

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Marvell Semiconductor, Inc.

Petitioner,

v.

Credo Technology Group Ltd.,

Patent Owner.

Case No. IPR2025-01220

U.S. Patent 11,032,111

Declaration of Michael S. Chen

Table of Contents

I.	Overview of the Technology	6
A.	Equalization.....	8
B.	Channel Loss & Transmission Performance.....	13
II.	The '111 Patent.....	16
A.	Prosecution History	17
B.	Priority Date	19
C.	Challenged Claims	19
III.	Level of Ordinary Skill in the Art	20
IV.	Claim Construction.....	22
V.	Invalidity Grounds	23
A.	Ground 1	23
1.	U.S. Patent Application Publication No. 2019/0288881 ("Berke") (Ex-1004)	24
2.	U.S. Patent No. 8,516,238 ("Cornelius") (Ex-1005)	26
3.	Independent Claim 1	27
4.	Dependent Claim 3.....	49
5.	Dependent Claim 4.....	51
6.	Dependent Claim 6.....	53
7.	Dependent Claim 7.....	53
8.	Independent Claim 8	55
9.	Dependent Claim 10	66
10.	Independent Claim 11	67

11.	Dependent Claim 13	69
12.	Dependent Claim 14	69
13.	Dependent Claim 15	70
14.	Independent Claim 16	72
15.	Dependent Claim 17	72
16.	Dependent Claim 19	73
B.	Ground 2	73
1.	U.S. Patent No. 10,715,357 (“Ran”) (Ex-1006)	73
2.	Claims 1, 3, 4, 6, and 7	77
3.	Independent Claim 8	78
4.	Dependent Claim 10	87
5.	Independent Claim 11	88
6.	Dependent Claims 13-15	89
7.	Independent Claim 16	90
8.	Dependent Claims 17 and 19	90
C.	Grounds 3 and 4	90
1.	High Speed Serdes Devices and Applications (“Stauffer”) (Ex-1008)	91
2.	Dependent Claim 2	92
3.	Dependent Claim 5	97
4.	Dependent Claim 9	102
5.	Dependent Claim 12	102
6.	Dependent Claim 18	102

D.	Ground 5.....	103
1.	U.S. Patent Application Publication No. 2014/0237301 ("Mejia") (Ex-1009).....	103
2.	Independent Claim 1	105
3.	Dependent Claim 2.....	128
4.	Dependent Claim 3.....	131
5.	Dependent Claim 4.....	134
6.	Dependent Claim 5.....	135
7.	Dependent Claim 6.....	139
8.	Dependent Claim 7.....	139
9.	Independent Claim 8	140
10.	Dependent Claim 9.....	144
11.	Dependent Claim 10	145
12.	Independent Claim 11	145
13.	Dependent Claim 12	147
14.	Dependent Claim 13	147
15.	Dependent Claim 14	148
16.	Dependent Claim 15	149
17.	Independent Claim 16.....	150
18.	Dependent Claim 17	150
19.	Dependent Claim 18	150
20.	Dependent Claim 19	151
VI.	Background and Qualifications	151

A.	Compensation.....	154
B.	Materials and Other Information Considered	154
VII.	Understanding of the Law	154
VIII.	Reservation of Rights	155
Appendix 1:	Curriculum Vitae.....	156
Appendix 2:	Materials Considered	171
Appendix 3:	Challenged Claims	172
Appendix 4:	Understanding of the Law.....	176

I, Michael Shuo-Wei Chen, declare as follows:

1. My name is Michael Shuo-Wei Chen. I am a Professor in the Department of Electrical Engineering at the University of Southern California. A summary of my background and qualifications can be found in §VI below. My full Curriculum Vitae is attached to this declaration as Appendix 1.

2. I have prepared this report as an expert witness retained by Marvell Semiconductor, Inc. In this report I give my opinions as to whether certain claims of U.S. Patent No. 11,032,111 (“the ’111 patent”) are invalid. I provide technical bases for these opinions as appropriate.

3. This report contains statements of my opinions formed to date and the bases and reasons for those opinions. I may offer additional opinions based on further review of materials in this case, including opinions and/or testimony of other expert witnesses. I make this declaration based upon my own personal knowledge and, if called upon to testify, would testify competently to the matters contained herein.

I. Overview of the Technology

4. High-speed serial communications became popular and widespread in the 1990s. In contrast to parallel architectures where multiple lines send different data signals in parallel, serial communications send a single signal down a single channel. As transmission speeds increased, parallel data lines suffered from

limitations from having to route numerous data lines next to each other (in parallel) along with clock signals. Ex-1008, 000011-000012. Parallel signals must be synchronized so that all signals and their shared clock arrived at the receiver within a certain window, which became increasingly difficult as bus widths and clock speeds increased. *See* Ex-1008, 000011-000012, 000013-000014. Serial communications techniques avoid these issues while saving pin count on the integrated circuit (“IC”).

5. By the 2000-2005 timeframe, computer bus and networking protocols had largely transitioned to serial connections, where parallel data within a chip was converted into a serial format, sent over a serial link, and then converted back into a parallel format. This process of serialization and then de-serialization, called “SerDes,” was used in various standardized bus and networking protocols. For example, the parallel ATA standard was replaced by Serial ATA, or “SATA” for short. The parallel PCI standard was replaced by a serial version called PCIe. The Universal Serial Bus (“USB”) replaced various parallel interfaces.

6. Ethernet was a popular networking standard going back to the 1990s. By the 2010s, Ethernet was widely used for high-speed serial communications in enterprise and data-center networks, scaling from Gigabit Ethernet in the early 2000s to the widespread adoption of 10 GbE by 2010. And the ubiquitous Ethernet

standard had long been implemented using serial cables and CAT-5 plugs. Ex-1008, 000058-000061.

A. Equalization

7. As connection protocols shifted to serial links, they adopted established serial communications techniques. One widely used technique was equalization, which adjusted for channel losses. Equalization performed at the transmitter was called “pre-equalization” because it adjusted the signal at the transmitter, before a signal is sent over the channel, to compensate for anticipated signal loss. Equalization performed at the receiver adjusted for channel losses after the signal crossed the channel. The use of these types of equalization allowed robust and reliable transmission over noisy channels and at high data rates. Ex-1008, 000020-000024.

8. Pre-equalization was a well-established serial communication technique. By 2018, and for many years prior, it was one of several common components taught in undergraduate courses on high-speed serial communications. Ex-1008, 000020-000024. Textbooks on serial communications typically included a chapter or section explaining pre-equalization. Ex-1008, 000020-000024; Ex-1022, §1.1.1.

9. Because a channel affects different frequencies unevenly, pre-equalization boosts the signal at frequencies where greater losses were expected.

This compensation counteracted channel losses, so that after traversing the channel, the received signal more closely matched the original. The resulting signal at the receiver had a signal spectrum more constant over frequency, resulting in less waveform distortion and reduced inter-symbol interference (“ISI”). Pre-equalization was most commonly performed using “pre-emphasis” to boost certain frequencies; it is illustrated below, where the red boxes show the original waveform with and without pre-emphasis, and the waveforms on the right show the result after passing through the channel. Ex-1008, FIGS. 1.9-1.10 at 000021:¹




Fig. 1.9 Signal distortion for a typical channel application




Fig. 1.10 Typical channel application with equalization at the transmitter

¹ All emphasis and annotations added unless stated otherwise.

10. Pre-equalization at transmitter is often implemented using a finite impulse response (“FIR”) filter, including a specific implementation called a Feed-Forward Equalizer (“FFE”). Ex-1008, 000050. An FIR filter essentially performs a weighted sum between adjacent data symbols by multiplying coefficients with adjacent data (also called “taps”) and summing the results. For pre-equalization, the FIR filter coefficients are chosen so as to boost frequencies that the channel diminishes. This way, when the signal goes through both the filter and the channel, their effects cancel out, helping the signal reach the receiver closer to its original form. Typically, the FIR filter provides more signal gain at higher frequencies (so-called pre-emphasis) and attenuates the signal more at lower frequencies (so-called de-emphasis) to compensate for the frequency-dependent channel loss. Ex-1008, 000020-000021; Ex-1022, 000025-000027.

11. On the other side of the channel, equalization was used at the receiver to further adjust for channel loss. Some popular equalizers at the receiver end include continuous time linear equalization (“CTLE”), Feed-Forward Equalizer (“FFE”), and Decision Feedback Equalizers (“DFE”). The equalizers here have a similar objective as pre-equalization at the transmitter, which is to compensate for the channel loss such that it equalizes the frequency response distortion of the channel. Since the digital communication signal would go through the transmitter, physical channel, and then the receiver, the overall equalization can be jointly co-

designed between transmitter pre-equalization and receiver equalization to compensate for the channel loss and hence minimize ISI or data-dependent jitter (“DDJ”). Therefore, receiver equalization is useful when transmitter pre-equalization alone cannot sufficiently compensate for channel loss, or when the physical channel loss varies beyond what transmitter pre-equalizer has anticipated.

12. Many serial communications systems use a coefficient training process to configure equalization for the channel, where the transmitter and the receiver communicates to jointly determine equalization coefficients to use for data transmission. During coefficient training, the channel was first configured with initial equalization coefficient values and a known reference pattern or training sequence was sent through the channel. The receiver compared the received pattern to the expected pattern to evaluate how it was distorted by the channel. The coefficients were iteratively adjusted to improve link performance. As the coefficients were refined, the output signal more closely matched the original transmitted data. Training was common because it improved signal integrity and transmission reliability, especially as channel conditions changed due to temperature, aging components, or other environmental factors. *See generally* Ex-1008, 000067-000069; Ex-1004, ¶3; Ex-1006, Abstract, 3:29-50, 15:45-50; Ex-1009, ¶32. The effectiveness of the training process depended on the quality of the initial coefficient values, where coefficients that were closer to the optimal values

reduced the number of training iterations required and increased the likelihood of converging to a stable equalization setting without overcompensation or oscillation. It was common practice for serial communications systems to store multiple sets of baseline equalization coefficient values for use as the starting point or initial values for coefficient training. Ex-1004, ¶¶30-32; Ex-1006, 10:6-14, Fig. 3B:

Bit 13	Bit 12	Description
1	1	Preset3
1	0	Preset2
0	1	Preset1
0	0	Co-efficient Control

FIG. 3B

The baseline values were referred to as preset or whitelisted values. *Id.* Different sets of baseline equalization coefficient values corresponded to different channel characteristics such as layout, design, configurations, settings, and environmental conditions, such that one of the sets could be selected for use as the starting point of training based on its performance for a particular channel. Ex-1004, ¶¶30-33, Ex-1009, ¶¶31-32, ¶36. The storage of multiple sets of baseline equalization coefficient values allowed the system to accommodate channels with different characteristics.

B. Channel Loss & Transmission Performance

13. In the area of wireline communication, data (typically in the form of digital binary bits: “0” or “1”) is conveyed from one end (transmitting end) to the other end (receiving end) via a physical channel, such as a cable or a trace on a printed circuit board (“PCB”). For example, “0” can be conveyed using low voltage, while “1” can be conveyed using high voltage. If the physical channel is perfect without any loss, the receiver simply needs to detect whether the received voltage level is above or below a certain threshold to recover the transmitted bits. However, in practice, physical channels are not perfect and include various loss mechanisms, such as parasitic resistance, capacitance, skin effects, dielectric loss, connectors, vias, etc., which attenuate and distort signals transmitted through the physical channel.

14. If such a channel loss were strictly constant over frequency, the lost signal energy will result in smaller signal swing at the receiving end but the transmitted signal waveform can be retained. However, in reality, the channel loss is frequency-dependent and can distort the signal waveform at the receiving end and interfere between data symbols, resulting in ISI. Ex-1008, 000073 (“This type of jitter [DDJ], also known as pattern dependent jitter or intersymbol interference (ISI), is observed on nonclocklike waveforms since the frequency spectrum of such signals is continually changing ... high baud rates still expose these signals to a

significant amount of signal attenuation resulting in DDJ contribution.”), 000071-000072, 000074-000075. The ISI effect can cause the received waveform to be smeared to anywhere between the high and low voltage. Sometimes, the supposedly high voltage may be attenuated to a much lower voltage. As a consequence, the receiver is unable to detect the intended data reliably and can cause bit error (e.g. detected as “0” while the transmitted data is “1” and vice versa), often characterized as bit error rate (“BER”).

15. When evaluating performance of high-speed serial communications systems, BER, and opening of an eye diagram were common criteria for assessing the signal quality and reliability of the system. *See generally* Ex-1008, 000014-000015, 000020-000021.

16. BER measures the number of bits that end up with the wrong value. It is a ratio of the number of bits that are received in error to the total number of bits transmitted over a given period of time. BER quantifies how often errors occur due to noise, interference, or distortion in the transmission channel. A lower BER indicates a more reliable communication link, while a higher BER suggests that the system is experiencing more frequent errors. *See generally* Ex-1008, 000014-000015.

17. Eye diagrams overlay multiple bits of the signal at the receiver. For simple eye diagrams, the high-voltage value was the top line, the low-voltage value

was the bottom, and the sides showed what the waveform looked like as it transitioned between high and low values:




FIG. 2

Ex-1004, FIG. 2, ¶22 (“Receiver eye diagram 200 shows an eye height 202 and an eye width 204.”). The opening of an eye diagram was a well-known metric taught in undergraduate courses for assessing the overall link quality in high-speed serial

communications, including the impact of equalization. The “eye opening” is the empty, open area in the center of the eye diagram. A large and well-defined eye opening indicates that the signal is clean, with clear differences between the high and low values of the received signal, allowing accurate data recovery at the receiver. Conversely, a smaller or partially closed eye shows muddled differences between high and low values, which make bit errors more likely because the difference between a 1 and 0 are difficult to distinguish. *See generally* Ex-1008, 000014-000015.

