim 16	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.

1306

Claim 16	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.

1307

Claim 16	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.

1308

Claim 16	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.

1309

Claim 16	Disclosure

1310

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	

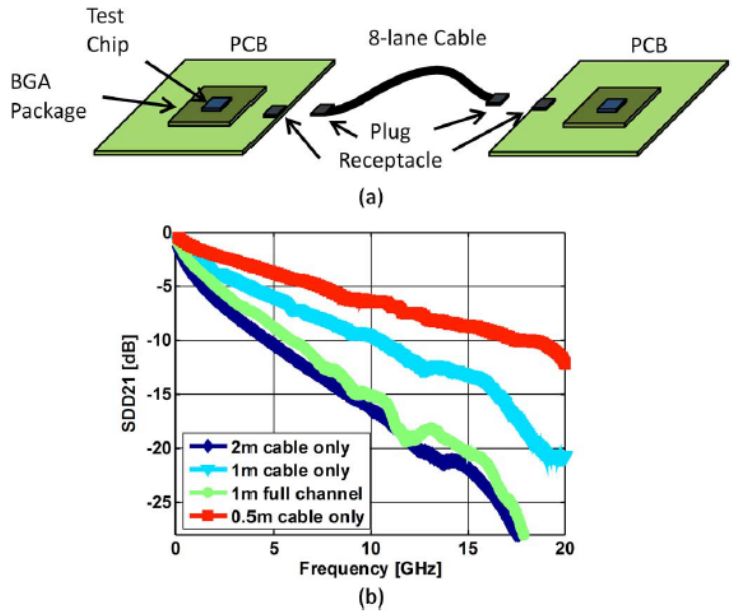
CMIS, 71.

Claim 16	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

1311

Claim 16	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.

1312

Claim 16	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

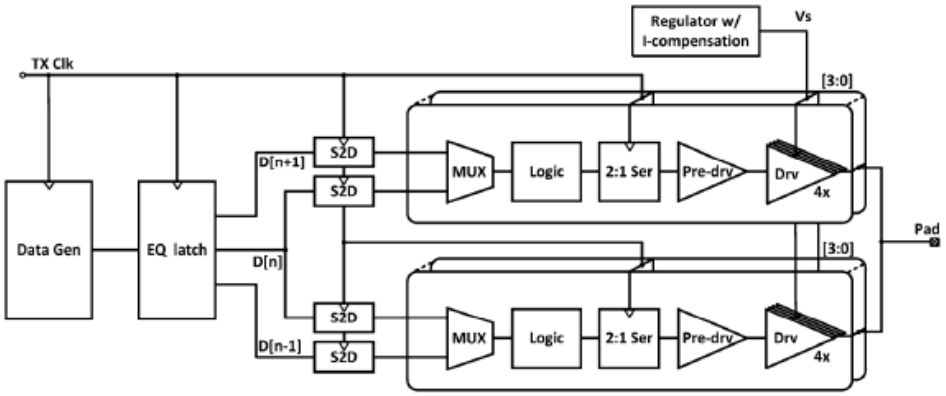
1313

Claim 16	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.

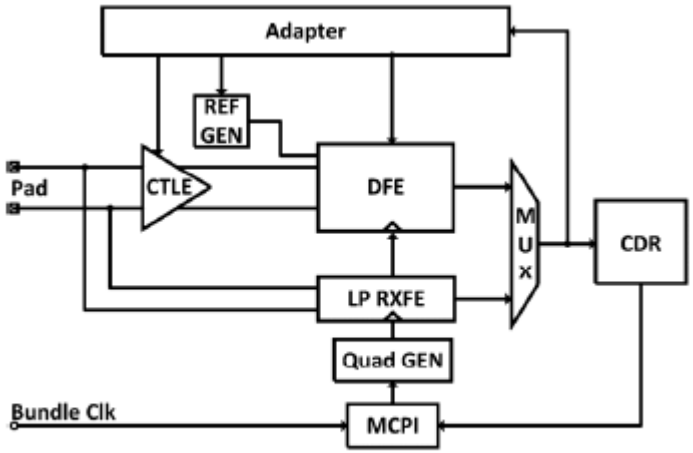
1314

Claim 16	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.

1315

Claim 16	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

1316

Claim 16	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.

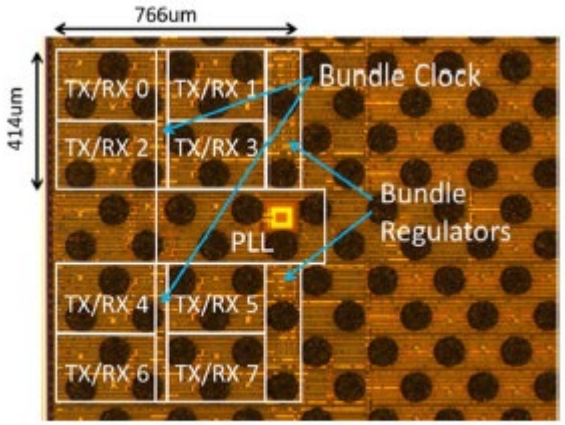
1317

Claim 16	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.

1318

Claim 16	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.

1319

Claim 16	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.

1320

Claim 16	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. 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

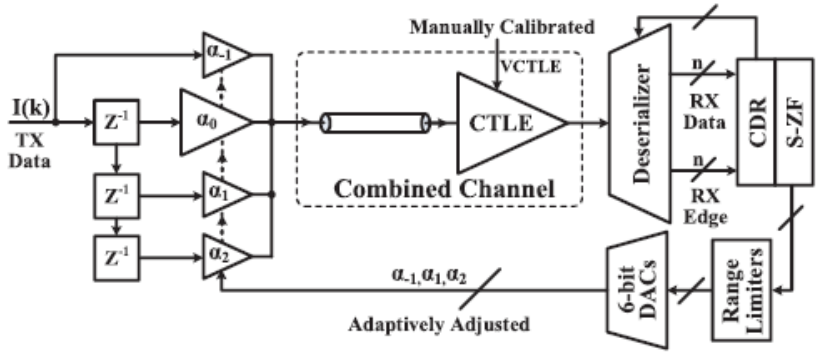
1321

Claim 16	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.

1322

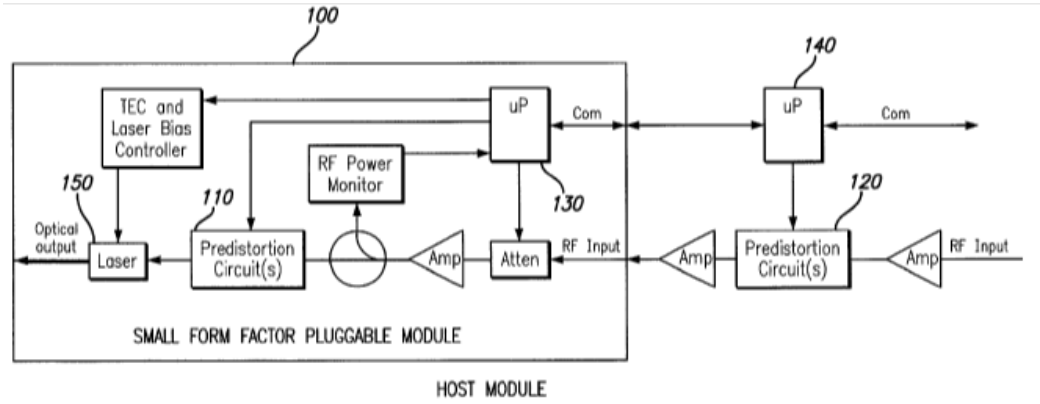
Claim 16	Disclosure
	III. RECEIVER CHIP *A. Overall Architecture* The main task of the RX is to extract the transmitted data from the received signal using appropriate equalization and CDR techniques [26]–[29]. Fig. 8 shows the block diagram of the RX chip. It consists of a two-stage continuous-time 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.

1323

Claim 16	Disclosure
1324	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 16	Disclosure
1325	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 16	Disclosure
1326	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. 7,233,617 to Gorecki* Gorecki teaches "[I]nformation representative of… the coefficients of the tap(s)… may be stored… in… for example… ROM or EEPROM . In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up," which is a power-on event. *See* Gorecki at 7:18-33, 16:47-53. In addition, Gorecki describes a "controller" that "distribute[s] to the transmitter" the filter coefficients thereby configuring a filter for pre-emphasis at this power-on event. *See* Gorecki at 7:4-17; 7:26-33. Gorecki describes a "controller" that "distribute[s] to the transmitter" the filter coefficients thereby configuring a filter for pre-emphasis at this power-on event. *See* Gorecki at 7:4-17; 7:26-33.

Claim 16	Disclosure
	 FIG. 1 Ishaug, Figure 1.

Claim 16	Disclosure
	In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4.

1328

Q. DEPENDENT CLAIM 17

Claim 17	Disclosure
17. The method of claim 15, wherein the first and second DRR devices are programmed to use the transmit coefficient values each time power is supplied to the first and second end connector plugs.	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.

1329

Claim 17	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.

1330

Claim 17	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.

1331

Claim 17	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.

1332

Claim 17	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.

1333

Claim 17	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.

1334

Claim 17	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.

1335

Claim 17	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.

1336

Claim 17	Disclosure
1337	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.

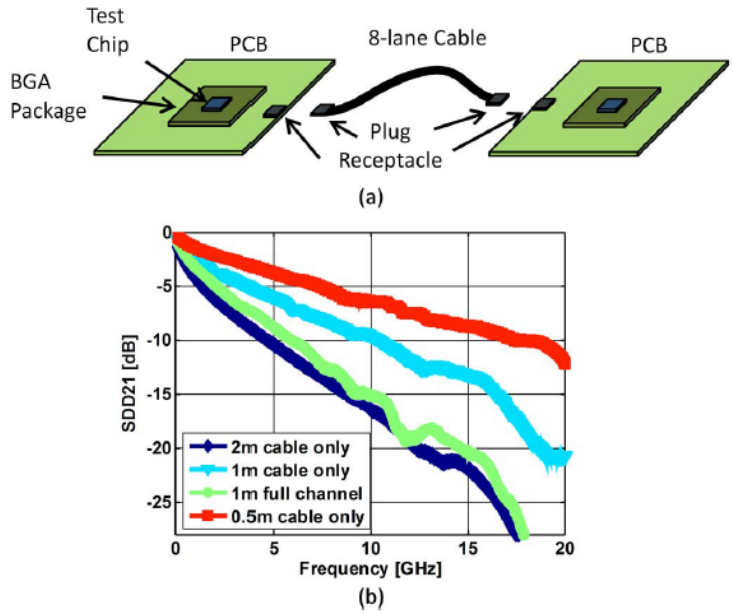
Byte	Bit	Name	Description	Type
161	7	Reserved		RO RQD
	6-5	Tx Input Eq Store/Recall buffer count	00b=Tx Input Eq Store/Recall not implemented 01b=Tx Input Eq Store/Recall buffer count=1 10b=Tx Input Eq Store/Recall buffer count=2 11b=reserved	RO RQD
	4	Tx Input Eq Freeze implemented	0b=Tx Input Eq Freeze 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 17	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

1338

Claim 17	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.

1339

Claim 17	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

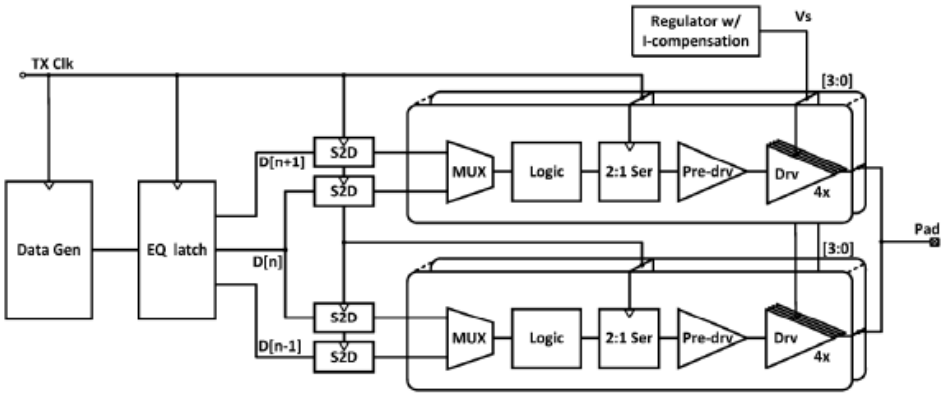
1340

Claim 17	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.

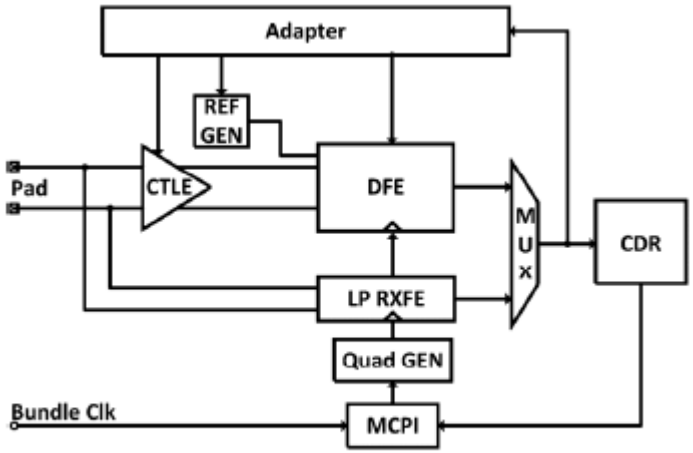
1341

Claim 17	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.

1342

Claim 17	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

1343

Claim 17	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.

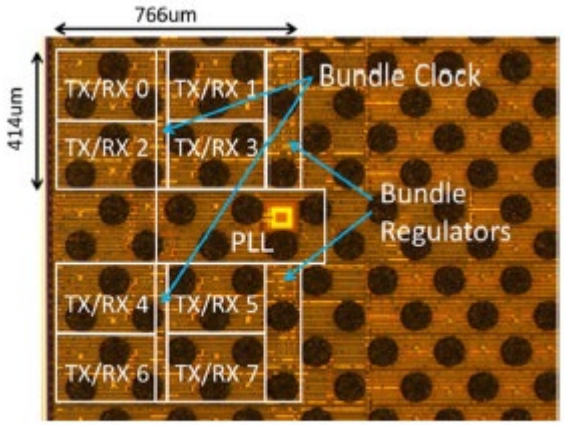
1344

Claim 17	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.

1345

Claim 17	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.

1346

Claim 17	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.

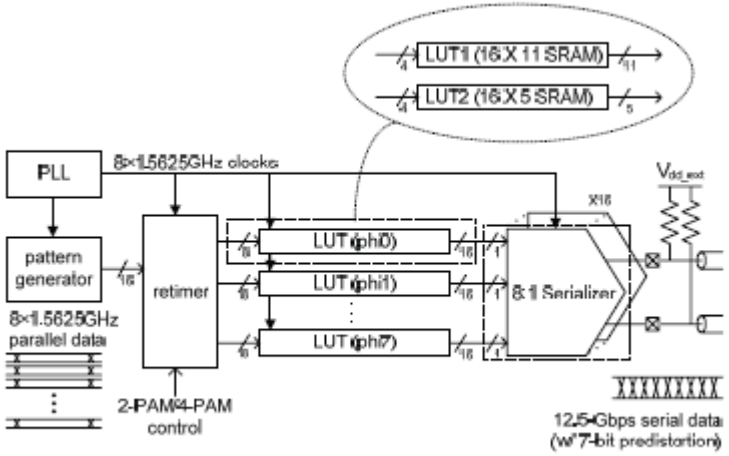
1347

Claim 17	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.

1348

Claim 17	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.

1349

Claim 17	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1350

Claim 17	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

1351

Claim 17	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.

1352

Claim 17	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

1353

Claim 17	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.

1354

Claim 17	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.

1355

Claim 17	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.

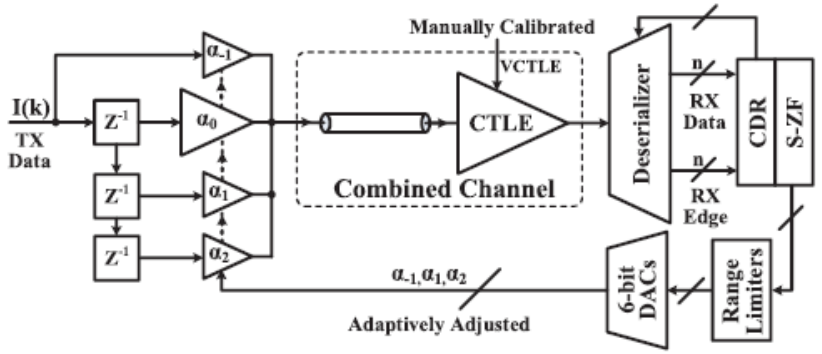


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1356

Claim 17	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.

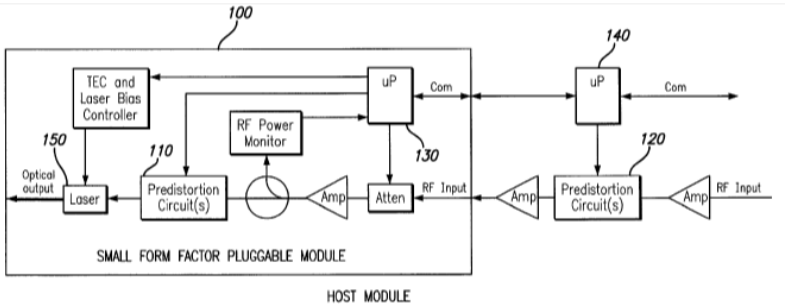
1357

Claim 17	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. 7,233,617 to Gorecki* For the limited purpose of this analysis, the phrase "transmit coefficient values" is interpreted as referring to "transmit filter coefficient values," (*see, e.g.*, claim limitation 1[f]) as the term otherwise has no antecedent basis. Like the '233 Patent's specification, Gorecki discloses storing of coefficient values in memory as programming. *See* Gorecki at 7:1-33 ("the adjustment or control of the taps of the equalization circuitry may be by the user via programming… coefficients of the tap(s), and/or the pulse durations of the tap(s) may be pre-programmed"); *see also*'233 Patent at 11:3-6 ("the automated tester equipment 'burns' the flash memory or otherwise programs the initial default values of the filter coefficients[.]").

1358

Claim 17	Disclosure
	Gorecki teaches that the programmed memory "may reside on the integrated circuit containing the transmitter, receiver or transceiver." *See* Gorecki at 7:43-49. This programming is done to use the transmit coefficient values when power is supplied to the first and second end connector plugs. Gorecki discloses that stored filter coefficients are retrieved from the non-volatile memory at start-up/power-up. *See*, *e.g.*, Gorecki at 7:18-33; *see also* claim 2. Power-up indicates that power is supplied to the first and second end connector plugs. An ordinarily skilled artisan would have understood that power is supplied from a host device when the enclosing connector's connector plug was electrically connected to the host interface port in a corresponding electronic device 101a/b as described claim limitation 1[a] because this was conventionally how standardized connectors like QSFP or SFP worked . *See, e.g.*, SFF-8436 at 21 ("A host board together with the QSFP+ module(s) forms an integrated power system. The host supplies stable power to the module."), 22 Fig. 4 ("Recommended Host Board Power Supply Filtering"). Regardless, an ordinarily skilled artisan would have found it conventional and obvious to implement the QSFP connectors powered by a connected host device. An ordinarily skilled artisan would have understood that Gorecki uses the transmit coefficient values stored in memory each time power is supplied to the first and second end connector plugs, and teaches that the programmed values can be stored "permanently, semi-permanently or temporarily[.]" *See* Gorecki at 7:21. When stored permanently, an ordinarily skilled artisan would have understood them to be used each time power is supplied to the first and second end connector plugs.