II. The '111 Patent

18. According to the '111 patent, the IEEE 802.3 standard provided a common specification for communications over “coaxial cable, twisted wire pair cable, fiber optic cable, and electrical backplanes, with various channel signal constellations.” Ex-1001, 1:19-28. “As demand continues for ever-higher data rates, extensions to the standard must account for increased channel attenuation and dispersion even as the equalizers are forced to operate at faster symbol rates.” Ex-1001, 1:28-34. To address channel loss and the need for equalizers to operate at faster symbol rates, the '111 patent purports to select initial pre-equalizer coefficients, stored in registers, and then use link training to refine them. Ex-1001, 1:39-58.

A. Prosecution History

19. The application for the '111 patent was allowed after two Office Actions. In the non-final Office Action dated July 9, 2020, the Examiner rejected original claims 1, 6-10, and 17-19 as being obvious under §103 in view of prior art references including U.S. Patent Application Publication No. 2013/0223506 (“Kolze”) (Ex-1017), U.S. Patent Application Publication No. 2010/0232493 (“Thirumoorthy”) (Ex-1018), and U.S. Patent Application Publication No. 2013/0201316 (“Binder”) (Ex-1019). Ex-1002, 000331-000345. The Examiner relied on Kolze for selecting pre-equalizer settings and updating pre-equalizer coefficient values, Thirumoorthy for storing pre-equalizer coefficients in registers, and Binder for suggesting a port connector coupling a port transceiver to a pluggable module transceiver. *Id.*, 000334-000340. The Examiner indicated that claims 2-5, 11-16, and 20 recited allowable subject matter. *Id.*, 000340.

20. In response, the Applicant did not dispute the Examiner’s rejection under §103 but sought to overcome the Office Action by incorporating allowable subject matter into the rejected claims. Ex-1002, 000370-000375. The Applicant amended independent claim 1 to incorporate limitations of original claim 5 and recite “determining a performance characteristic for initial pre-equalizer coefficient values in each of multiple registers.” *Id.*, 000371. The applicant amended claim 2 into independent form and amended independent claims 9 and 19 to purportedly

incorporate the limitations of claims 14 and 20, respectively. *Id.*, 000371-000375. Original claims 2, 14, and 20 each recited that “each of the multiple registers correspond[s] to a different channel model.” *Id.*, 000371. In addition, the applicant amended claim 11, which originally recited “wherein the port transceiver iteratively selects different ones of the multiple registers and determines a performance characteristic for each,” into independent form. *Id.*, 000372. In the final Office Action dated February 4, 2021, the Examiner rejected claims 1 and 6-8 as obvious in view of the combination of Kolze and Thirumoorthy. Ex-1002, 000389-000396. Again, the Applicant did not dispute the Examiner’s rejection. *Id.*, 000413-000418. Instead, the Applicant canceled claim 1 and amended claims 6-8 to depend from allowed independent claim 2. *Id.*, 000414-000416. Allowed independent claims 2, 9, 11, and 19 were issued as claims 1, 8, 11, and 16, respectively. I understand that the Examiner did not compare the claims identified to recite allowable subject matter with any prior art reference in either Office Action. Nor did the Examiner or the Applicant explain why these claims were non-obvious in view of the prior art.

21. None of the references used in the invalidity grounds discussed below were considered during examination. These references are not listed on the face of the patent, nor have I seen them addressed in the prosecution history.

B. Priority Date

22. The application that issued as the '111 patent was filed on August 27, 2019, claiming priority to Provisional Application No. 62/723,701 ("111-Provisional") filed on August 28, 2018.

23. I have reviewed 111-Provisional (Ex-1021) and compared it with the '111 patent. The specification of the '111 patent includes disclosure that is not in 111-Provisional. For example, 111-Provisional does not mention "channel model," "chip-to-module (C2M) communications link," or "pluggable module," which are limitations recited in the independent claims of the '111 patent. The '111 patent is not entitled to the priority date of its provisional application because each of its independent claims is not adequately supported by the written description of the provisional application.

24. Nonetheless, in this Declaration, I have performed my analysis using the assumption that the claims of the '111 patent are entitled to claim priority to August 28, 2018. My analysis does not change for obviousness as of August 27, 2019 because the Challenged Claims would be even more obvious a year later, and each of the prior art references would still be prior art.

C. Challenged Claims

25. I understand that Petitioner is challenging the validity of claims 1-19 of the '111 patent in the Petition for Inter Partes Review to which this declaration

will be attached. Those claims are reproduced in Appendix 3. I have considered all claims 1-19 of the '111 patent, as well as portions of the '111 patent prosecution history in forming my opinions.

III. Level of Ordinary Skill in the Art

26. I have analyzed the '111 patent and determined that the field of the patent is serial communications. *See, e.g.*, Ex-1001, Abstract, 1:19-35. The patent is directed to methods for sending data in serial form over a channel. *Id.*, Abstract (“a corresponding serial symbol stream being transmitted to the port transceiver”), 1:49-54, cls.11, 16. For example, the specification states that the communications method of the '111 patent is implemented on a transceiver “as application-specific integrated circuitry for very high-rate serial data transmission and reception.” *Id.*, 5:10-16, Abstract, 1:39-41, cls.1-20. Figures 4-5 describe a transmit and receive chain of communication channels. *Id.*, 5:24-34, 10:4-11. On the transmit chain, serializers “each accept a stream of transmit data blocks and convert the stream of blocks into a (higher-rate) stream of channel symbols.” *Id.*, 6:34-37. The stream of channel symbols are transmitted via the channel from the transmitter to the receiver. *Id.*, FIGS. 4-5. On the receive chain, deserializers “group the digital receive data stream bits or symbols into blocks.” *Id.*, 10:39-40. Independent claims 8, 11, and 16 all recite “serial” communications, and claim 1 recites “SerDes,” which includes serialization for serial communications. *See above* §I.A.

27. In determining the characteristics of a hypothetical person of ordinary skill in the art (“POSITA”) of the ’111 patent at the time of the claimed invention, I considered several things, including various prior art techniques relating to the serial communications, the type of problems that such techniques gave rise to, and the rapidity with which innovations were made. I also considered the sophistication of the technologies involved, and the educational background and experience of those actively working in the field at the time. I also considered the level of education that would be necessary to understand the ’111 patent. Finally, I placed myself back in the relevant period of time and considered the engineers and programmers that I have worked with and managed in the field of serial communications.

28. I came to the conclusion that a person of ordinary skill in the field of art of the ’111 patent would have been a person with a bachelor’s degree in electrical or computer engineering with at least three years of experience in serial communications. A person with more education but less practical experience may also meet this standard.

29. In the relevant timeframe (*see* above §II.B), a POSITA with the above-described level of skill would have been able to:

- Understand common components in high-speed serial communications systems, including but not limited to transceivers, equalization, pre-equalization, FIR filters, and controllers;
- Evaluate channel loss, compensation, and performance of high-speed serial communication systems using various indicators such as BER and eye diagrams;
- Design and implement high-speed serial communication systems;
- Design and implement equalization and pre-equalization filters, including by using an FFE and optimizing tap coefficients;
- Design and implement training methods for equalization coefficients;
- Program a digital signal processor (“DSP”) to perform common functions in transceivers, including pre-equalization and equalization;
- Implement protocols for calibrating and setting equalization parameters; and
- Configure host port connector and cable/card plugs with embedded transceivers.

IV. Claim Construction

30. For purposes of this inter partes review, I have considered the claim language, specification, and portions of the prosecution history, to determine the meaning of the claim language as it would have been understood by a person of

ordinary skill in the art at the time of the invention. The “plain and ordinary meaning” or *Phillips* standard has traditionally been applied in district court litigation, where a claim term is given its plain and ordinary meaning in view of the specification from the view point of a person of ordinary skill in the art.

31. I have applied the *Phillips* standard in my analysis. Unless otherwise stated, I have applied the plain and ordinary meaning to claim terms.

V. Invalidity Grounds

32. The reasons for my conclusions are explained more fully below.

Grounds	Claims	Statutory Basis	Prior Art
1	1, 3, 4, 6-8, 10, 11, 13-17, 19	§ 103	Berke and Cornelius
2	1, 3, 4, 6-8, 10, 11, 13-17, 19	§ 103	Berke and Ran
3	2, 5, 9, 12, 18	§ 103	Berke, Cornelius, and Stauffer
4	2, 5, 9, 12, 18	§ 103	Berke, Ran, and Stauffer
5	1-19	§ 103	Ran, Mejia, and Stauffer

A. Ground 1

33. It is my opinion that a POSITA would have found Claims 1, 3, 4, 6-8, 10, 11, 13-17, and 19 obvious based on the combination of Berke and Cornelius. Berke teaches techniques in a serial communication system for selecting initial pre-equalizer coefficient values and updating them in a training process, while

Cornelius teaches serial communications between a host and an active cable with a pluggable module. The motivation to combine is explained below. *See below* §V.A.8[8pre].

1. U.S. Patent Application Publication No. 2019/0288881 (“Berke”) (Ex-1004)

34. I have reviewed the Berke reference. Berke was filed on March 19, 2018 and published on September 19, 2019. I understand Berke is prior art to the ’111 patent under §102(a)(2). *See below* Appendix 4 (Understanding of the Law).

35. Berke is directed to “a system and method to optimize equalization coefficients in a high-speed serial interface.” Ex-1004, ¶1. Berke teaches a high-speed serial data interface comprising transceivers having pairs of transmitter and receiver connected via serial channels. Ex-1004 ¶¶1, 12-13, FIG. 1. Berke recognizes a problem that “[a]s the speed of serial interfaces increases, variations in circuit design, component manufacture, environmental conditions, and other factors make it increasingly difficult to ensure highly reliable data transmission.” Ex-1004, ¶3.

36. To address the problem, Berke teaches techniques for configuring a serial channel, including equalization coefficients for a feed-forward equalization (FFE) module on the transmitter side that performs pre-equalization. Ex-1004, ¶¶12-14, FIG. 1. Equalization coefficient settings known to provide robust

performance characteristics are stored in a whitelist that includes a group of registers. Ex-1004, ¶¶12-15, 30, FIG. 1. A set of equalization coefficients in the whitelist is selected based on the performance they provide for a channel. Ex-1004, ¶33. The selected set of equalization coefficients are used as a starting point for adaptive back channel training that determines a satisfactory set of coefficients. Ex-1004, ¶30, ¶13, ¶¶24-26.

37. Berke is in the same field of endeavor as the '111 patent because Berke, like the '111 patent, discloses systems and methods for serial communications, including those comply with serial communication standards including PCIe, SATA, and Gigabit Ethernet. *See* above §II; *e.g.*, Ex-1004, Title (“System and method to optimize equalization coefficients in a high-speed serial interface”), ¶¶1-4, 13, 35, 39. Berke is also reasonably pertinent to the problem with which the inventors of the '111 patent were involved because both Berke and the '111 patent relate to ensuring reliable data transmission through a communication channel. *Compare* Ex-1001, 1:19-35 (“As demand continues for ever-higher data rates, extensions to the standard must account for increased channel attenuation and dispersion even as the equalizers are forced to operate at faster symbol rates.”) *with* Ex-1004, ¶3 (“As the speed of serial interfaces increases, variations in circuit design, component manufacture, environmental

conditions, and other factors make it increasingly difficult to ensure highly reliable data transmission.”).

2. U.S. Patent No. 8,516,238 (“Cornelius”) (Ex-1005)

38. I have reviewed the Cornelius reference. Cornelius was filed on June 30, 2011 and issued on August 20, 2013. I understand Cornelius is prior art to the ’111 patent under §§102(a)(1)-(2). *See* below Appendix 4 (Understanding of the Law).

39. Cornelius is directed to “[c]ircuits, methods, and apparatus that allow signals that are compliant with multiple standards to share a common connector on an electronic device.” Ex-1005, Abstract. Cornelius teaches a configuration where a host device is connected to another device via a high-speed active cable. *Id.*, 4:39-54. The active cable includes active plugs to be connected to a host device. Ex-1005, 6:11-20, 8:20-26, FIGS. 5-6. A connector on the host device couples a transceiver on the active plug to the host device’s transceiver. Ex-1005, 4:22-31, 4:55-65, 6:16-18, 6:21-24, FIGS. 1-3. Cornelius further teaches that “the reliability and accuracy of data connections may be improved by calibrating or training circuitry in the hosts, cables, and other devices,” such as equalization for channel compensation. *Id.*, 11:45-51.

40. Cornelius is in the same field of endeavor as the ’111 patent because Cornelius, like the ’111 patent, discloses devices and methods for serial

communications, including cable connectors for high-speed host-to-device data transmission, such as those comply with serial communication standards including “Universal Serial Bus (USB), Peripheral Component Interconnect Express (PCIe), and Display Port (DP).” *See* above § II; *e.g.*, Ex-1005, Abstract, 1:15-22. Cornelius is also reasonably pertinent to the problem with which the inventors of the ’111 patent were involved because both Cornelius and the ’111 patent relate to ensuring reliable data transmission through a serial communications channel. *Compare* Ex-1001, 1:19-35 *with* Ex-1005, 11:45-62 (“[T]he reliability and accuracy of data connections may be improved by calibrating or training circuitry in the hosts, cables, and other devices.”).