1359

Claim 17	Disclosure
1360	 FIG. 1 Ishaug, Figure 1. In a preferred embodiment, there is no predistortion circuitry in the pluggable module and the host module predistortion settings required for optimum performance are stored in the non-volatile memory of the pluggable optical module. These predistortion setting may be unique for each individual module or for each type of module and may be set by the pluggable module manufacturer using appropriate techniques to determine the optimum predistortion settings. At power up, the host module predistortion settings stored in the pluggable optical module are read by the host module and implemented by the predistortion

Claim 17	Disclosure
	circuit in the host module, typically by providing appropriate bias voltages or currents to the predistortion circuit. The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:51–5:4. **[0033]** In example embodiments a low speed digital connection or interface **152** is provided between host SERDES **110** and the interface module **100** to allow back channel communications between transmitting elements of the SERDES module **110** and corresponding receiving elements of the interface module **100** and between transmitting elements of the interface module **100** and receiving components of the SERDES Module **110**. In some implementations, the interface module **100** includes a uC (microcontroller) module **150** that facilitates communications over the low speed digital interface **152**. By way of non-limiting example, the low speed digital interface **152** could include an I2C ("Inter-Integrated Circuit") interface, or SPI ("Serial Peripheral Interface"), MDIO ("Management Data Input/Output") interface, or similar interface. In some example embodiments, the low speed data connection **152** is used for data rates under 10 MB/s. Tremblay, paragraph 33.

1361

R. DEPENDENT CLAIM 18

Claim 18	Disclosure
18. The method of claim 17, wherein the transmit filter coefficient values are determined and stored in the nonvolatile memories after assembly of the 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.

1362

Claim 18	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.

1363

Claim 18	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.

1364

Claim 18	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.

1365

Claim 18	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.

1366

Claim 18	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.

1367

Claim 18	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.

1368

Claim 18	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.

1369

Claim 18	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.

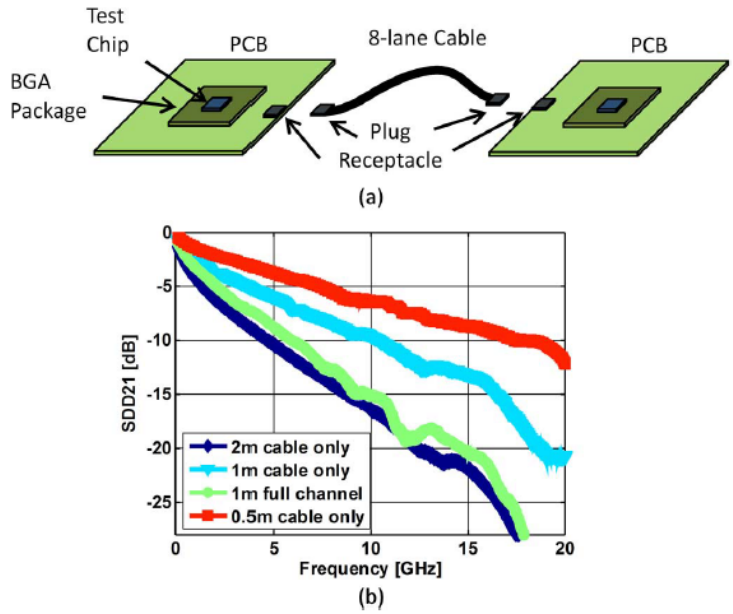
1370

Claim 18	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

1371

Claim 18	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.

1372

Claim 18	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

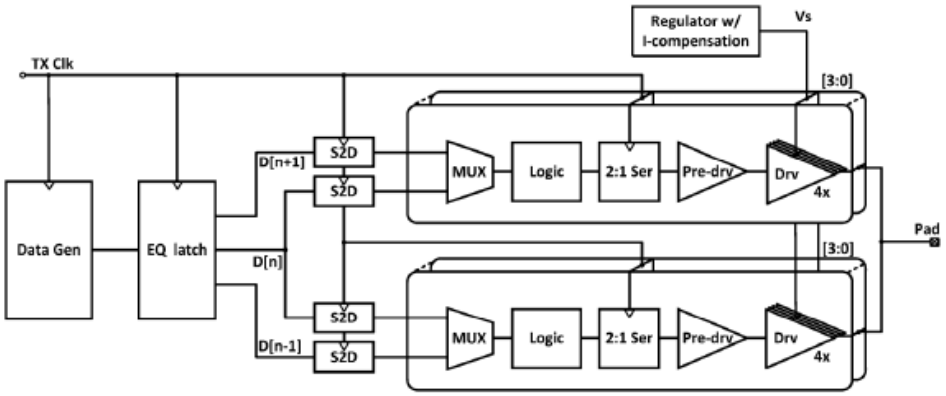
1373

Claim 18	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.

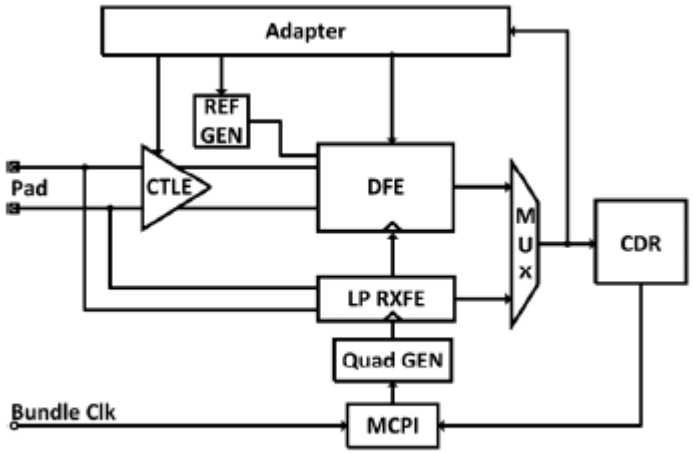
1374

Claim 18	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.

1375

Claim 18	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

1376

Claim 18	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.

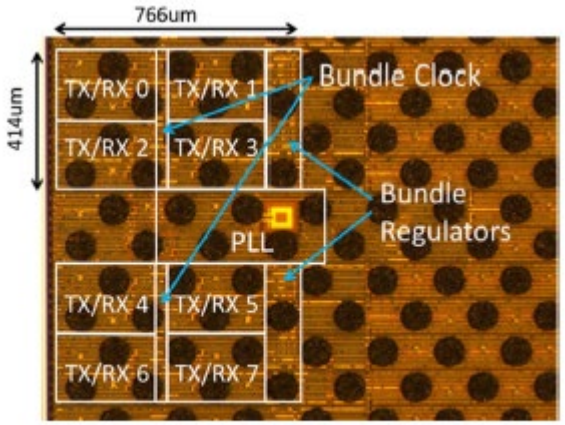
1377

Claim 18	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.

1378

Claim 18	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.

1379

Claim 18	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.

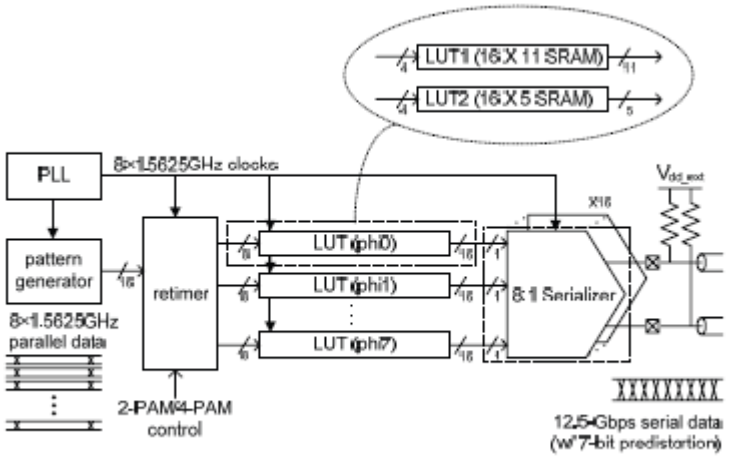
1380

Claim 18	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.

1381

Claim 18	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.

1382

Claim 18	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1383

Claim 18	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

1384

Claim 18	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.

1385

Claim 18	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

1386

Claim 18	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.

1387

Claim 18	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.

1388

Claim 18	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.

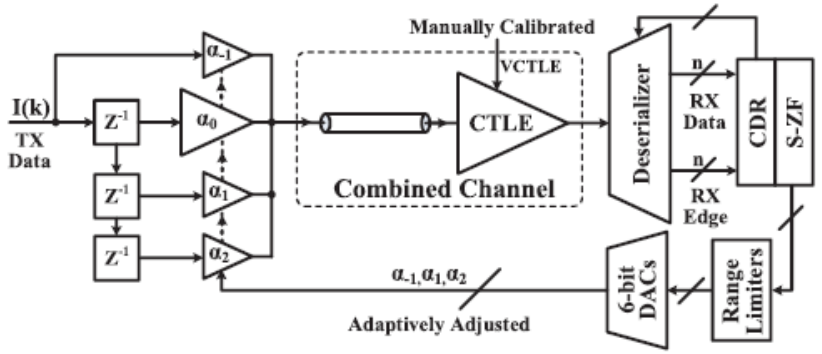


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1389

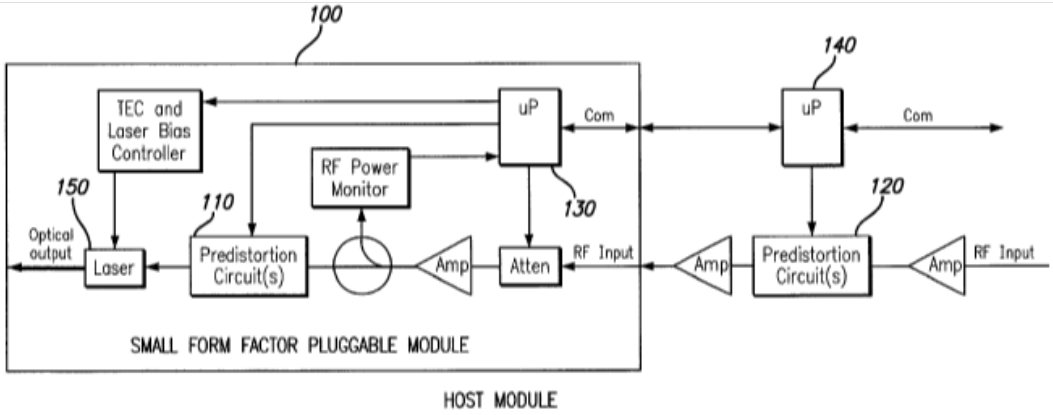
Claim 18	Disclosure
	Zheng, 2970. *D. Implementation of the Edge-Data Correlation-Based S-ZF* Fig. 19 shows the implementation of our S-ZF adaptation algorithm, which contains three identical paths to process the quantized data/edge sequences to produce the desired bias voltages for TX-FFE taps. Here, 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.

1390

Claim 18	Disclosure
1391	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. Gorecki discloses determining and then storing the transmit filter coefficient values in the non-volatile memories after assembly of the cable. The values should be selected to offset effects of a communication channel. *See* Gorecki at 9:8-10 (The equalizers are "configured to introduce compensation that is ideally the inverse of the effects caused by the communications channel."). An ordinarily skilled artisan would have found it obvious to determine and then store filter coefficients, such as by measuring the effects in that communication channel after assembly. Such an approach was known in the art . *See, e.g.*, U.S. Patent App. Pub. No. 2014/0119425 at [0033] ("The programming is preferably performed with the cable installed between the actual source and actual sink it will be connecting."). Gorecki describes "the coefficients of the tap(s)… are adjusted or controlled during operation of the transmitter, for example using an adaptive algorithm." *See* Gorecki at 10:27-51, 6:50-53, 17:41-51. Determining values "during operation" happens after cable assembly. The adaptive algorithm—which

Claim 18	Disclosure
	determines transmit filter coefficient values—is performed "after (or during) the performance of an initialization or re-initialization process." *See* Gorecki at 7:50-61. Gorecki discloses that stored coefficient values can be "fine-tuned to enhance the system performance" (Gorecki at 7:51-53), and that the values can be "re-programmed" and, for that reason, are stored in writable non-volatile memory. *See, e.g.*, Gorecki at 7:18-33, 10:37-51. An ordinarily skilled artisan would have had reasons to use NVM including writable NVM ("EEPROM or the like"). *See* Gorecki at 7:24-25, 10:46-47. EEPROM is electrically erasable programmable read only memory, e.g., rewritable NVM. An ordinarily skilled artisan would have known that flash memory was another well-known rewritable non-volatile memory (e.g., "EEPROM or the like") conventionally used with transmit FIR filters. *See, e.g.*, U.S. Patent No. 6,975,140 at 3:42-4:2 ("Any type of NVRAM can be used as the rewritable non-volatile storage, so long as [it] is of the type suitable for integration on an integrated circuit."). An ordinarily skilled artisan would have understood such memories would be used when the content of those memories is both re-programmed and preserved through "startup/power-up, initialization or re-initialization," as described in Gorecki at 7:18-33. Goreki discloses rewritable non-volatile memory that could be updated with new coefficient values that fine-tuned pre-equalization or equalization with new coefficient values learned through training at manufacture or during usage with a particular device in a particular environment . *See, e.g.*, Gorecki '716 at 7:50-61, 10:37-51; *see also* U.S. Patent No. 6,975,140 at 3:42-4:2, 9:21-28. Preserving such adapted pre-equalization filter coefficients in rewritable non-volatile memory like an EEPROM would have ensured that the fine-tuned coefficients remained available after an AEC transceiver—or the host powering it—was power cycled, initialized, or re-initialized. *See* Gorecki at 7:31-33 ("In this way, the transmitter may access the memory to retrieve the necessary information during start up/power-up, initialization or re-initialization."). Regardless, an ordinarily skilled artisan would have known that storing updated filter coefficients in the L-G combination's non-volatile memory was conventional and desirable so that the adapted coefficient values would be available after power cycling the active electric cable transceiver by unplugging it from a host device powering it, or power cycling the host device powering the active element cable transceiver. *See* U.S. Patent No. 6,975,140 at Abstract ("[D]ata transmitter includes a rewriteable non-volatile storage, operable to be rewritten with control information representing the values of coefficients updated during

Claim 18	Disclosure
	operation off the FIR driver."), 2:46-49 (same), 3:45-52 ("Updated values of the coefficients used in the taps of the FIR transmitter 30 are determined and stored to the flash memory 20 during the operation of the FIR transmitter . When the FIR transmitter is powered off and back on again, the coefficient values as last updated are available to be retrieved from the flash memory 20 and applied to the taps of the FIR transmitter again.").

Claim 18	Disclosure
	 FIG. 1 Ishaug, Figure 1.

Claim 18	Disclosure
	The host module may also use input from a user interface and/or other "smart" equipment to determine, in conjunction with the parameters stored in the pluggable module or elsewhere, the optimum predistortion settings to achieve lowest noise and distortion at the receiver. Some examples of additional user input or parameters that may come from other "smart" equipment include, but are not limited to, fiber link length, transmitter optical mux dispersion and receiver optical mux dispersion. Ishaug, 4:63–5:4.

S. DEPENDENT CLAIM 19

Claim 19	Disclosure
19. The method of claim 18, wherein the first and second DRR devices employ receiver-based equalization of the electrical transit signals using coefficient values stored in the nonvolatile memories.	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.

1396

Claim 19	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.

1397

Claim 19	Disclosure
1398	**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 19	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.

1399

Claim 19	Disclosure
1400	**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 19	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.

1401

Claim 19	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.

1402

Claim 19	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.

1403

Claim 19	Disclosure

1404

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	

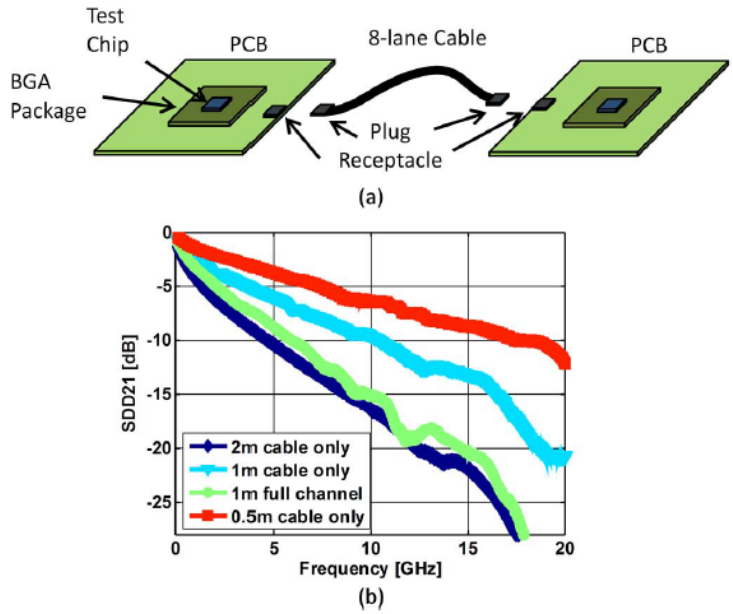
CMIS, 71.

Claim 19	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

1405

Claim 19	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.

1406

Claim 19	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

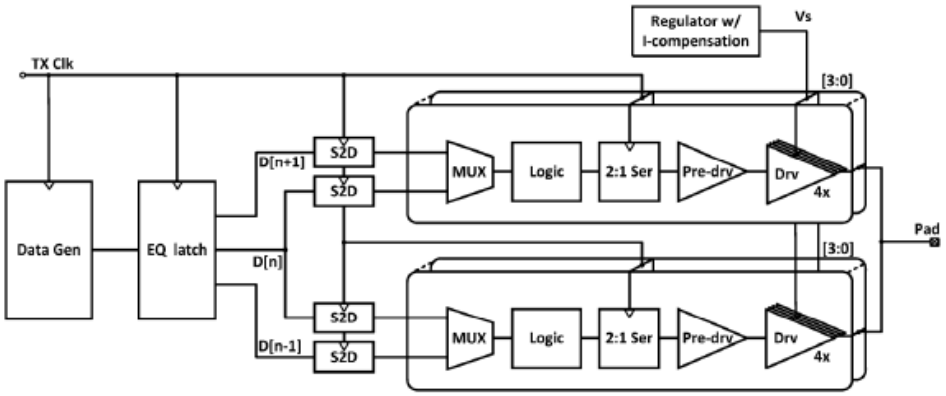
1407

Claim 19	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.