3. Independent Claim 1

[1pre]. A SerDes communications method that comprises, in a transceiver:

41. Ground 1 teaches a **SerDes communications method**. Berke is directed to “[a] high-speed serial data system.” Ex-1004, Abstract. Berke teaches “a high speed serial channel 100 of an information handling system,” where two components communicate over “a bi-directional serial data link[.]” Ex-1004, ¶12. Berke teaches or suggests a common serial architecture where data is serialized in the transmitter, sent across the high-speed serial channel, and then deserialized in the receiver at the other end. Ex-1004, ¶1 (“a system and method to optimize equalization coefficients in a high-speed serial interface”), ¶4 (“High-speed serial

data interface ... may include a transmitter and a receiver.”), ¶12 (“Serial channel 100 includes a transmitter 110, a transmission channel 120, and a receiver 130.”), FIG. 1.




FIG. 1

42. Berke further applies its teachings to standards with SerDes architectures, such as PCIe and SATA. According to Berke, “[a]n example of serial channel 100 includes a PCI-Express (PCIe) channel ... a Serial ATA (SATA) channel ..., or another high speed serial channel.” Ex-1004, ¶12. The PCIe and SATA standards included serialization and deserialization, which a POSITA would have understood given the widespread use of both standards.² See, e.g., Ex-1010 (PCI Express® Base Specification Revision 3.0), 000192-000193; Ex-1011 (Serial

² PCIe and SATA were used in nearly every PC (including Apple computers) by 2018.

ATA Revision 3.1 Specification), 000621 (“A number of state machine, impedance compensation, and serializer/deserializer (SerDes) circuits make up a typical Serial ATA interface Phy.”); *see also* Ex-1008, 000058-61 (listing SATA and PCI Express as “protocol standards for which HSS [High-Speed SerDes] cores are used.”), 000018. Therefore, a POSITA would have understood Berke to be teaching serial communication techniques in the context of common SerDes architectures. Berke teaches or suggests a method for SerDes communications.

43. Likewise, Cornelius teaches PCIe for its active cables (Ex-1005, 4:15-17), citing the PCIe Base Specification. Ex-1005, 000002 (Other Publications), 1:15-22. Moreover, Cornelius uses its active cable to provide high-speed serial connections between computers. Ex-1005, 2:65-3:53. “This figure illustrates computer 110 communicating with computer or display 220 over high-speed connection 225.” *Id.*, 3:32-35, FIGS. 2, 5.





FIGURE 2



Cornelius converts parallel data from each computer into a high-speed serial connection, with a single channel in each direction. *Id.*, FIGS. 2, 5. In this way, Cornelius's architecture funnels data like water from a wide pipe into a narrow one: normal parallel data in a computer is serialized into a narrow, high-speed stream, like a high-speed flow of water in a narrow pipe, and then deserialized at the other end back to its original parallel form. *Id.*, 3:32-35. Cornelius teaches a well-known SerDes architecture that a POSITA would have recognized from its teachings and the context of high-speed serial communications between computers using cables. Therefore, the combination of Berke and Cornelius teaches or suggests a SerDes communications method.

44. **Motivation to Combine.** A POSITA would have been motivated, and found it obvious, to combine Berke and Cornelius's teachings by using Berke's high-speed serial communications in a SerDes context, for example to connect two computers together. Cornelius provides express motivation by teachings this usage for its active cable. Ex-1005, 2:65-3:53, Fig. 2. Additionally, a POSITA would have recognized Berke's teachings to be in the context of SerDes communications, given Berke's overall teachings and the PCIe/SATA context as explained above. Ex-1004, ¶12. Berke's teachings suggest serialization and deserialization because it teaches techniques that enable high-speed communications over a channel. These types of serial communications channels presume serialization to force data into a high-speed stream and then deserialization to restore the original data, similar to flow through pipes as explained above, which a POSITA would have recognized when looking at Berke in the context of Cornelius's active cable teachings. Because computer chips generally use parallel data internally—for example using 8 parallel lines to move 8 bits in and out of an 8-bit register—the combination of Berke and Cornelius teaches or suggests serialization of parallel data in one computer, sent across the high-speed channel, and then deserialization on the other side to restore the original parallel data in the recipient computer. A POSITA would have understood Berke's serial channel teachings to be provided in the context of serialization on the transmitter side (where the transmitting chip

converts parallel data into serial format for transmission over the channel) and deserialization at the receiver side (where the receiving chip converts the serialized data back into parallel format). This routine process took advantage of high-speed communications with reduced pin count and easier routing. *See* above §I. A POSITA would have known this because that is how serial communications channels had been used for decades, and serial communications were fundamentally associated with SerDes well before 2018.

45. Berke further teaches a **transceiver** performing operations of a communications method. *See* below §V.A.3[1a-1c]. As shown in Figure 1, Berke’s “[s]erial channel 100 includes a transmitter 110, a transmission channel 120, and a receiver 130.”





FIG. 1

Ex-1004, ¶12, FIG. 1.

46. Berke next explains that Figure 1 illustrates “one half” of a bi-directional data link, with the “other half” including a similar channel going the other direction, with the receiver on the left and transmitter on the right:

Serial channel 100 represents one half of a bi-directional serial data link for communicating data from transmitter 110 located at a first component to receiver 130 located at a second component. The other half of the bi-directional serial data link is similar to serial channel 100, but with a receiver in the first component, and a transmitter in the second component, for communicating data back from the second component to the first component.

Ex-1004, ¶12. Therefore, each side of the channel includes a transmitter and receiver. *Id.* Berke thus teaches a transceiver at each end of the communication link, the transceiver comprising a transmitter and a receiver and coupled to two unidirectional channels. I have illustrated Berke’s teachings below:



[1a] selecting one of multiple registers to specify initial pre-equalizer coefficient values, each of the multiple registers corresponding to a different channel model;
--

47. Berke teaches **pre-equalizer coefficient values**. For example, Berke's transceiver performs pre-equalization using coefficient values: the transmitter 110 includes "a feed-forward equalization (FFE) module 114," which "operates to proactively provide compensation to a transmitted signal by pre-distorting or shaping the data over several bit periods in order to invert the channel loss/distortion introduced by transmission channel 120." Ex-1004, ¶¶14-15. Berke's FFE module performs pre-equalization because it is configured on the transmitter side and it provides compensation to a signal before it is transmitted to the receiver. *Id.* Also, by "pre-distorting" data to "invert the channel loss/distortion introduced by transmission channel 120," the FFE module compensates for at least some of the channel loss before the data is transmitted via the channel to reduce or eliminate the needs for compensation at the receiver. *Id.* Berke teaches that "transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values within the transmitter and the receiver to compensate for the insertion loss and other signal degradations of transmission channel 120." Ex-1004, ¶13. Here, Berke's compensation values refer to equalization coefficients. For example, Berke discusses "select[ing] a

compensation value from among a plurality of compensation values” and “stor[ing] a whitelist value from among the compensation values” (Ex-1004, ¶4) and later explains that this involves “select[ing]” a “set of equalization coefficients” (Ex-1004, ¶23) and “equalization coefficient whitelists” (Ex-1004, ¶30). Ex-1004, ¶22 (teaching “configur[ing]” serial channel “with the equalization coefficients”). Those skilled in the art would have understood that equalization compensates for “insertion loss and other signal degradations” and are configured using equalization coefficients. Ex-1004, ¶13. The compensation values or equalization coefficients include those that configure Berke’s FFE. Ex-1004, ¶15 (“The amount of compensation is determined by enabling a number of circuit feed-forward taps.”); *see also* Ex-1008, 000020-000021, 000050. Therefore, Berke teaches that the pre-equalization (e.g., FFE) is performed based on pre-equalizer coefficient values (e.g., referred to as “compensation values” or “equalization coefficients”). A POSITA would have understood Berke’s equalization coefficients that are used to configure the transmitter’s pre-equalization process to be pre-equalizer coefficient values.

48. Berke teaches **selecting one of multiple registers** (e.g., two or more of the registers that store pre-equalizer coefficients in the whitelist) **to specify initial pre-equalizer coefficient values**. Berke teaches that the “[t]ransmitter 110 includes ... an equalization coefficient whitelist 118.” Ex-1004, ¶14, FIG. 1.

“[C]ertain equalization coefficient settings or combinations of equalization coefficient settings are known to produce one or more of a particularly robust eye height or eye width, and such known coefficient settings or combinations of coefficient settings are stored in equalization coefficient whitelists.” Ex-1004, ¶30, ¶32 (“[C]ontrol logic module 140 determines a whitelist set of equalization coefficients based upon the characterizations, and stores the whitelist set of equalization coefficients to equalization coefficient whitelists 118 and 144.”).

49. Berke teaches implementing the whitelists with different types of storage media, among which is “a register or group of registers of transmitter 110 and receiver 130.” Ex-1004, ¶30. Because Berke teaches storing multiple sets of whitelisted coefficient settings (Ex-1004, ¶33) in a group of registers (Ex-1004, ¶30), Berke teaches or suggests using each of the registers in the group to store one set of pre-equalizer coefficients. This was a simple and reliable way to organize the sets of coefficients taught by Berke. Moreover, Berke teaches using an addressable register design, where each register is structured in a certain bit width (e.g. 8-bit register, 16-bit register, 32-bit register, or higher), and all the registers are addressed in order to be properly accessed, e.g. write and read. The total number of registers and their bit width depend on the needs of the system. Typical protocols to access such addressable registers included Inter-Integrated Circuit (I²C) interface and System Packet Interface (SPI) among others, as Berke

mentions. Ex-1004, ¶36. In view of the typical width of multi-bit registers in transceiver circuits, storing a set of three or more FFE tap coefficients per register was a straightforward and efficient design choice; this would have been readily apparent to a POSITA with experience in serial communications. I²C and SPI were popular protocols to access local registers such as those located in transceivers, which allowed reading or writing one addressed multi-bit register in one clock cycle. In fact, this practice—storing a set of equalization coefficients in one register so that they can be accessed in one clock cycle—was a long-standing practice in the field. *See, e.g.*, Ex-1030, 1:10-25 (“FIR filters used in channel equalizers”), 2:23-51 (“FIR filter 10 comprises four multipliers 20, 22, 24 and 26... Input data source 28 may provide a sequence of 8-bit data words to each of the multipliers at the rate of a clock signal CLK. A coefficient register 14 supplies a coefficient C3 to multiplier 20, a coefficient C2 to multiplier 22, a coefficient C1, to multiplier 24 and a coefficient C0 to multiplier 26. Each multiplier 20, 22, 24 and 26 multiplies a current input data word with its respective coefficient C3, C2, C1 and C0...”), FIG. 1.

50. Berke teaches two options for implementing the whitelist: use registers or memory. Ex-1004, ¶30. A POSITA would have been motivated and found it obvious to use registers in the transmitter because it was simple and effective. As explained above, using registers would have allowed each set of

coefficients to be stored in a multi-bit register, meaning the set of coefficients would be selected by simply accessing that register and loading the coefficients stored inside. This would have been advantageous because it allowed a set of pre-equalizer coefficients to be loaded in one clock cycle, with a single register read, as compared to more complicated schemes involving reads to registers or memory or other operations just to select a set of pre-equalizer coefficients. Moreover, the use of registers for this purpose—with each register storing a set of coefficients—was a long-standing practice, and the prior art provides express motivation to do this. Ex-1030, 1:10-25, 2:23-51. In addition to the advantage in terms of clock-cycle efficiency, this register-based approach simplified the control logic for selecting and accessing coefficients, streamlined the configuration process, saved power, and reduced overhead. For example, it allowed control logic to uniquely identify a set of whitelisted pre-equalizer coefficients using one identifier or address corresponding to the register that stores the set of coefficients. A POSITA would therefore have been motivated to apply Berke's register teachings in this manner.

51. A POSITA would have had a reasonable expectation of success given the simplicity of a register-based approach. Reading and writing to registers was a basic undergraduate skill. Since implementations of well-known protocols such as I²C and SPI allowed designers to specify their preferred bit width and total number of registers, it would have been obvious to configure registers designed for storing

pre-equalizer coefficients such that each register store one set of such coefficients. Accordingly, using multiple registers—each holding a set of pre-equalizer coefficients—not only falls squarely within the teachings of Berke (Ex-1004, ¶¶30-33) but also represents an obvious implementation that a POSITA would have been motivated to use. This would have applied a known technique to a known device ready for improvement to yield predictable results (a register-based whitelist, as Berke teaches). *See* Ex-1004, ¶30.

52. Berke further teaches selecting one of the multiple registers to specify initial pre-equalizer coefficient values. Berke selects a whitelist set of equalization coefficients from multiple sets by configuring the serial channel with each set of coefficients for training and monitoring the performance of the channel, until a whitelist set of equalization coefficients that satisfies a desired performance level is selected:

[A]fter a first whitelist set of equalization coefficients is utilized for training serial channel 100, control logic module 140 operates to monitor the performance of the serial channel. ... If the performance of serial channel 100 falls below the desired performance level, then control logic module 140 selects a second whitelist set of equalization coefficients from equalization coefficient whitelists 118 and 144 and utilizes the second whitelist set of equalization coefficients to retrain the serial channel.

Ex-1004, ¶33.³ The pre-equalizer coefficients stored in the whitelist at the transmitter are initial pre-equalizer coefficient values because they are used as the starting point for the training process. Ex-1004, ¶30 (“[I]n running the training algorithm for the serial interface, control logic module 140 utilizes the coefficient settings or combinations of coefficient settings stored in equalization coefficient whitelist[] 116 [sic] ... as a starting point for running the training algorithm.”).