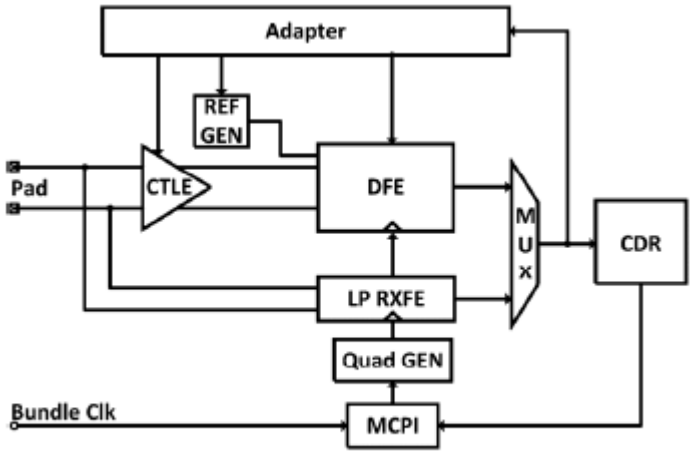
1408

Claim 19	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.

1409

Claim 19	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

1410

Claim 19	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.

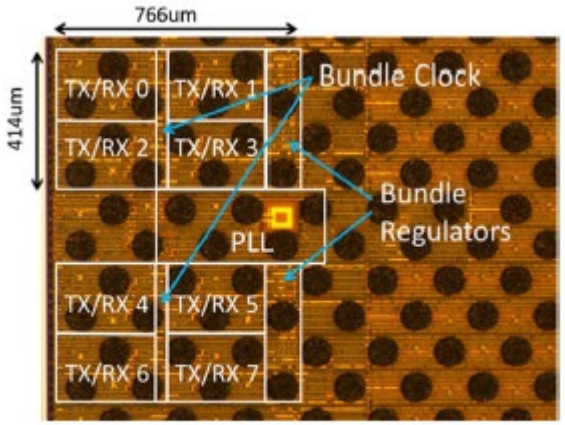
1411

Claim 19	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.

1412

Claim 19	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.

1413

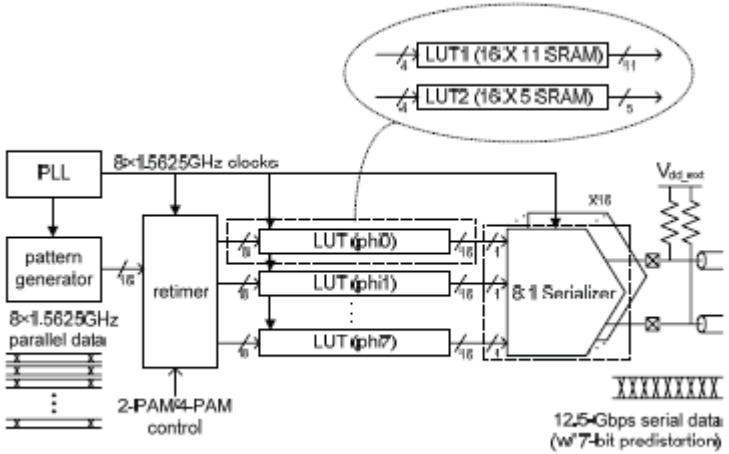
Claim 19	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

Claim 19	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.

1415

Claim 19	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.

1416

Claim 19	Disclosure
	 Fig. 1. Overall transmitter block diagram Chung, Figure 1.

1417

Claim 19	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

1418

Claim 19	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.

1419

Claim 19	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

1420

Claim 19	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.

1421

Claim 19	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.

1422

Claim 19	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.

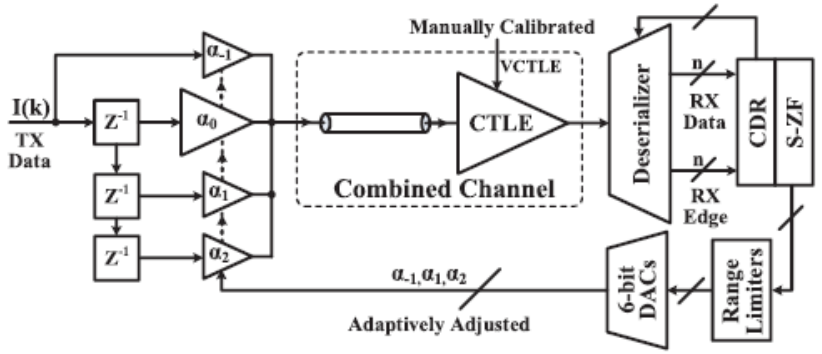


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

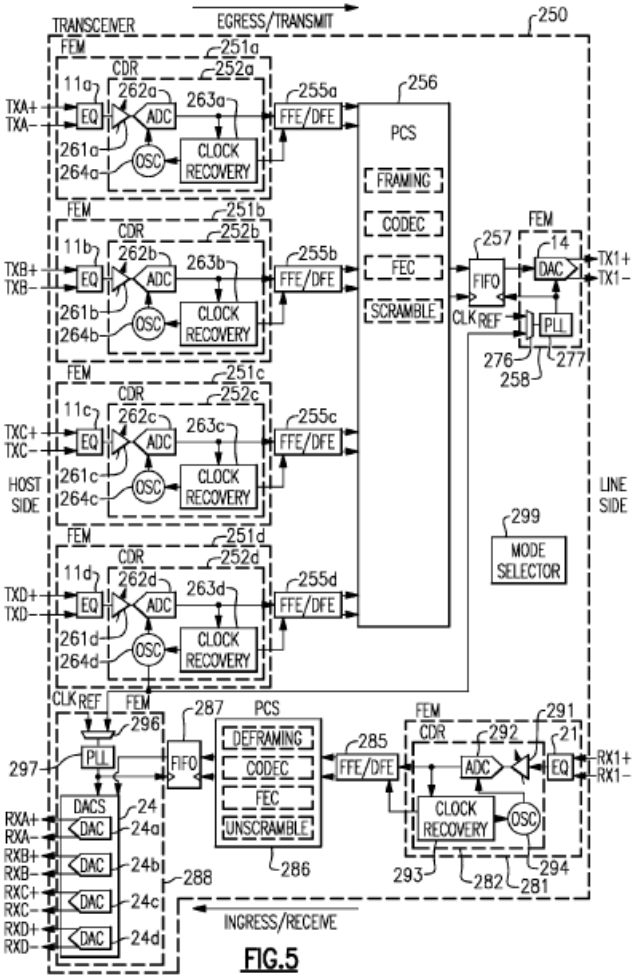
1423

Claim 19	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.

1424

Claim 19	Disclosure
1425	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 '233 Patent specification repeats the "receiver-based equalization" language of dependent claim 5, but fails to provide any explanation of what the phrase means. *See, e.g.*, '233 Patent at 2:42-45, 2:61-63. Nevertheless, "receiver-based equalization" is met by equalizing a received signal to compensate for distortion incurred during transmission. *See, e.g.*, U.S. Patent No. 9,178,542 at 10:22-35; U.S. Patent App. Pub. No. 2005/0078758 at [0004]; U.S. Patent No. 7,570,708 at 2:17-21; Pavan Kumar Hanumolu, *et al.*, "Equalizers for High-Speed Serial Links," International Journal of High Speed Electronics and Systems, vol. 15, no. 2, 2005 ("Hanumolu") at p.185 (describing "receive-side equalization" suppressing channel loss). Gorecki suggest that equalizers are, like filters, configured to compensate for signal distortion over the cable. *See* Gorecki at 1:59-67 ("[C]onventional high speed digital baseband communications systems often employ circuitry, for example, a finite impulse response filter ('FIR filter'), in the receiver…. Such

Claim 19	Disclosure
	circuitry typically includes [sic] one or more taps having fixed or pre-programmed 'positions' and coefficients."); *see id.* at 2:17-21 ("Regardless of where the equalization circuitry is implemented, the duration of the equalization signal and the relative position or placement of the tap(s) of the equalization circuitry are selected or designed to avoid interference with the signal representative of the transmitted information.").
	The equalization functionality depends on filter coefficient settings. *See, e.g.,* Hanumolu at pp. 185-198 (describing different receive-side equalization architectures configured with coefficients [C-1, C0, C1]); U.S. Patent App. Pub. No. 2014/0281068 at [0069] (describing applying coefficients to a receiver's settings to minimize communication data loss).
	Though Gorecki describes implementing its equalization structure and techniques at a transmitter "may have an advantage," (*see, e.g.*, Gorecki at 18:47-53), an ordinarily skilled artisan would have recognized from this statement that implementation in the receiver was also an option, particularly as the combination explains implementation more generally in connection with equalization circuitry more generally. *See* Gorecki at 2:17-21, 6:14-23 ("The circuitry and techniques include leading and/or trailing taps to reduce, minimize, mitigate or effectively eliminate precursor and/or post-cursor intersymbol interference due to, for example, bandwidth limitations and reflections in high-speed digital communication systems. In this way, the equalization circuitry and techniques may reduce, minimize or eliminate non optimum (e.g., over-equalization) at the boundaries of the data signal (i.e., the symbol)."). An ordinarily skilled artisan would have considered such an approach would have been applicable to equalizers 21 as they provide equalization in the same way as transmit filter 17.
	Additionally, an ordinarily skilled artisan would have understood that the equalization performed at both a transmit filter and a receive path equalizer should collectively compensate for the effects introduced the communication channel and would have had a reason to set equalization parameters for both in the same way.

Claim 19	Disclosure
	 FIG.5 Lugthart, Figure 5.

Claim 19	Disclosure
	The illustrated transmit path line-side FEM 258 includes a transmit path DAC 14, a multiplexer 276 and a phase-locked loop (PLL) 277. Additionally, the receive path line-side FEM 281 of FIG. 5 includes a receive path equalizer 21 and a receive path CDR circuit 282, which includes a thresholder 291, an ADC 292, a clock recovery circuit 293, and a controllable oscillator 294. Furthermore, the illustrated receive path host-side FEM 288 includes a multiplexor 296, a PLL 297, and a plurality of receive path DACs 24 including a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. Lugthart, 24:24–34. 10. A transceiver comprising: an equalizer configured to receive a data signal and to generate a clock and data recovery (CDR) input signal based on equalizing the data signal; and a CDR circuit comprising: a data recovery analog-to-digital converter (ADC) configured to receive the CDR input signal and to generate a first digital output signal; a clock recovery ADC configured to receive the CDR input signal and to generate a second digital output signal, wherein the data recovery ADC has a higher resolution than the clock recovery ADC; a timing control circuit configured to receive the second digital output signal and to generate a clock control signal based on the second digital output signal; and a clock generation unit configured to control a timing of the data recovery ADC and a timing of the clock recovery ADC based on the clock control signal. Lugthart, 58:64–59:13.

1428

Claim 19	Disclosure
	[0029] Example embodiments will now be described with reference to FIG. **1** which illustrates a host **102** that is connected through a host channel **103** to a pluggable network transceiver interface module **100**, which in turn is connected to a transmission medium **104**. In the illustrated embodiment, the transmission medium **104** includes fibre optic cable **124** for transmitting high speed data from the interface module **100** and fibre optic cable **126** for transmitting high speed data to the interface module **100**, however embodiments described herein can also be applied to non-fibre transmission mediums such as copper for example. In the illustrated embodiment, the host **102** includes a network interface card or board **106** that allows data to be exchanged between the host **102** and the network interface module **100**. In at least one example embodiment, the host interface board **106** includes data processing circuitry, which could for example be implemented an ASIC **108** ("Application Specific Integrated Circuit"), and a SERDES ("SERializer-DESerializer") chip or module **110**. SERDES module **110** receives parallel data from the host ASIC **108** over a parallel data channel **170** and converts it into high speed serial data for transmission over the host channel **103** to the network interface module **100**, and receives high speed serial data from the network interface module **100** through the host channel **103** which it then converts into parallel data and passes on to the ASIC **108** over the parallel data channel **170**. Tremblay, paragraph 29.

1429

Claim 19	Disclosure
	[0038] With reference to the drawings, FIG. **2** is a flow chart of the operation of an example embodiment, wherein the LA/CDR unit **144** in conjunction with the microcontroller **150** performs pre-equalization and optimizes power consumption of the transmission over host channel **103** via physical conductors **114** based on the feedback received from the SERDES module **110** via the low-speed digital interface **152**. In FIG. **2**, the operational flow of the LA/CDR unit **144** and the microcontroller **150** begins at step **202** when the microcontroller receives feedback via the low speed digital interface **152**. At step **204**, the microcontroller **150** instructs the LA/CDR unit **144** to equalize its transmitted data signal in response to the feedback, applying sufficient transmit pre-emphasis and other pre-equalization effects to the transmitted data signal to achieve predetermined target eye characteristics at the SERDES module **110**. Once the feedback from the SERDES module **110** indicates that the predetermined target eye characteristics have been achieved, the microcontroller **150** at step **206** instructs the LA/CDR unit **144** to begin reducing the launch amplitude of the transmitted data signal as much as it can without letting the received signal at the SERDES module **110** fall short of the predetermined target eye characteristics. The process is iterative, with feedback being monitored at step **202** and equalization being adjusted at step **204** while the launch amplitude is being reduced at step **206**. In alternative embodiments, the process may operate according to a different sequence of steps while achieving the same result, namely the optimization of signal quality and power consumption over the host channel **103**. Tremblay, paragraph 38.

1430

T. DEPENDENT CLAIM 20

Claim 20	Disclosure
20. The method of claim 15, wherein the first and second DRR devices do not perform pre-equalization of the multi-lane data streams provided to the first and second host interface ports.	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.

1431

Claim 20	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.

1432

Claim 20	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.

1433

Claim 20	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.

1434

Claim 20	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.

1435

Claim 20	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.

1436

Claim 20	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.

1437

Claim 20	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.

1438

Claim 20	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.

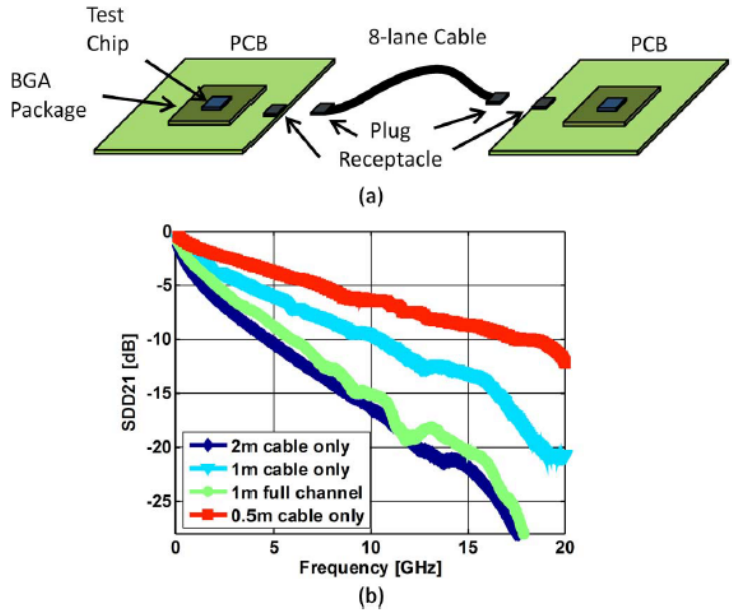
1439

Claim 20	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

1440

Claim 20	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.

1441

Claim 20	Disclosure
	 Fig. 1. (a) Interconnect topology. (b) Channel frequency response. Musah, 3080.

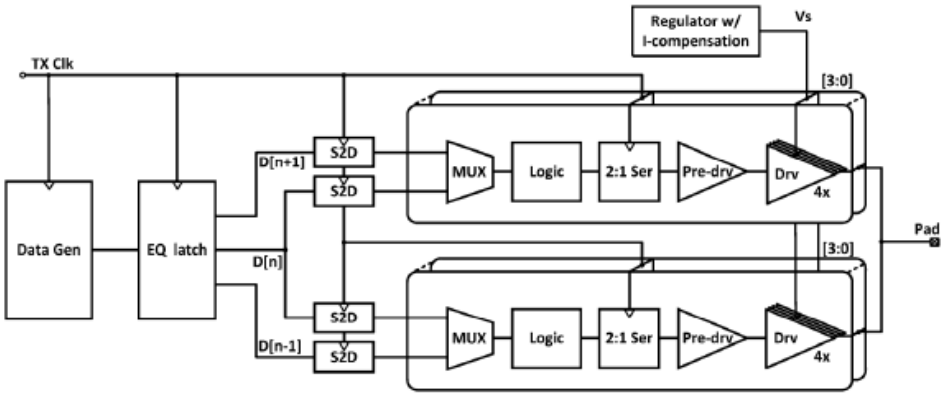
1442

Claim 20	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.

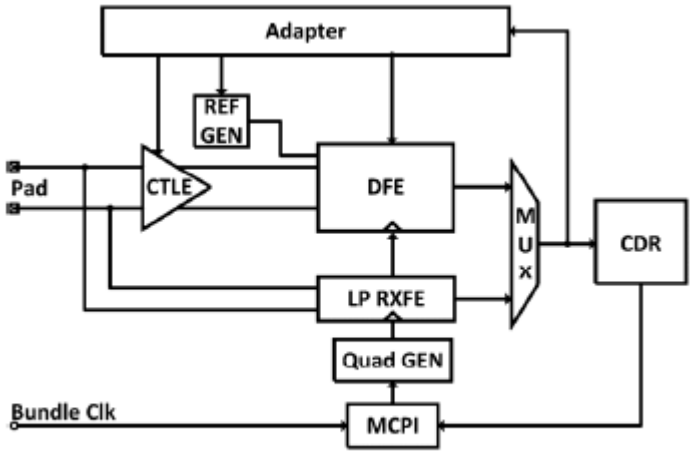
1443

Claim 20	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.

1444

Claim 20	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

1445

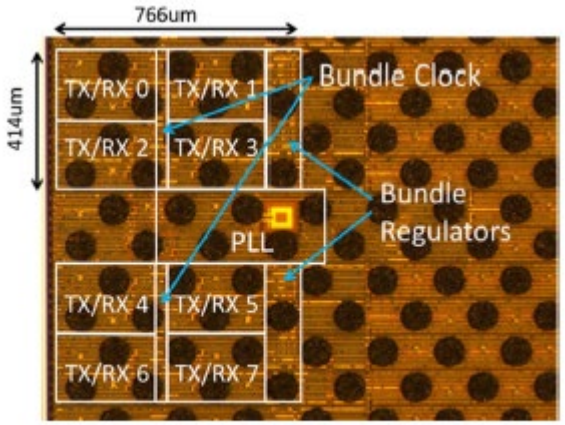
Claim 20	Disclosure
	simple comparator that contains a track-and-hold, sense amplifier and RS latch. Offset correction is included in the input of the sense amplifier. Since the LP RX is employed only at data rates below 12 Gb/s, the TX FFE is sufficient to mitigate channel induced inter-symbol interference (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.

1446

Claim 20	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 20	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.