³ While Berke uses language such as “[i]n a particular embodiment” and “[i]n another embodiment” to refer to different parts of its disclosure, a POSITA would have understood that these disclosures of Berke are intended to be implemented together in a flexible manner. For example, paragraph 30’s disclosure of storing equalization coefficient settings in registers and paragraph 33’s disclosure of selecting sets of equalization coefficients for training are directed to different aspects of Berke’s system and are compatible to be used together. Berke confirms the flexible applicability of its disclosure, stating that “[a]lthough only a few exemplary embodiments have been described in detail herein, those skilled in the art will readily appreciate that many modifications are possible in the exemplary embodiments without materially departing from the novel teachings and advantages of the embodiments of the present disclosure.” Ex-1004, ¶40.

Because each set of coefficients is stored in a register (explained above), Berke teaches selecting the register with that set of coefficients.

53. Berke teaches **each of the multiple registers corresponding to a different channel model**. Berke stores sets of pre-equalizer coefficients that are known to provide “robust” performance for respective channel characteristics, so that one of them can be selected to serve as a good starting point for the equalizer training process. Ex-1004, ¶¶30-33. The whitelist registers contain coefficients used “to compensate for channel loss.” Ex-1004, ¶3. In particular, a set of coefficients (for multiple taps) are used by the FFE module to “proactively provide compensation to a transmitted signal by ... invert[ing] the channel loss/distortion introduced by transmission channel 120.” Ex-1004, ¶15. As a POSITA would have understood, variations in channel loss or distortion arise from differences in channel characteristics—such as design, settings, and environmental conditions—where different channel characteristics correspond to different channel models.

54. Berke determines equalization coefficients (which, as explained above, include pre-equalizer coefficient values) to include in its whitelist based on different channel characteristics. For example, “the whitelisted coefficient settings can be determined based upon knowledge gained during the design, development, or manufacturing of serial channel 100.” Ex-1004, ¶30. Such knowledge would have informed one about “variations in circuit design, component manufacture”

(Ex-1004, ¶3), which would have included characteristics of circuit traces, connectors, and cables (e.g., length of a trace) (Ex-1004, ¶12), characteristics regarding the channel's compliance with standards (e.g., PCIe, SATA) (Ex-1004, ¶12), characteristics such as "insertion loss and other signal degradations" (Ex-1004, ¶13), and channel characteristics such as "impedance ... based upon a design target for [the] transmission channel" (Ex-1004, ¶19).

55. In addition, whitelisted equalization coefficients can also be "determined based upon knowledge gained during normal operation of serial channel." Ex-1004, ¶31. Such knowledge would have informed one about settings that were used during operation such as receiver equalization settings (Ex-1004, ¶16, ¶18; *see* below §V.A.4) and environmental conditions such as temperature and noise levels (Ex-1004, ¶21), all of which are channel characteristics that impact channel loss or distortion. For example, because transmitter pre-equalization and receiver equalization were conventionally jointly co-designed to compensate for the overall channel loss (*see* above §I.A), the receiver equalization settings impact the overall channel performance and inform the decision as to what pre-equalizer coefficients to use.

56. Each of the aforementioned channel characteristics impact the channel loss and frequency response characteristics. Ex-1004, ¶3, ¶¶12-13, ¶¶16-21. As Berke includes a set of pre-equalizer coefficients in the whitelist based on

knowledge about at least some of these channel characteristics (Ex-1004, ¶¶30-31), this set of whitelisted pre-equalizer coefficients corresponds to a channel model that represents or assumes the known and relied-upon channel characteristics. Berke's multiple sets of whitelisted pre-equalizer coefficient values (Ex-1004, ¶33) therefore are designed to compensate for different channel loss/distortion properties resulting from channel characteristics that correspond to different channel models. This includes, for example, channel models with different amounts of insertion loss or other signal degradations, different trace lengths, connectors, impedance targets, temperatures, types and amounts of receiver equalization, and high noise versus low noise. Ex-1004, ¶¶12-13, ¶¶16-21.

57. To account for these different channel models, Berke selects between different sets of pre-equalizer coefficients from the whitelist. Ex-1004, ¶33. As explained previously in this section, Berke teaches or suggests storing different sets of pre-equalizer coefficients in different multi-bit registers at the transmitter side, where each register stores one set of coefficients. By selecting a register storing a suitable set of pre-equalizer coefficients, Berke starts its training process with a set of pre-equalizer coefficients that meets a desired performance level. Ex-1004, ¶¶30-33. Therefore, Berke teaches or suggests that each of the multiple registers with FFE coefficients in the whitelist correspond to a different channel model.

[1b] updating the initial pre-equalizer coefficient values during a training phase; and
--

58. Berke teaches a training process that updates the pre-equalizer coefficient values. The training process adjusts compensation values within the transmitter, including the pre-equalizer coefficients, via communications between the transmitter and receiver. Ex-1004, ¶13 (“Serial channel 100 operates to provide back channel adaptation where transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values within the transmitter and the receiver to compensate for the insertion loss and other signal degradations of transmission channel 120.”).

59. Berke describes a training phase during which variations of pre-equalizer coefficients of the transmitter (e.g., equalization setting, FFE tap settings) are run and performance characteristics (e.g., eye height and eye width, BER) of data transmission are recorded:

[C]ontrol logic module 140 trains transmitter 110 and receiver 130 by running through the various combinations of settings for the equalization setting, the gain settings, the DFE taps settings, and the FFE tap settings, and recording for each combination, an associated eye height and eye width, or an associated BER, as needed or desired.

Ex-1004, ¶20.

60. For the training process, Berke teaches using a “best-effort” approach to update pre-equalizer coefficient values. Under this approach, “transmitter and receiver equalization mechanisms to compensate for channel loss are calibrated on a best-effort basis, where settings that result in a ‘good enough’ compensation solution are quickly obtained, in favor of iterative processes that might yield a more optimal solution but which require an inordinate amount of time for such link training.” Ex-1004, ¶3; *see also* ¶¶24-29, FIG. 3.

61. To do so, first, “control logic module 140 operates to determine if the variability in equalization coefficients derived over different runs of the training algorithm is greater than an equalization variation threshold.” Ex-1004, ¶24; *see also* ¶¶27-28. If so, a new set of training runs are commenced to pick out a small number of variables in equalization coefficients⁴ (i.e., a narrowed set) to which the performance of serial channel is most sensitive:

If the variability is greater than the equalization variation threshold, then control logic module 140 commences a set of training runs that are targeted to determining a small number of variables in the equalization coefficients that, when altered from one setting to another, produce a

⁴ As explained above, Berke’s equalization coefficients that are used to configure the transmitter’s pre-equalization process are pre-equalization coefficient values. *See* above §V.A.3[1a].

wide variance in the measured performance of serial channel 100, in terms of eye height and width, or in terms of BER.

Id., ¶24; *see also* ¶¶25, 28, FIG. 3 (“318 Determine Sensitive Coefficients,” “320 Tighten Training Range For Sensitive Coefficients”). After the narrowed set of variables (“knobs”/“coefficients”) are determined, “control logic module 140 determines a narrowed set of knob settings for each knob, the set of knob settings being associated with the best channel performance, and then, in future runs of the training algorithm, control logic module 140 only utilizes the narrowed set of knob settings for the selected knobs to converge on a best set of equalization settings.”


Id., ¶26. “[I]f the variability in equalization coefficients derived over different runs of the training algorithm with the narrowed set of settings for the selected knobs is still greater than the equalization variation threshold, the control logic module then commences another set of training runs.” *Id.*; *see also* ¶29. The training concludes once it is determined that the variability is not greater than the equalization variation threshold, and “the set of equalization coefficients associated with the best performance of serial channel 100, that is, the set of equalization coefficients that resulted in the highest eye height and eye width, or that resulted in the lowest BER, is selected.” *Id.*, ¶24.

62. As explained above, Berke teaches using the whitelist set of equalization coefficients “as a starting point for running the training algorithm.”

Ex-1004, ¶30. Accordingly, Berke teaches a training phase that updates the initial pre-equalizer coefficients to a converged set of pre-equalizer coefficients.

[1c] using the updated pre-equalizer coefficient values to convey a transmit data stream.

63. Berke discloses a training process in which “[s]erial channel 100 operates to provide back channel adaptation where transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values within the transmitter and the receiver to compensate for the insertion loss and other signal degradations of transmission channel 120.” Ex-1004, ¶13. Here, “adaptation” refers to dynamically adjusting transmission parameters based on “test runs” or training. As demonstrated in Figure 1, in the transmitter 110, the FFE module 114 uses the pre-equalizer coefficients to process a serial data stream, outputs the pre-equalized signals to a channel output module 112, which in turn outputs the transmit data stream for transmission on channel 120 to be received by receiver 130. *Id.*, FIG. 1, ¶14.



100

FIG. 1

64. As discussed for limitation [1b] above, Berke teaches a training phase during which the initial equalizer coefficients are updated to a converged set of equalization coefficients with desirable performance. *See* above §V.A.3[1b]. Berke teaches subsequently using the updated pre-equalizer coefficients to convey a data stream:

[T]he set of equalization coefficients associated with the best performance of serial channel 100, that is, the set of equalization coefficients that resulted in the highest eye height and eye width, or that resulted in the lowest BER, is selected, and normal data transmission operation is commenced on serial channel 100.

Ex-1004, ¶24.

4. Dependent Claim 3

3. The method of claim 1, wherein at least some of the different channel models presume different types of receiver equalization.

65. Ground 1 renders obvious the method of claim 1. *See* above §V.A.3.

Ground 1 further renders obvious the additional limitation of claim 3.

66. Berke teaches **different types of receiver equalization**. Berke's receiver 130 includes a "continuous time linear equalization (CTLE) module 134" and a "decision feedback equalization (DFE) module 138." Ex-1004, ¶16. The CTLE and DFE modules provide different types of receiver equalization. *Compare id.*, ¶16 ("[T]he CTLE module operates to provide compensation for inter-signal interference (ISI) in order to open the signal eye of the received signal.") *with* ¶18 ("DFE module 138 operates to provide feedback based compensation to the received signal."). Berke teaches that the CTLE module supports different amount of equalization, or alternatively, could be turned off. *Id.*, ¶16. ("[R]eceiver 130 can support 21 CTLE equalization settings which each prescribe a different amount of equalization, from 0 dB to 10 dB, in 0.5 dB steps."). For the CTLE module, "[t]he amount of compensation is determined based upon an equalization setting." *Id.*, ¶16. Similarly, the DFE also supports different amount of equalization, or alternatively, could be turned off. *Id.*, ¶18 ("DFE module 138 can support up to 16 taps that provide compensation based upon up to 16 previous data points. ...DFE module 138 can be turned off. ... [O]ne or more tap of DFE module 138 can be

turned on based upon the taps setting.”). For the DFE module, “[t]he amount of compensation is determined by enabling a number of circuit feedback taps.” Ex-1004, ¶18. Accordingly, Berke teaches different types of receiver equalization, corresponding to using CTLE, DFE, or both for receiver equalization. Furthermore, the use of CTLE and/or DFE at different settings and/or taps constitutes different types of receiver equalization.

67. As explained above, Berke teaches a whitelist with registers that store sets of pre-equalizer coefficients, each of which corresponds to a channel model of channel characteristics such as the receiver equalization. *See* above §V.A.3[1a]. Berke teaches that the transmitter communicates with the receiver to optimize compensation values for the communication channel. Ex-1004, ¶13 (“Serial channel 100 operates to provide back channel adaptation where transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values within the transmitter and the receiver to compensate for the insertion loss and other signal degradations of transmission channel 120.”). The type of receiver equalization used impacts the channel performance. Ex-1004, ¶29 (“[C]hanges to the CTLE settings may result in widely varying performance on the serial channel.”). Because Berke teaches jointly deciding pre-equalization settings on the transmitter side and the equalization settings on the receiver side, the transmitter-side pre-equalization settings are a function of the equalization type on

the receiver side. In other words, for a set of equalization settings for pre-equalization to achieve high performance (e.g., robust eye heights or eye widths, low BER) during operation of a serial channel, they must be suitable for the receiver equalization type. *See* above §V.A.3[1a]. For example, when less compensation is available at the receiver, greater compensation is required at the transmitter, and vice versa. Therefore, Berke teaches or suggests that **at least some of the different channel models presume** different types of receiver equalization.

5. Dependent Claim 4

4. The method of claim 1, wherein said selecting includes using the initial pre-equalizer coefficient values to determine a performance characteristic for each of the multiple registers.

68. Ground 1 renders obvious the method of claim 1. *See* above §V.A.3.

Ground 1 further teaches the additional limitation of claim 4.

69. As explained with respect to limitation [1a], Berke selects a set of equalization coefficients to use for training based on performance characteristics, where equalization coefficients that meets a desired performance level are selected for use in training. *See* above §V.A.3[1a]; Ex-1004, ¶33. Berke's equalization coefficients that are used to configure the transmitter's pre-equalization process are pre-equalizer coefficient values. *See* above §V.A.3[1a]. Berke explains configuring the channel using each of multiple sets of whitelisted equalization coefficients (e.g., "a first whitelist set of equalization coefficients," "a second whitelist set of

equalization coefficients”) and evaluating its performance level using a “receiver eye diagram” with “eye height and eye width” and/or “a BER of serial channel 100.” Ex-1004, ¶33 (“For example, control logic module 140 can monitor the receiver eye diagrams to monitor that the eye height and eye width remain within specified thresholds, or can monitor that a BER of serial channel 100 remains below a specific BER threshold.”); *see also* ¶¶13, 20, 24-26, 31. Eye diagram and BER are characteristics commonly used to evaluate equalization performance and thus constitute **performance characteristics** for pre-equalizer coefficient settings. *See* above §I; *see also* Ex-1008, 000014-000015.

70. Berke teaches monitoring performance characteristics (e.g., eye height, eye width, BER) when each of multiple sets of whitelisted equalization coefficients are utilized to configure the channel for training purposes. Ex-1004, ¶33. Since the whitelisted equalization coefficients include initial pre-equalizer coefficient values, Berke **uses the initial pre-equalizer coefficient values to determine a performance characteristic**. As explained above, Berke teaches or suggests using different registers to each store a set of pre-equalizer coefficients. *See* above §V.A.3[1a]. It follows that the performance characteristics determined using each set of equalization coefficients are determined for the corresponding register. Berke thus teaches using the initial pre-equalizer coefficient values to determine a performance characteristic **for each of multiple registers**.

6. Dependent Claim 6

6. The method of claim 4, wherein the performance characteristic is a bit error rate.

71. Ground 1 renders obvious the method of claim 4. *See* above §V.A.5. As explained with respect to claim 4, Berke teaches using the bit error rate (BER) as a “performance characteristic.” *See* above §V.A.5. Therefore, Ground 1 teaches claim 6.

7. Dependent Claim 7

7. The method of claim 4, further comprising updating the selected register with the updated pre-equalizer coefficient values.

72. Ground 1 renders obvious the method of claim 4. *See* above §V.A.5. Ground 1 further teaches the additional limitation of claim 7.

73. As explained with respect to limitation [1b], Berke teaches updating the initial pre-equalizer coefficient values from those stored in the selected register of the whitelist to updated pre-equalizer coefficient values during a training phase. *See* above §V.A.3[1b]. Further, Berke teaches storing the results of the training process in the whitelist:

In another embodiment, control logic module 140 operates to providing narrowed equalization coefficients for training serial channel 100 in accordance with the method shown in FIG. 3. ... Then, when all of the knobs and settings have been characterized, control logic module 140 determines a whitelist set of equalization coefficients based upon the

characterizations, and stores the whitelist set of equalization coefficients to equalization coefficient whitelists 118 and 144.

Ex-1004, ¶32.⁵ A POSITA would have found it obvious and been motivated to replace the initial pre-equalizer coefficient values that were used as the starting point of the training process with the updated pre-equalizer coefficient values resulting from the training process. Berke teaches that “coefficient settings or combinations of coefficient settings that result in particularly robust eye heights or eye widths can be determined based upon knowledge gained during normal operation of serial channel 100.” *Id.*, ¶31. A POSITA would have recognized that the updated pre-equalizer coefficients obtained during the training phase represent knowledge that is gained during normal operation and provide better performance than the starting point of the training. This would have motivated a POSITA to use the updated pre-equalizer coefficients to replace the initial pre-equalizer

⁵ As explained in footnote 3, while Berke uses language such as “[i]n a particular embodiment” and “[i]n another embodiment” to refer to different parts of its disclosure, a POSITA would have understood that these disclosures of Berke are intended to be implemented together in a flexible manner. *See* above §V.A.3[1a] n.3.

coefficients by updating the selected register with the updated pre-equalizer coefficient values.

8. Independent Claim 8

[8pre] A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

74. Cornelius teaches a **chip-to-module communications link** between a host and a module in an active cable, to which a POSITA would have found it obvious and been motivated to apply Berke's equalization techniques. Berke applies its teachings broadly to serial communications over circuit traces on printed circuit boards, connectors, and/or transmission cables. Ex-1004, ¶12. Cornelius applies serial communications techniques to active cables and teaches a path from a host to either end of a cable that includes a module, which is a chip-to-module communications link. Ex-1005, 11:59-62. The host (e.g., a computer) communicates with an endpoint (e.g., an external device such as a display or disk drive) via the active cable. Ex-1005, 5:33-45 ("In line 450, an HSIO source or host communicates with an HSIO sink or endpoint. In this configuration, an active cable is required between the HSIO source or host and the HSIO sink or endpoint."), FIGS. 1-2:




FIGURE 2

Cornelius contemplates communicating data signals according to various standards, including, for example, PCIe and USB. Ex-1005, 1:15-22, 4:55-65. The chip-to-module communications link is further explained below.

75. Cornelius's host has **a port connector**, e.g., a “connector on computer 110” (Ex-1005, 3:43-53, 1:15-22, 4:55-65, FIGS. 1-3) that is connected to **a port transceiver**, as evidenced by its pins for transmitting and receiving high-speed signals, such as “[h]igh-speed transmit pins 3 and 5” and “high-speed receive pins 4 and 6.” *Id.*, 4:22-31, FIG. 3. The circuitry that transmits the Tx signals and receives the Rx signals is the transceiver. *Id.*

76. Cornelius's active cable comprises two **pluggable modules** that are “two active plugs 500 and 505, one on each end of cable 507.” Ex-1005, 6:11-20, FIG. 5; *see also* 8:20-26, FIG. 6. Each active plug includes a transceiver (e.g., dual clock and data recovery (CDR) circuits). Ex-1005, 6:11-20 (“Each active plug includes dual clock and data recovery circuits.”), 6:25-29 (“The clock and data recovery circuits ... may include optical receivers and transmitters.”). Cornelius's CDR circuits perform a range of functions, including clock and data recovery, retiming data, performing equalization and pre-equalization, and transmitting and

receiving data. Ex-1005, 6:11-20 (“Each active plug includes dual clock and data recovery circuits for retiming data. Specifically, active plug 500 provides high-speed transmission signals on pins 3 and 5, and receives high-speed signals on pins 4 and 6.”), 6:21-24, 6:30-42 (“[T]he clock and data recovery circuits may employ equalizer circuits, buffers, emphasis, and de-emphasis circuits as appropriate.”), FIGS. 5-6. Therefore, the dual CDR circuitry in each plug is a transceiver because it receives a signal and transmits a signal; this applies to an individual CDR component, which receives on one side and transmits on the other, and also to the dual CDR circuitry as a whole, which transmits and receives on each side. *Id.*, 6:11-29.⁶

⁶ To the extent Patent Owner argues that the transmitter and receiver must be a single IC or package, Cornelius already teaches that the transmitter and receiver functions of the dual CDR/retimer circuitry are combined physically into a single active plug. 6:11-29, 8:20-26, FIGS. 5-6. Integrating them onto a single IC would have been an obvious implementation detail because cable plugs were small form factors, and it would have made sense to integrate the transceiver functions onto a single IC to fit into Cornelius’s active plug. The integration of transmitting and receiving circuitry into a single unit was a well-known and routine design choice by 2018. *See, e.g.*, Ex-1020, cl.1.

77. **The port connector**, through its transmit and receive pins, **couples the port transceiver to the pluggable module transceiver**. Ex-1005, 1:15-22 (“Electronic devices often include connectors to provide ports where power and data signals can be shared with other devices.”), 3:44-53, 6:16-18 (“Specifically, active plug 500 provides high-speed transmission signals on pins 3 and 5, and receives high-speed signals on pins 4 and 6.”), 6:21-24. The connector’s transmit and receive pins couple the transceiver in Cornelius’s active plug with the host chip that transmits and receives those signals. *Id.* Therefore, when Cornelius’s active cable is plugged in the port connector on a host (e.g., a computer), it forms a **chip-to-module communications link** interconnecting a chip in the host and a pluggable module.

78. **Motivation to Combine**. Berke teaches techniques for serial communications that apply to a wide variety of channels including circuit traces on printed circuit boards, connectors, and/or transmission cables. Ex-1004, ¶12. Cornelius teaches an active cable with three serial communication segments: (1) the path from the host chip to the transceiver at the near end of the cable; (2) the path through the cable; and (3) the path from the far end of the cable to a device or other host chip. Ex-1005, 11:59-62. These paths were precisely the types of board traces and cables that Berke applied its teachings to. Therefore, a POSITA would have been motivated to apply Berke’s teachings to the serial communication

paths for Cornelius's active cable because it would have simply used Berke's known serial techniques to improve similar devices (e.g., Cornelius's serial channels) in the same way as taught by Berke.

79. Active cables retime signals, which improves the signal and reduces jitter, thereby enabling longer cable lengths. Active cables were widely known by 2018, and the prior art provides motivation to use them. *See* Ex-1025 (datasheet for TE SFP+ cables, supporting 1 and 10 Gigabit Ethernet); Ex-1026, 000001, 000005, 000008-000009 (various Cisco transceiver and cable modules supporting 40 Gigabit Ethernet). In view of the commercial use of numerous types of active cables, a POSITA would have been motivated to apply Cornelius's active cable teachings to connect computers to each other and external devices, as also taught by Berke.

80. Furthermore, Cornelius teaches pre-equalization for its communications link. *See, e.g.,* Ex-1005, 6:30-42, 6:50-57, 7:7-17, 11:45-62. Berke's techniques for configuring pre-equalizer coefficient values provided an efficient and effective process for configuring a serial communications channel to compensate for the insertion loss and other signal degradations. Ex-1004, ¶13. Berke's best-effort link training techniques were advantageous in that it obtains satisfactory equalization configurations in a shorter amount of time than systems that exhaust all possible configurations. *See* above §V.A.3[1b]; Ex-1004, ¶3.

Applying Berke's techniques to Cornelius would have provided such benefits and help Cornelius achieve its goal of providing the desired equalization for its communications link.

81. Additionally, the teachings of Berke and Cornelius were readily compatible with each other. First, both are in the same field of serial communications. *See* above §§V.A.1-2. Both apply their teachings to circuit boards and cables, for computers and external devices. *See, e.g.*, Ex-1004, ¶12 (“The components can also be understood to include devices of an information handling system, such as a hard drive, a storage array, and the like, that are separate from the printed circuit board of the information handling system, where transmission channel 120 can include one or more transmission cables.”), ¶37; Ex-1005, FIGS. 1-2. Both references teach support for common standards, including PCIe and USB. Ex-1004, ¶¶12, 37; Ex-1005, 1:15-22, 4:55-65. The references apply their teachings to similar devices. Berke teaches serial communications links between a computer and external devices such as various input and output (I/O) devices including, for example, video display. Ex-1004, ¶34. Likewise, Cornelius teaches an active cable for connecting a computer with a display device. Ex-1005, 3:17-23. The combination would not have changed the principle of operation for either reference because it uses Berke's serial communications techniques for exactly the purpose taught by Berke: to improve pre-equalization for serial

communications. In addition to the reasons explained above, a POSITA would have had a reasonable expectation of success in the combination because it would have been a straight-forward combination that simply applies Berke's serial communications techniques to Cornelius's serial links. As explained above, the teachings were readily compatible and used for the same purposes taught in the references themselves. Active cables, pre-equalization, and serial communications had been widely known and used for many years before 2018. *See* above §I.

[8a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

82. As explained with respect to limitation [1a], Berke teaches a “feed-forward equalization (FFE) module 114” that performs pre-equalization in a transceiver based on pre-equalizer coefficient values. *See* above §V.A.3[1a]. As serial channels transmit data symbols sequentially, the data going through a serial channel is a serial symbol stream. Berke's FFE module is configured at the transmitter and pre-equalizes a serial symbol stream to be transmitted on the chip-to-module communications link. *Id.* This teaches a transmit filter because a feed-forward equalizer is a type of FIR filter commonly used in serial communications. *See, e.g.,* Ex-1008, 000020 (referring to a three-tap FFE as “the filter” and each tap's tap weight value as “a filter coefficient”). Therefore, Berke

teaches **one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted.**

83. Cornelius teaches performing pre-equalization at the transceiver of its pluggable module (e.g., dual CDR circuits). Ex-1005, 6:11-20 (“Each active plug includes dual clock and data recovery circuits.”), 6:30-42 (“[T]he clock and data recovery circuits may employ equalizer circuits, buffers, emphasis, and de-emphasis circuits as appropriate.”), 11:45-62. As explained above, a POSITA would have been motivated to apply Berke’s pre-equalization techniques to Cornelius’s system. *See* above §V.A.8[8pre] (motivation subsection explaining how and why Berke and Cornelius would have been combined). Under this combination, FFE functionalities, as taught by Berke, are implemented in Cornelius’s pluggable module transceiver to pre-equalize a corresponding serial symbol stream being transmitted **to the port transceiver.**

[8b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.
--

84. As explained above with respect to limitation [1a], Berke teaches multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model. *See* above §V.A.3[1a]. In addition, Berke teaches selecting one of the multiple registers to

specify initial pre-equalizer coefficient values for the FFE module. *Id.* In this way, the selected register is used to set initial coefficient values for the one or more transmit filters. *See* above §V.A.8[8a] (explaining that the FFE module is a transmit filter).

85. The combination of Berke and Cornelius further teaches **a controller having the multiple registers and using one of them to set initial coefficient values**. Berke teaches whitelist and channel management functionalities on the transmitter side (e.g., the functionalities of the channel management module and the equalization coefficient whitelist) that control the selection and training of pre-equalizer coefficients. Ex-1004, ¶14 (“Transmitter 110 includes a channel output module 112, feed-forward equalization (FFE) module 114, **a channel management module 116**, and **an equalization coefficient whitelist 118**.”), ¶19, ¶33, FIG. 1. The components of Berke’s transceiver that implement these functionalities on the transmitter side constitute a controller.

86. Berke’s controller controls the pre-equalizer settings on the transmitter side based on instructions from the receiver side. The control logic module 140, configured on the receiver side, selects a whitelist set of pre-equalizer coefficient values and its corresponding register by determining their corresponding performance characteristics:

[A]fter a first whitelist set of equalization coefficients is utilized for training serial channel 100, control logic module 140 operates to monitor the performance of the serial channel. ... If the performance of serial channel 100 falls below the desired performance level, then control logic module 140 selects a second whitelist set of equalization coefficients from equalization coefficient whitelists 118 and 144.