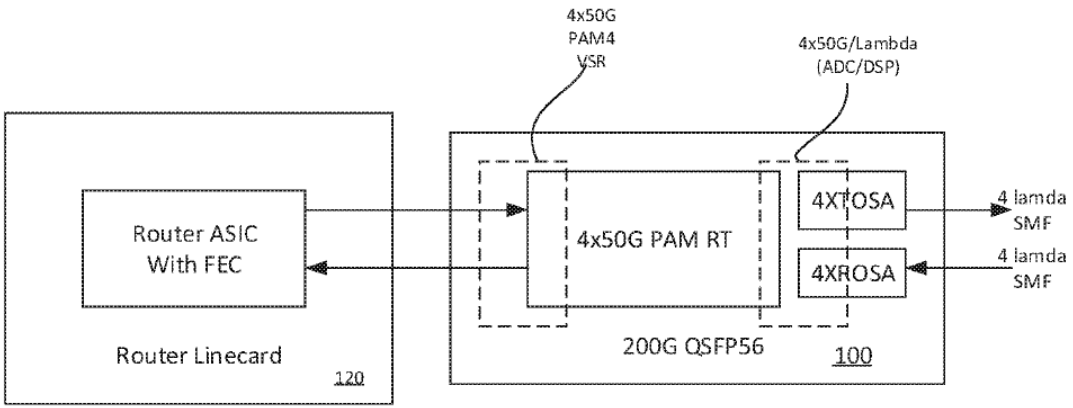
1448

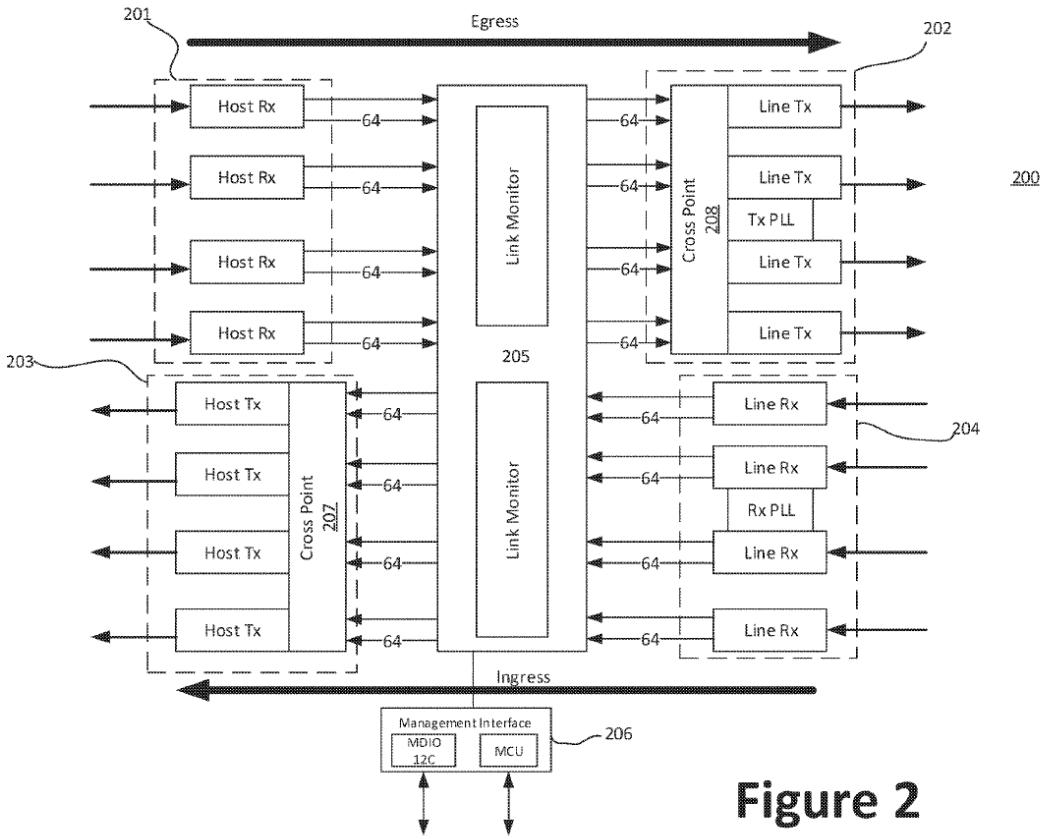
Claim 20	Disclosure
	Two 4-lane bidirectional transceivers, shown in Fig. 15, were fabricated in a 22 nm CMOS process. A 4-lane transceiver occupies 0.317 mm^2. In addition to the cables listed in Section I, a combination of cables and FR-4 PCB traces were employed to realize different channel losses to characterize the link performance at lower data rates. Musah, 3086.  Fig. 15. 22 nm CMOS test chip micrograph. Musah, 3086.

1449

Claim 20	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.

1450

Claim 20	Disclosure
1451	 **Figure 1** Riani, Figure 1.

Claim 20	Disclosure
	 Figure 2 Riani, Figure 2.

Claim 20	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.

1453

Claim 20	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 20	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.

1455

Claim 20	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.

1456

Claim 20	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.

1457

Claim 20	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.

1458

Claim 20	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.

1459

Claim 20	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.

1460

Claim 20	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.

1461

	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.	

1462

Claim 20	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.

1463

Claim 20	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

1464

Claim 20	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.

1465

Claim 20	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.

1466

Claim 20	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.

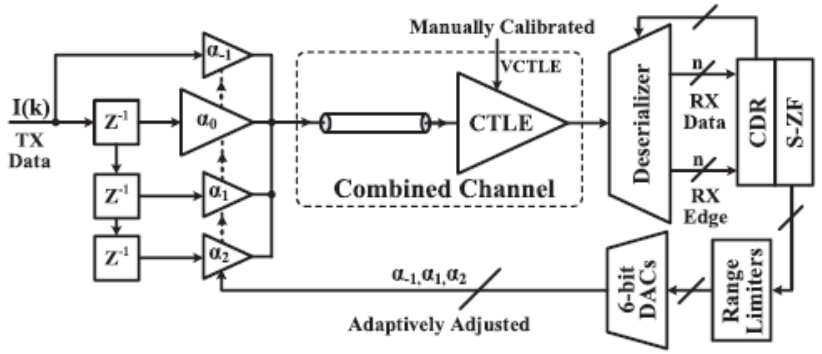


Fig. 17. Implemented equalization scheme with the proposed S-ZF algorithm.

1467

Claim 20	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.

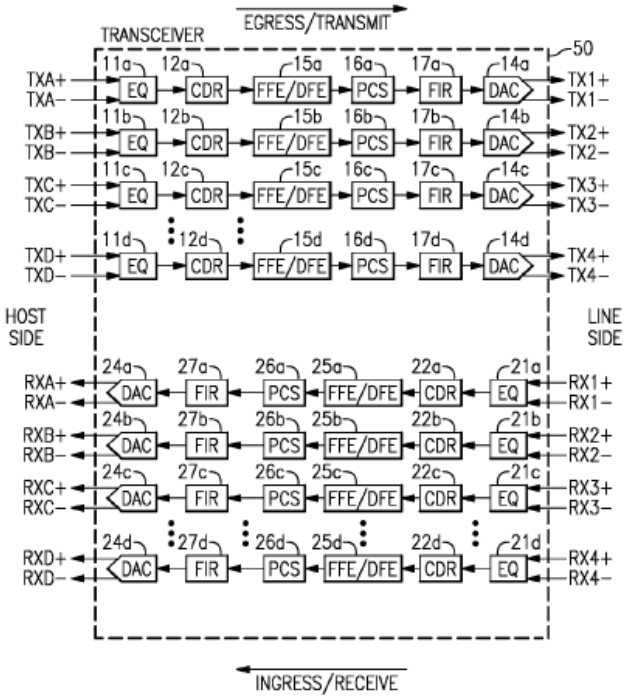
1468

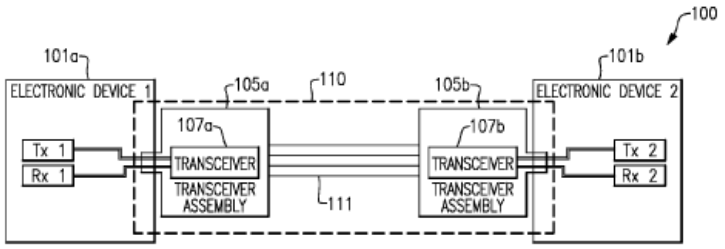
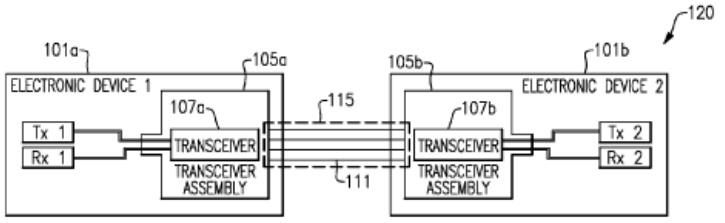
Claim 20	Disclosure
	This paper implements a 40-Gb/s TX and RX chipset over a >16-dB loss PCB channel using a 65-nm CMOS process. The TX utilizes a bandwidth-enhanced 4:1 MUX and an interleaved-retiming latch array to obtain wide operation range, high power efficiency, and small area occupation. By introducing 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. *IEEE 802.3-2015* The IEEE 802.3-2015 standard discloses setting transmit filter coefficients to disable pre-equalization. As was known, an ordinarily skilled artisan would have disabled pre-equalization "if not needed to consume minimum power while satisfying the performance or margin objective" for a system. *See* U.S. Patent No. 9,137,063 to Zerbe at 27:24-37. Thus, for instances where signal transmission between host devices using the active cable meets performance objectives without pre-equalization of the multi-lane data streams provided to the first and second host interface ports, a user would disable those pre-equalization based on the IEEE 802.3-2015 standard.