Ex-1004, ¶33; *see above* §V.A.3[1a].

87. The control logic module 140 communicates adaptation instructions to channel management module 116 that is on the transmitter side:

In a particular embodiment, management module 116 and channel management module 142 represent a separate side-band communication channel for communicating adaptation instructions between transmitter 110 and receiver 130.

Ex-1004, ¶19. Pre-equalization on the transmitter side is configured based on the adaptation instructions. Ex-1004, ¶13 (“Serial channel 100 operates to provide back channel adaptation where transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values within the transmitter and the receiver to compensate for the insertion loss and other signal degradations of transmission channel 120.”). A POSITA would have understood the adaptation instructions include information indicating the selected set of whitelist pre-equalizer coefficient values since Berke teaches “the back channel adaptation mechanism may operate to provide different sets of compensation

values,” which, as explained above, include pre-equalizer coefficient values. *Id.*, ¶¶13, 20; *see above* §V.A.3[1a]. The selected set of whitelist pre-equalizer coefficient values are stored in the equalization coefficient whitelist 118. Ex-1004, ¶30, ¶32; *see above* §V.A.3[1a]. By receiving adaptation instructions from the control logic module 140 and causing the selected set of whitelist pre-equalizer coefficient values to be used for the transmitter side, the channel management module 116 performs the controller function of using one of the registers to set initial coefficient values for the one or more transmit filters.

88. To the extent a controller requires executable software, Berke teaches or suggests implementing its teachings using processors, software, and combinations of hardware and software. Ex-1004, ¶¶41-42, ¶2, ¶34. Therefore, an obvious implementation would have included a controller to perform Berke’s teachings, including the whitelist and channel management functionalities in Berke’s transmitter. As explained for [1a], the equalization coefficient whitelist 118 includes multiple registers that store pre-equalization coefficients. *See above* §V.A.3[1a]; Ex-1004, ¶¶12, 30. Processors included numerous registers, and it would have been obvious to include Berke’s registers in a processor, particularly given Berke’s teachings of customizable hardware such as a “Programmable Gate Array (FPGA)” “or a device embedded on a larger chip[.]” Ex-1004, ¶41.

89. Furthermore, Cornelius teaches that its active plug includes a “cable microcontroller” that is “used to configure clock and data recovery circuits,” which include pre-equalization. Ex-1005, 6:11-42. When combined with Berke (*see* above §V.A.8[8pre] (motivation subsection explaining how and why Berke and Cornelius would have been combined)), the microcontroller would obviously have performed configuration functions such as the configuration of pre-equalization coefficients; therefore, the microcontroller would obviously have included the registers taught by Berke and performed the Berke’s teachings for selecting and using those registers.

9. Dependent Claim 10

10. The communications link of claim 8, wherein at least some of the different channel models presume different types of receiver equalization.
--

90. Ground 1 renders obvious the communications link of claim 8. *See* above §V.A.8. Ground 1 further renders obvious the additional limitation of claim 10 for the same reasons as explained for claim 3. *See* above §V.A.4.

10. Independent Claim 11

[11pre] A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

[11a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

91. Limitations [11pre]-[11a] are identical to [8pre]-[8a]. Ground 1 renders obvious [11pre]-[11a] for the same reasons explained for [8pre]-[8a]. *See* above §§V.A.8[8pre]-[8b].

[11b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values for which the port transceiver determines a performance characteristic, the controller using one of the registers selected by the port transceiver to specify the initial coefficient values for the one or more transmit filters.

92. As explained with respect to limitation [8b], the combination of Berke and Cornelius teaches a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values, the controller using one of the registers to specify the initial coefficient values for the one or more transmit filters. *See* above §V.A.8[8b].

93. As discussed with respect to claim 4, Berke also teaches **determining a performance characteristic** for each set of initial coefficient values contained in the multiple registers. *See* above §V.A.5.

94. Furthermore, the combination of Berke and Cornelius teaches a controller. Berke teaches the controller on the transmitter side controlling the pre-equalizer settings based on adaptation instructions from the control logic module 140, configured on the receiver side. *See* above §V.A.8[8b]; Ex-1004, ¶13, ¶¶19-20. Berke teaches that the control logic module 140, configured on the receiver side, selects a whitelist set of pre-equalizer coefficient values and its corresponding register by determining their corresponding performance characteristics. *See* above §V.A.8[8b]; Ex-1004, ¶33.

95. The channel management module 116 is part of the transceiver that teaches the pluggable module transceiver. *See* above §V.A.8[8b]. It follows that the control logic module 140, being on the receiver side, is part of the transceiver that teaches the **port transceiver**. *See* above §V.A.8[8pre] (motivation subsection explaining how and why Berke and Cornelius would have been combined). The combination of Berke and Cornelius thus teaches each of the multiple registers containing a set of initial coefficient values **for which the port transceiver determines a performance characteristic**, the controller using one of the registers **selected by the port transceiver** to specify the initial coefficient values for the one or more transmit filters.

11. Dependent Claim 13

13. The communications link of claim 11, wherein the performance characteristic is a bit error rate.

96. Ground 1 renders obvious the communications link of claim 11. *See* above §V.A.10. Ground 1 further renders obvious the additional limitation of claim 13 for the same reasons as explained for claim 6. *See* above §V.A.6.

12. Dependent Claim 14

14. The communications link of claim 11, wherein the port transceiver generates updates for the coefficient values of the one or more transmit filters.
--

97. Ground 1 renders obvious the communications link of claim 11. *See* above §V.A.10. Ground 1 further renders obvious the additional limitation of claim 14.

98. As explained with respect to limitation [1b], Berke teaches a training process that updates the pre-equalizer coefficient values. *See* above §V.A.3[1b]. A pre-equalizer is a transmit filter. *See* above §V.A.8[8a].

99. The control logic module 140 conducts the training process that generates the updates:

With the tuning setting of channel output module 112 set to the minimum offset setting, control logic module 140 trains transmitter 110 and receiver 130 by running through the various combinations of settings for the equalization setting, the gain settings, the DFE taps

settings, and the FFE tap settings, and recording for each combination, an associated eye height and eye width, or an associated BER, as needed or desired.

Ex-1004, ¶20.

If the variability is greater than the equalization variation threshold, then control logic module 140 commences a set of training runs that are targeted to determining a small number of variables in the equalization coefficients that, when altered from one setting to another, produce a wide variance in the measured performance of serial channel 100, in terms of eye height and width, or in terms of BER.

Ex-1004, ¶24. As noted above with relation to limitation [11b], the control logic module 140 is part of the transceiver that teaches the port transceiver. *See* above §V.A.10[11b]. The combination of Berke and Cornelius (*see* above §V.A.8[8pre] (motivation subsection explaining how and why Berke and Cornelius would have been combined)) thus teaches that the port transceiver generates updates for the coefficient values of the one or more transmit filters.

13. Dependent Claim 15

15. The communications link of claim 14, wherein the controller saves updated coefficient values in one of the registers.
--

100. Ground 1 renders obvious the communications link of claim 14. *See* above §V.A.12. Ground 1 further renders obvious the additional limitation of claim 15.

101. As explained with respect to claim 14, Berke's control logic module 140 conducts the training process that generates updated coefficient values. *See* above §V.A.12. The control logic module 140 communicates adaptation instructions to channel management module 116 for effectuation on the transmitter side:

Control logic module 140 communicates with transmitter 110 via a communication channel established between channel management module 116 and channel management module 142. In a particular embodiment, management module 116 and channel management module 142 represent a separate side-band communication channel for communicating adaptation instructions between transmitter 110 and receiver 130.

Ex-1004, ¶19; *see* above §V.A.8[8b]. A POSITA would have understood the “adaptation instructions” include updated pre-equalizer coefficient values, since Berke teaches “the back channel adaptation mechanism may operate to provide different sets of compensation values.” Ex-1004, ¶¶13, 20.

102. Upon receipt of the “adaptation instructions” from the receiver, the transmitter then effectuates the changes locally. As explained with respect to claim 7, Berke teaches updating the selected register with the updated pre-equalizer coefficient values. *See* above §V.A.7. Berke further teaches that the “control logic module 140 determines a whitelist set of equalization coefficients based upon the characterizations, and stores the whitelist set of equalization coefficients to

equalization coefficient whitelists 118 and 144.” Ex-1004, ¶32. Accordingly, a POSITA would have understood that Berke’s channel management module 116 receives updated pre-equalizer coefficient values and causes them to be stored in the equalization coefficient whitelist 118. Accordingly, a POSITA would have understood that the combination of Berke and Cornelius renders obvious that the controller saves updated coefficient values in one of the registers.

14. Independent Claim 16

[16pre] A pluggable module transceiver including:

[16a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to a port transceiver; and

[16b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.

103. For the same reasons explained with respect to limitations [8pre]-[8b], Ground 1 renders obvious [16pre]-[16b]. *See* above §§V.A.8[8pre]-[8b].

15. Dependent Claim 17

17. The pluggable module transceiver of claim 16, wherein the port transceiver determines a performance characteristic for each register.

104. Ground 1 renders obvious the pluggable module of claim 16. *See* above §V.A.14. Ground 1 further teaches the additional limitation of claim 17 for the same reasons explained for limitation [11b]. *See* above §V.A.10[11b].

16. Dependent Claim 19

19. The pluggable module transceiver of claim 17, wherein the performance characteristic is a bit error rate.

105. Ground 1 renders obvious the pluggable module of claim 17. *See* above §V.A.15. Ground 1 further teaches the additional limitation of claim 19 for the same reasons as explained for claim 6. *See* above §V.A.6.

B. Ground 2

106. It is my opinion that a POSITA would have found claims 1, 3, 4, 6-8, 10, 11, 13-17, and 19 are obvious based on the teachings of Berke and Ran. Berke teaches techniques in a serial communication system for selecting initial equalization coefficient values and updating them in a training process, while Ran teaches serial communications between a network interface controller and a pluggable optical module. The motivation to combine is explained below. *See* below §V.B.3[8pre].

1. U.S. Patent No. 10,715,357 (“Ran”) (Ex-1006)

107. I have reviewed the Ran reference. Ran was filed on April 30, 2019 and claims priority to provisional Patent Application No. 62/722,517 (Ex-1007, “Ran-Provisional”), filed August 24, 2018. I understand Ran is prior art to the ’111 patent under §102(a)(2). *See* below Appendix 4 (Understanding of the Law).

108. Based on my review, Ran-Provisional provides clear and unambiguous support for Ran’s teachings that I cite in this Declaration. I have

analyzed both Ran-Provisional and Ran. The following table shows how cited disclosures of Ran are supported by corresponding disclosures from Ran-Provisional.

<u>Citations to in Ran (Ex-1006)</u>	<u>Corresponding Disclosure from Ran-Provisional (Ex-1007)</u>
Abstract	000022
1:23-41	000002
2:26-39	000004
2:66-3:13	000004
3:14-28	000005, 000007
3:29-50	000005
4:5-13	000006
4:23-30	000006
5:4-14	000007
5:15-25	000007
5:26-33	000007-08
5:39-49	000010
5:50-61	000010
6:23-44	000012-13
6:52-62	000011
7:36-59	FIG. 2, 000001, 000007, 000010
7:60-67	FIG. 2, 000001, 000007, 000010
8:1-10	FIGS. 2-3, 000007-08
8:43-50	000017
9:45-59	FIGS. 2-3, 000008
10:6-14	FIG. 3A, 000008
11:45-55	000011
11:65-12:9	000011-12
12:10-14	000012
12:15-21	000012
12:22-31	000012
12:41-50	000013
13:24-30	FIGS. 2-6, 000011-13
13:31-48	000013
14:59-67	000015
15:1-22	000015-16

15:45-50	000016
19:64-67	000017
20:59-62	000017
FIG. 1	FIG. 1
FIG. 2A	FIG. 2
FIG. 2B	FIG. 2, 000001, 000007, 000010
FIG. 3A	FIG. 3
FIG. 3B	FIG. 3A
FIG. 4A	FIG. 4
FIG. 5A	FIG. 5
FIG. 6A	FIG. 6

109. Based on my review, Ran-Provisional also provides clear and unambiguous support for at least independent claim 13 of Ran. The following table identifies exemplary support in Ran-Provisional for each limitation of Ran's claim 13.

Claim 13 of Ran	Exemplary Support from Ran-Provisional
A method comprising:	Claim 8, Title
receiving, by a receiver in a first device, signals over a lane of a communications link from a transmitter in a second device, the first device including a plurality of management registers and the second device including a plurality of management registers;	Claim 8, 000007, FIG. 2
communicating, by a management entity, equalization settings from the first device to the second device by copying content of a management register of the first device to a management register of the second device and during link operation;	Claim 8, 000009-13
selecting, by the management entity, an equalization coefficient for the transmitter of the lane from the second device to configure equalization settings of an equalizer tap coefficient independent of change of a coefficient of another tap; and	Claim 8, 000005, 000022

reconfiguring equalization for the communications link after the communications link is established.	Claim 10
--	----------

110. Ran is directed to “adaptation of a transmit equalizer using management registers.” Ex-1006, Title. Ran teaches that “[s]election of equalization coefficients to configure a communications link between a receiver in a host system and a transmitter in an optical or electrical communication module is performed by a management entity with access to management registers in the receiver and transmitter.” Ex-1006, Abstract.

111. Ran further teaches a training process where “[c]ontinuous modification of the selected equalization coefficients is enabled on the communications link after the communications link is established to handle varying operating conditions such as temperature and humidity.” Ex-1006, Abstract. Ran teaches implementing its serial communications link between a network interface controller and a pluggable optical module, each having a transceiver and coupled together via a port. Ex-1006, 1:23-33, 4:5-13, 4:23-30, 5:26-33, FIGS. 1, 2A, 2B. Ran’s serial communications link supports communication standards including the IEEE 802.3-2015 Ethernet standard, PCIe, Universal Serial Bus (USB), SATA, etc. Ex-1006, 2:66-3:13, 13:31-40, 14:59-67, 15:1-22.

112. Ran is in the same field of endeavor as the ’111 patent because Ran, like the ’111 patent, discloses systems and methods for serial communications. *See*

above §II; *e.g.*, Ex-1006, 1:34-41 (“[s]erial communication between a transmitter and a receiver at very high data rates”). Ran is also reasonably pertinent to the problem with which the inventors of the ’111 patent were involved because both Ran and the ’111 patent relate to ensuring reliable data transmission through a communications channel. *Compare* Ex-1001, 1:19-35 (“As demand continues for ever-higher data rates, extensions to the standard must account for increased channel attenuation and dispersion even as the equalizers are forced to operate at faster symbol rates.”) *with* Ex-1006, Abstract, 1:23-41 (“Serial communication between a transmitter and a receiver at very high data rates requires equalization to mitigate frequency dependent signal attenuation.”).

2. Claims 1, 3, 4, 6, and 7

113. Ground 2 renders obvious the methods of claims 1, 3, 4, 6, and 7 based on the teachings of Berke for the same reasons as explained above for Ground 1. *See* above §§V.A.3-7. As explained above, Berke teaches a SerDes communications method. *See* above §V.A.3[1pre]. Ran additionally teaches SerDes communications as explained below. *See* below §V.D.2[1pre]. A POSITA would have been motivated, and found it obvious, to combine Berke and Ran’s teachings by using Berke’s high-speed serial communications in a SerDes context, such as a connection between a network interface controller and an optical module, as taught by Ran. *See* below §V.B.3[8pre]. Berke’s teachings suggest serialization

and deserialization because it teaches techniques that enable high-speed communications over a channel. *See* above §V.A.3[1pre]. Ran teaches a system with such high-speed communications (Ex-1006, 3:25-28) that a POSITA would have found Berke's teachings to be suitable for.

3. Independent Claim 8

[8pre] A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

114. Ran teaches a **chip-to-module communications link** between a chip and an optical module, to which a POSITA would have found it obvious and been motivated to apply Berke's equalization techniques. As explained above for limitation [8pre] in Ground 1, Berke applies its teachings broadly to serial communications over circuit traces on printed circuit boards, connectors, and/or transmission cables. *See* above §V.A.8[8pre]; Ex-1004, ¶12. Berke teaches its communications link being implemented with network interfaces such as "a NIC disposed within information handling system 400," which includes, for example, "Gigabit Ethernet channels." Ex-1004, ¶¶39, 35.

115. Ran teaches an optical module being connected to a chip via a serial link, forming a "chip-to-module interconnection[]."

An optical module is a hot-pluggable optical transceiver that is used in high-bandwidth data communications applications. The optical module typically has an electrical connection to a chip (for example, a host

application specific integrated circuit (ASIC)). This electrical connection is a high data rate serial link (referred to as xGAUI-n where AUI stands for Attachment Unit Interface, x is the data rate per lane in Giga bits per second (Gbps), and n is the number of 30 lanes), with current rates of 50 Gb/s per lane (50GAUI-1). 50GAUI-1 is a one lane serial link used for chip-to-chip or chip-to-module interconnections.

Ex-1006, 1:23-33; 5:4-14 (“The multi-lane Attachment Unit Interface (AUI) Chip-to-Module forms a single link between the network interface controller 100 and an optical module 120 by transferring data in parallel over multiple lanes.”). Ran demonstrates the above described chip-to-module communications link in Figures 1 and 2A.




FIG. 2A

Ex-1006, FIG. 2A.

116. Ran teaches a **port connector** on a host device, such as an Ethernet port on the network interface controller (NIC). Ran illustrates in Figure 1, “a block diagram illustrating Ethernet port logic in a network interface controller 50.” Ex-1006, 4:5-13, FIGS. 1, 2A, 2B. The port connector is connected **to a port transceiver**. Ex-1006, 5:26-33 (“The network interface controller 100 also includes a host receiver 206 and a host transmitter 208 for one lane of the xGAUI-1 electrical link between the network interface controller 100 and the optical module 120.”), 8:1-10 (“Host 250 can include transceiver 252 for communication with module 232.”); FIGS. 2A-B.

117. Ran teaches a **pluggable module transceiver** that is configured on the module. Ex-1006, 1:23-33 (“An optical module is a hot-pluggable optical transceiver.”), 5:26-33 (“The optical module 120 also includes a module receiver 212 and module transmitter 210 for one lane of the xGAUI-1 electrical link between the network interface controller 100 and the optical module 120.”),⁷ FIGS. 2A, 2B.

⁷ Ran illustrates a transceiver as a transmitter and a receiver in the same device. For example, Figure 2A represents the “hot-pluggable optical transceiver” (Ex-1006, 1:23-33) as receiver 212 and transmitter 210.

118. Ran's Ethernet port acts as a port connector that couples the host device with the optical module. Ex-1006, 4:23-30 ("The network interface controller 50 is communicatively coupled via an xGAUI-n (a high data rate serial link) electrical connection to an optical module 70."), FIG. 1 (showing Ethernet port logic):




FIG. 1

See also Ex-1006, 5:4-14 ("The multi-lane Attachment Unit Interface (AUI) Chip-to-Module forms a single link between the network interface controller 100 and an optical module 120 by transferring data in parallel over multiple lanes. In addition, the system can be used for C2C intercommunications."). Because the port

connector (e.g., Ethernet port) connects the host device and the optical module for the purpose of transmitting and receiving data, it couples the port transceiver to the pluggable module transceiver. *See, e.g.*, Ex-1006, 5:26-33, 7:36-59 (“Transceiver 238 can be used for electrical signal transmission and receipt between module 232 and host network interface device 250. Transceiver 238 can provide multiple transmit and receive lanes for electrical signal communication between module 232 and host device 250.”), 7:60-67, FIGS. 2A-2B (showing transceiver on host and module are coupled).

119. **Motivation to Combine.** Berke teaches techniques for serial communications that apply to a wide variety of channels including circuit traces on printed circuit boards, connectors, and/or transmission cables. Ex-1004, ¶12. Ran teaches a chip-to-module communications link that couples a host device with an optical module. Ex-1006, 1:23-33; 4:23-30, 5:4-14. Ran’s type of link precisely uses the board traces, connectors, and cables that Berke applies its teachings to. Therefore, a POSITA would have been motivated to apply Berke’s teachings to the serial link of Ran because it would have simply used Berke’s known serial techniques to improve similar devices (e.g., Ran’s serial link) in the same way as taught by Berke.

120. The references apply their teachings to similar devices. Berke teaches using a network interface to connect its information handling system to external

devices and implementing the network interface using a network interface card (NIC). Ex-1004, ¶¶38-39. Ran teaches a NIC that includes a port connector for connecting with external modules, such as a pluggable optical module. Ex-1006, 1:23-33, 4:5-13, 5:26-33, FIGS. 1 and 2A-2B.

121. A POSITA would have understood that Ran's configuration of coupling a host device and an external module via an Ethernet port connector was commonly used and well known well before 2018 in applications such as data center networking, high-speed Ethernet communications, and server interconnects, where pluggable optical modules were routinely used to enable modular, scalable, and serviceable network interface configurations. Accordingly, a POSITA would have understood Berke's network interface to contemplate configurations that involve connecting a NIC with an external pluggable optical module, as taught by Ran. A POSITA would have been motivated to implement Berke's network interface according to the teachings of Ran given the prevalence of such modules in enterprise and data center applications.

122. Furthermore, Berke's technique of configuring pre-equalization coefficient values provided an efficient and effective process of transmitter and receiver to communicate and configure their communication channel to compensate for the insertion loss and other signal degradations. Ex-1004, ¶13. This

addresses a problem recognized by Ran. Ran teaches that the increase in data rate makes it necessary to optimize equalization settings:

Fixed equalization as done in existing C2M AUIs may not work well in 100 Gb/s per lane, as the variations between modules will likely require different equalization settings, and variable operating conditions will have a higher effect on circuits and boards.

Ex-1006, 2:26-39; *see also* Ex-1006, 1:34-41; *see below* §§V.B.3[8a]-[8b]. A POSITA would have recognized that Berke's technique was readily applicable to Ran in solving the problem that Ran recognizes. Berke's equalization was advantageous in that it obtains satisfactory equalization configurations in a shorter amount of time than systems that exhaust all possible configurations. Ex-1004, ¶3. Applying Berke's equalization to Ran would have provided such benefits and help Ran achieve its goal of providing the desired equalization for its communications link. A POSITA would have recognized that this ability to quickly obtaining satisfactory equalization configurations was particularly useful for the applications that Ran contemplates, where continuous modification of equalization coefficients is needed to "handle varying operating conditions such as temperature and humidity." Ex-1006, Abstract.

123. Additionally, the teachings of Berke and Ran were readily compatible with each other. Ran and Berke are in the same field of serial communications. *See above* §§V.A.1, V.B.1. Both apply their teachings to circuit boards and cables, for

computers and external devices. *See, e.g.*, Ex-1004, ¶12 (“The components can also be understood to include devices of an information handling system, such as a hard drive, a storage array, and the like, that are separate from the printed circuit board of the information handling system, where transmission channel 120 can include one or more transmission cables.”), ¶37; Ex-1006, 1:23-33, 4:23-30 (“The network interface controller 50 is communicatively coupled via an xGAUI-n (a high data rate serial link) electrical connection to an optical module 70.”), 5:4-14, 5:33-45, FIGS. 1, 2A-2B. Berke and Ran both contemplated connections according to common standards such as PCIe and SATA. Ex-1004, ¶12; Ex-1006, 13:31-40, 14:59-67. Ran teaches chip-to-module communications, for example, for Ethernet. *See, e.g.*, Ex-1006, 5:4-14, FIGS. 1, 2A, 2B. Berke contemplated its systems/methods be used for “Gigabit Ethernet channels.” Ex-1004, ¶39. The combination would not have changed the principle of operation for either reference because it uses Berke’s serial communications techniques for exactly the purpose taught by Berke: to improve pre-equalization for serial communications. In addition to the reasons explained above, a POSITA would have had a reasonable expectation of success in the combination because it would have been a straightforward combination that simply applies Berke’s serial communications techniques to Ran’s serial link. As explained above, the teachings were readily compatible and used for the same purposes taught in the references themselves. Pluggable

modules, pre-equalization, and serial communications had been widely known and used for many years before 2018. *See* above §I.

[8a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and
--

124. As explained above in Ground 1 with respect to limitation [8a], Berke teaches **one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted**. *See* above §V.A.8[8a]. Ran teaches performing equalization on the chip-to-module communication link. *See, e.g.*, Ex-1006, 3:14-28. As explained above, a POSITA would have been motivated to apply Berke's pre-equalization techniques to Ran's system. *See* above §V.B.3[8pre] (motivation subsection explaining how and why Berke and Ran would have been combined). Under this combination, FFE functionalities, as taught by Berke, are implemented in Ran's pluggable module transceiver to pre-equalize a corresponding serial symbol stream being transmitted **to the port transceiver**.

[8b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.
--

125. As explained above in Ground 1 with respect to limitation [8b], Berke teaches this limitation. *See* above §V.A.8[8b]. Furthermore, Ran teaches using a management entity to control pre-equalization settings. *See* below §V.D.2[1a]; Ex-

1006, 5:4-25, 9:45-59, 11:65-12:31, FIGS. 2A, 3, 6. Ran teaches or at least suggests implementing the management entity on the optical module to configure pre-equalization for transmission from the optical module to the host. *See* below §V.D.9[8b]; Ex-1006, Abstract, 5:39-61, 6:52-62, 13:24-30. When combined with Berke (*see* above §V.B.3[8pre] (motivation subsection explaining how and why Berke and Ran would have been combined)), the management entity would obviously have performed configuration functions such as the configuration of pre-equalization coefficients; therefore, the management entity would obviously have included the registers taught by Berke and performed Berke's teachings for selecting and using those registers.

4. Dependent Claim 10

10. The communications link of claim 8, wherein at least some of the different channel models presume different types of receiver equalization.

126. Ground 2 renders obvious claim 10 in light of Berke's teachings for the same reasons as explained above for Ground 1. *See* above §V.A.9, §V.A.4.

5. Independent Claim 11

[11pre] A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

[11a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

127. Limitations [11pre]-[11a] are identical to [8pre]-[8a]. Ground 2 renders obvious [11pre]-[11a] for the same reason explained for [8pre]-[8a]. *See* above §§V.B.3[8pre]-[8a].

[11b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values for which the port transceiver determines a performance characteristic, the controller using one of the registers selected by the port transceiver to specify the initial coefficient values for the one or more transmit filters.

128. As explained with respect to limitation [8b], the combination of Berke and Ran teaches a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values, the controller using one of the registers to specify the initial coefficient values for the one or more transmit filters. *See* above §V.B.3[8b]. As discussed for Ground 1 with respect to claim 4, Berke also teaches **determining a performance characteristic** for each set of initial coefficient values contained in the multiple registers. *See* above §V.A.5.