1469

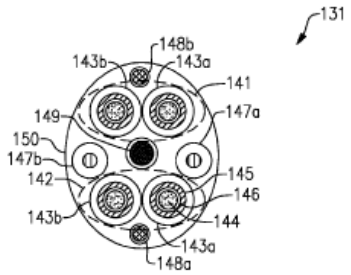
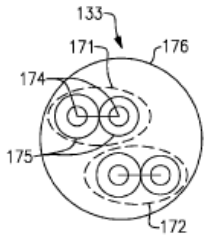
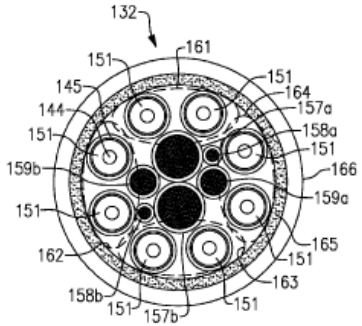
Claim 20	Disclosure
	72.6.10.3.4 Functions COEFF_UPDATE(coefficient, preset, initialize, inc, dec) Returns an updated coefficient based on the contents of the coefficient update field in the training frame. Sets a fixed coefficient value, or adds, or subtracts from the current coefficient value to create the updated coefficient. If multiple actions are requested in the coefficient update field, then the priority is: 1) preset 2) initialize 3) inc/dec Values: preset; If preset is TRUE then the function returns the coefficient value equivalent to no equalization [*c*(–1) and *c*(1) coefficients are set to zero, *c*(0) set to maximum]. initialize; If initialize is TRUE, then the function returns the coefficient value such that the transmit output meets the conditions defined in 72.6.10.4.2. inc; If inc is TRUE then the function returns (coefficient + step). dec; If dec is TRUE then the function returns (coefficient – step). The requirements for the value of step are defined in 72.7.1.10 and Table 72–7. IEEE 802.3-2015, Section 5 at 480.

1470

Claim 20	Disclosure
	 FIG.1D Lugthart, Figure 1D.

Claim 20	Disclosure
	 FIG.2A  FIG.2B Lugthart, Figures 2A and 2B.

1472

Claim 20	Disclosure
	 FIG.3A  FIG.3C  FIG.3B Lugthart, Figures 3A –3C. According to yet another aspect, an active cable comprises a first cable comprising a first pair of conductors including a first pluggable module. The first pluggable module can include a first transceiver

Claim 20	Disclosure
	including host-side analog-to-digital converter (ADC) circuitry configured to receive a first plurality of host-side signals and to generate a plurality of digitized egress signals. Lugthart, 3:66–4:5. The illustrated configuration operates using differential signaling on the host and line sides. Using differential signaling can improve performance in a variety of ways, including, for example, enhancing performance in the presence of common-mode noise and/or systematic offsets or errors. In certain configurations, a particular differential signal is transmitted or received over a pair of differential micro coaxial cables or conductors. In certain configurations, two or more pairs of differential micro coaxial conductors can be housed in a shared cable. Various examples of such cables will be described in detail further below with reference to FIGS. 3A-3C. Although various examples of cabling have been described, other configurations are possible, including, for example, other implementations of electrical cables and/or implementations using optical cables. Lugthart, 11:51–65. FIG. 1D is a schematic diagram of another embodiment of a transceiver 50. The transceiver 50 includes a first transmit path equalizer 11 a, a second transmit path equalizer 11 b, a third transmit path equalizer 11 c, a fourth transmit path equalizer 11 d, a first transmit path CDR circuit 12 a, a second transmit path CDR circuit 12 b, a third transmit path CDR circuit 12 c, a fourth transmit path CDR circuit 12 d, a first transmit path FFE/DFE block 15 a, a second transmit path FFE/DFE block 15 b, a third transmit path FFE/DFE block 15 c, a fourth transmit path FFE/DFE block 15 d, a first transmit path PCS block 16 a, a second transmit path PCS block 16 b, a third transmit path PCS block 16 c, a fourth transmit path PCS block 16 d, a first transmit path FIR 17 a, a second transmit path FIR 17 b, a third transmit path FIR 17 c, a fourth transmit path FIR 17 d, a first transmit path DAC 14 a, a second transmit path DAC 14 b, a third transmit path DAC 14 c, a fourth transmit path DAC 14 d, a first receive path equalizer 21 a, a second receive path equalizer 21 b, a third receive path equalizer 21 c, a fourth receive path equalizer 21 d, a first receive path CDR circuit 22 a, a second receive path CDR circuit 22 b, a third receive path CDR circuit 22 c, a fourth receive path CDR circuit 22 d, a first receive path FFE/DFE block 25 a, a second receive

1474

Claim 20	Disclosure
	path FFE/DFE block 25 b, a third receive path FFE/DFE block 25 c, a fourth receive path FFE/DFE block 25 d, a first receive path PCS block 26 a, a second receive path PCS block 26 b, a third receive path PCS block 26 c, a fourth receive path PCS block 26 d, a first receive path FIR 27 a, a second receive path FIR 27 b, a third receive path FIR 27 c, a fourth receive path FIR 27 d, a first receive path DAC 24 a, a second receive path DAC 24 b, a third receive path DAC 24 c, and a fourth receive path DAC 24 d. The illustrated transceiver 50 operates as a re-timer. For example, the transceiver 50 can retime one or more transmit signals received on the host side to generate one or more transmit signals on the line side. Similarly, the transceiver 50 can re-time one or more receive signals received on the line side to generate one or more receive signals on the host side. Re-timing signals on the host and/or line signals can reduce or remove high-frequency jitter in the signals, thereby improving signal fidelity and/or permitting communication over longer distances. In the illustrated configuration, the transceiver 50 receives four host side differential transit signals, which are retimed to generate four line side differential transmit signals. Additionally, the illustrated transceiver 50 receives four line side differential receive signals, which are retimed to generate four host side differential receive signals. Although the transceiver 50 is illustrated as retiming four transmit signals and four receive signals, the teachings herein are applicable to configurations in which a transceiver retimes more or fewer transmits signals and/or more or fewer receive signals. Lugthart, 12:1–53. FIG. 2A illustrates an embodiment in which the cable 110 is an active cable including actively powered componentry for improving performance of the cable 110. Thus, the cable 110 shown in FIG. 2A includes one or more conductive lines 111 as well as first and second transceiver assemblies 105 a, 105 b positioned at either end of the conductive lines 111. For example, the first transceiver assembly 105 a includes a host side that is electrically connected to the first electronic device 101 a and a line side that is electrically connected to a first end of the conductive lines 111. Additionally, the second transceiver assembly 105 b includes a host side that is electrically connected to the second electronic device 101 b and a line side that is electrically connected to a second end of the conductive lines 111. Although FIG. 2A illustrates a configuration including the conductive lines 111, in other configurations the conductive lines 111 can be omitted in favor of communication using an optical cable.

1475

Claim 20	Disclosure
	The first and second transceiver assemblies 105 a, 105 b can be implemented in a variety of ways. For example, the first and second transceiver assemblies include first and second transceivers 107 a, 107 b, respectively, which can be implemented using, for example, any of the embodiments of transceivers described earlier with respective to FIGS. 1A-1D. FIG. 2B is a schematic diagram of a high-speed point-to-point communication system 120 according to another embodiment. In contrast to the communication system 100 shown in FIG. 2A which includes an active cable 110, the communication system 120 of FIG. 2B includes a passive cable 115. As shown in FIG. 2B, active componentry including the transceivers 107 a, 107 b is situated externally to the cable 115. For example, the first transceiver assembly 105 a can be included in the first electronic device 101 a, and the second transceiver assembly 105 b can be included in the second electronic device 101 b. Lugthart, 14:3–38. In various implementations, the packaging can have a design and a form factor similar to the packaging defined by an existing interface standard (e.g., Thunderbolt Gen1/Gen2 or USB 2.0/3.0), although the packaging can have some other form factor. The packaging can comprise materials suitable for electronic and computer interfaces, such as plastic, metal or a composite material. Each of the transceiver assemblies 105 a and 105 b has an input port that is configured to mechanically and electrically connect, e.g., in a releasable fashion, to a corresponding port or other interface on the respective electronic device 101 a, 101 b, and an output port that is connected to the cable. The input port can include a connector similar to connectors of existing interface standards, including, but not limited to, Thunderbolt, USB 2.0/3.0, small form-factor pluggable (SFP), enhanced small form-factor pluggable (SFP+), and/or quad small form-factor pluggable (QSFP). Lugthart, 15:8–23.

1476

Claim 20	Disclosure
	[0037] Accordingly, the system of FIG. **1** provides a back-channel using a low speed interface to set the operating parameters of the serializer section **160** of SERDES module **110** and the operating parameters of the LA/CDR module **144** of the transceiver interface module **100**. In some applications, this allows the operating parameters to be change over time which may in turn facilitate adaptive optimization of high data speed transmission characteristics and power consumption used for the communications link. For example, some embodiments may optimize power consumption used for the communications link by reducing the launch amplitude of the signals transmitted across the host channel **103**. By monitoring the eye characteristics of the signal at the SERDES eye monitor **164** and the CDR eye monitor **142** and modifying the transmit equalization parameters used by the LA/CDR module **144** and the serializer section **160** of the SERDES module **110**, respectively, the predetermined target eye characteristics can be maintained through transmit equalization even while launch amplitude is decreased. Some embodiments may also optimize power consumption by deactivating the circuits that perform transmit equalization when the monitored eye characteristics indicate that pre-equalization is unnecessary for maintaining the predetermined target eye characteristics. Tremblay, paragraph 37.

1477