129. As further explained for Ground 1 with respect to limitation [11b], Berke's control logic module 140, configured on the receiver side, selects a

whitelisted set of pre-equalizer coefficient values and its corresponding register by determining their corresponding performance characteristics. *See above* §V.A.10[11b]; Ex-1004, ¶33. Berke's channel management module 116, on the transmitter side, is part of the transceiver that teaches the pluggable module transceiver. *See above* §V.B.3[8b], §V.A.8[8b]; Ex-1004, FIG. 1. It follows that the control logic module 140, being on the receiver side, is part of the transceiver that teaches the **port transceiver**. *See above* §V.B.3[8pre] (motivation subsection explaining how and why Berke and Ran would have been combined). The combination of Berke and Ran thus teaches each of the multiple registers containing a set of initial coefficient values **for which the port transceiver determines a performance characteristic**, the controller using one of the registers **selected by the port transceiver** to specify the initial coefficient values for the one or more transmit filters.

6. Dependent Claims 13-15

130. Claims 13-15 each depends directly or indirectly from claim 11. Ground 2 renders obvious the communications link of claim 11. *See above* §V.B.5. The additional limitations of claims 13-15 are further rendered obvious in light of Berke's teachings for the same reasons as explained above for Ground 1. *See above* §§V.A.11-13.

7. Independent Claim 16

[16pre] A pluggable module transceiver including:

[16a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to a port transceiver; and

[16b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.

131. For the same reasons explained with respect to limitations [8pre]-[8b], Ground 2 renders obvious [16pre]-[16b]. *See* above §§V.B.3[8pre]-[8b].

8. Dependent Claims 17 and 19

132. Claims 17 and 19 each depends directly or indirectly from claim 16. Ground 2 renders obvious the pluggable module of claim 16. *See* above §V.B.7. Ground 2 renders obvious the additional limitation of claims 17 and 19 for the same reasons explained for limitation [11b] and claim 13, respectively. *See* above §V.B.5[11b], §V.B.6.

C. Grounds 3 and 4

133. Claims 2 and 9 recite simple dependent limitations about insertion loss and package loss. These were well-known considerations taught by a textbook on serial communications (Stauffer). Claims 5, 12, and 18 recite using error signal energy as a performance characteristic. Stauffer teaches using root mean square as a performance characteristics, which is the square root of mean square error, a commonly used measure of error signal energy. Therefore, these dependent claims

are taught or suggested by Stauffer combined with Ground 1 (Berke and Cornelius) or with Ground 2 (Berke and Ran). The motivation to combine is explained below. *See below §§V.C.2-3.*

1. High Speed Serdes Devices and Applications (“Stauffer”) (Ex-1008)

134. I have reviewed the Stauffer reference. Stauffer is a book that was published in 2009 by Springer. I understand Stauffer is prior art to the ’111 patent under §102(a)(1). *See below Appendix 4 (Understanding of the Law).*

135. Stauffer is directed to “High Speed Serdes (HSS) Devices and Applications.” Ex-1008, 000001. Chapter 1 of Stauffer is intended to provide the reader with “a basic understanding of the rationale for using high-speed serializer/deserializer (Serdes) devices, and the inherent problems introduced by the high-speed operation of such devices.” Ex-1008, 000011. Stauffer covers features, functions, and design of HSS devices, as well as features and functions of protocol logic used to implement various network protocol interface standards, including those comply with serial communication standards including PCIe, SATA, and Ethernet. *Id.*, 000005-000010, 000058-000066.

136. Stauffer teaches performing equalization at the transmitter and/or receiver device. Ex-1008, 000020-000021. Stauffer notes: “Most Serdes transmitter implementations include a FFE.” *Id.* Stauffer teaches that signal integrity analysis is important for designing high-speed SerDes devices and

systems and “[t]he equalization scheme used by Serdes devices must compensate for the channel loss and other impairments in order to achieve the desired BER.”

Ex-1008, 000029.

137. Stauffer is in the same field of endeavor as the ’111 patent, because Stauffer, like the ’111 patent, is directed to serial communications. *See* above §II; *e.g.*, Ex-1008, 000011. Stauffer is also reasonably pertinent to the problem with which the inventors of the ’111 patent were involved because the ’111 patent is related to ensuring reliable data transmission through a communication channel and Stauffer addresses this problem by discussing signal integrity issues and solutions. Ex-1008, 000029-000033 (recognizing major channel impairments like insertion loss, reflections, and crosstalk, as well as loss over packages of transmitter and receiver chips, must be analyzed to preserve signal integrity).

2. Dependent Claim 2

2. The method of claim 1, wherein at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.
--

138. As explained above, the combination of Berke and Cornelius (Ground 1) and the combination of Berke and Ran (Ground 2) each renders obvious the method of claim 1. *See* above §§V.A.3, V.B.2. Grounds 3 and 4 each further renders obvious the additional limitation of claim 2.

139. Berke teaches storing sets of pre-equalizer coefficients, respectively corresponding to different channel models, in registers of a whitelist. *See* above §V.A.3[1a]. As explained with respect to limitation [8pre], the combination of Berke and Cornelius (Ground 1) and the combination of Berke and Ran (Ground 2) each teaches applying Berke's techniques to a chip-to-module (C2M) communications link. *See* above §§V.A.8[8pre], V.B.3[8pre]. Thus, according to the combinations, it would have been obvious for **at least some of the different channel models** to be directed to **chip-to-module (C2M) channels**.

140. As explained above, in Berke, for a set of pre-equalizer coefficients to be added to the whitelist, they must yield high performance (e.g., robust eye heights or eye widths, low BER) and fit a set of channel characteristics of a channel model so as to compensate for channel loss/distortion. *See* above §V.A.3[1a]; Ex-1004, ¶30, ¶32, ¶15. Berke teaches that "insertion loss and other signal degradations of transmission channel" are among the characteristics that have an impact on the performance and can be considered as part of a channel model:

Serial channel 100 operates to provide back channel adaptation where transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values within the transmitter and the receiver to compensate for the insertion loss and other signal

degradations of transmission channel 120. A determination is made as to whether or not a set of compensation values is satisfactory.

Ex-1004, ¶13; *supra* §V.A.3[1a]. A POSITA would have found it obvious for Berke's "other signal degradations of transmission channel" to include package loss, as package loss can be determined based on "knowledge gained during the design, development, or manufacturing of serial channel." Ex-1004, ¶30.

141. To the extent Berke does not literally mention "package loss," Stauffer expressly teaches package loss as a cause of signal degradation that should be considered in configuring equalization coefficients. Stauffer teaches that "[t]he equalization scheme ... must compensate for the channel loss and other impairments in order to achieve the desired BER." Ex-1008, 000029. In discussing "signal integrity analysis" (Ex-1008, 000029-000031), Stauffer explains "signal degradation due to the package" is a challenge to signal integrity and that package loss can be determined based on information from the manufacturer:

In addition to the channel as described above, the differential signal must also propagate through the chip package for the transmitting chip and the chip package for the receiving chip. ... Package models are supplied by the chip manufacturer which model the transfer function of the package. Analysis of the overall interconnect is performed by cascading the transmitter package model, the channel transfer function, and the receiver package model.

Ex-1008, 000031.

142. As shown by Stauffer's teachings, a POSITA would have understood Berke's "other signal degradations of transmission channel" to include signal degradation caused by package loss. A POSITA would have further found it obvious for Berke's "knowledge gained during the design, development, or manufacturing of serial channel" to include package models provided by the manufacturer, that were used to determine package loss. Thus, the combination of Berke, Cornelius, and Stauffer (Ground 3) and the combination of Berke, Ran, and Stauffer (Ground 4) each renders obvious that at least some of the different channel models are for chip-to-module (C2M) channels **with different insertion loss and package loss characteristics.**

143. **Motivation to Combine.** Berke teaches pre-equalization techniques for serial communication where whitelisted pre-equalizer coefficients are used to compensate for channel loss. *See* above §V.A.3[1a]; Ex-1004, ¶3, ¶13. Stauffer teaches that package loss is a cause of signal degradation that is separate from the insertion loss of the channel:

While it would be convenient from an analysis viewpoint to consider chip packages as part of the channel interconnect, it is not practical to access the connections inside the package in order to measure channel response in the lab. For this reason, channel measurements are performed from package pads as described previously, and the transfer function of the package is considered separately.

Ex-1008, 000031, 000029 (“Major channel impairments include insertion loss, reflections, and crosstalk.”). Berke provides express motivation to combine because it aims to configure compensation values “to compensate for ... signal degradations of transmission channel.” Ex-1004, ¶13. Indeed, signal degradation is exactly why Berke uses whitelisted coefficients: to compensate for known signal degradation and thereby provide “robust” performance. Ex-1004, ¶30. For high-speed scenarios where robust performance is needed—for example, enterprise applications (*see* Ex-1004, ¶39) that require high data throughput or extended cable lengths—a POSITA would have been motivated to maximize signal quality. While a POSITA might also be motivated to address other factors beyond package loss, that would not preclude them from specifically addressing package loss, which was a known source of signal degradation. Therefore, a POSITA would have found it obvious and been motivated to select sets of pre-equalizer coefficients to address the package loss characteristics that Stauffer teaches, as part of the channel model. *See* Ex-1008, 000031.

144. As explained above, a POSITA would have been motivated and found it obvious to combine the teachings in Berke and Cornelius (Ground 1) and Berke and Ran (Ground 2). *See* above §§V.A.8[8pre] (motivation subsection explaining how and why Berke and Cornelius would have been combined), V.B.3[8pre] (motivation subsection explaining how and why Berke and Ran would have been

combined). A POSITA would have had a reasonable expectation of success in considering package loss, as taught by Stauffer, in applying the teachings of Berke and including package loss as part of the channel model. This is because Stauffer shows that it was commonly known that package loss was a contributor to signal degradation and methods for measuring and addressing package loss were readily known, and Berke expressly teaches configuring equalization to address signal degradations.

145. A POSITA would have understood that the teachings of Berke and Stauffer are compatible with each other. Stauffer and Berke are in the same field of serial communications. Berke contemplates that its high speed serial communications link is compliant with different standards, including PCIe and SATA. Ex-1004, ¶¶12, 37. Stauffer is directed towards SerDes devices supporting a variety of protocols, including PCIe and SATA, making it compatible with Berke. Ex-1008, 000058, 000061. Applying Stauffer's teachings would have merely involved applying well-known understanding of a POSITA and would not have changed the principle of operation of Berke, Cornelius, Ran, or Stauffer.

3. Dependent Claim 5

5. The method of claim 4, wherein the performance characteristic is an error signal energy.
--

146. As explained above, the combination of Berke and Cornelius (Ground 1) and the combination of Berke and Ran (Ground 2) each renders obvious the

method of claim 4. *See* above §§V.A.5, V.B.2. Grounds 3 and 4 each further renders obvious the additional limitation of claim 5.

147. As explained above, Berke teaches using the initial pre-equalizer coefficient values to determine a performance characteristic for each of the multiple registers. *See* above §V.A.5. Berke does not limit what performance characteristic to determine, but provides examples including the “eye height and eye width” of “the receiver eye diagrams” and the “BER of serial channel.” Ex-1004, ¶33.

148. Stauffer teaches various methods for determining metrics of channel performance, including calculation of the “root-mean-square value.” Ex-1008, 000031-32. The root mean square (RMS) value was a well-known concept in signal processing taught in undergraduate-level courses, which was calculated by taking the square root of the mean square error (MSE)—the mean of the squares of error signals. The mean square error was a well-known measure of the error signal energy. Ex-1027, 000002 (“[MSE] has a clear physical meaning—it is the natural way to define the energy of the error signal.”). A POSITA would have understood that root mean square and mean square error are mathematically related metrics that rise and fall together, since RMS is simply the square root of MSE. As such, both metrics reflect the same underlying signal characteristics, and using one instead of the other does not change the result of a relative performance

comparison. By teaching assessing channel performance based on RMS, Stauffer teaches or suggests using the error signal energy, as measured by MSE, as a performance characteristic.

149. A POSITA would have found it obvious to apply Stauffer's teachings to Berke and use error signal energy as measured by MSE as one of the metrics of channel performance. Determining error signal energy for Berke's communications channel would have been obvious because the same signal data used to generate Berke's eye diagram and BER would have been available for calculating the error signal energy. Error signal was a known measurement that compared two signals: the received signal and the expected signal. *See generally* Ex-1024, cls.1-2 (teaching "determin[ing] an error signal value between a measured received signal and a desired signal"), FIGS. 6A-6B, ¶17, FIG. 7. During training, Berke uses the received signal to create an eye diagram: "[r]eceiver eye diagram 200 is derived from a training run." Ex-1004, ¶22. Berke uses an expected training signal to calculate the BER. First, Berke "trains transmitter 110 and receiver 130 by running through the various combinations of settings ... and recording for each combination, an associated eye height and eye width, or an associated BER[.]" Ex-1004, ¶20. BER is calculated by comparing received bits versus expected bits; in order to determine which bits were erroneous, the transmitter sends a pre-defined sequence of bits, called a training signal, which the receiver looks for and

compares with the received signal. This was the conventional training approach that a POSITA would have understood from Berke's references to training algorithms (Ex-1004, ¶¶24-30), and it would have been obvious because, without a known training signal, the receiver would not be able to determine which bits are correct and which were flipped. Therefore, Berke teaches or suggests the use of training sequences (Ex-1004, ¶¶24-26) and already compares the received signal to the expected training signal to calculate BER (Ex-1004, ¶¶24-26). The error signal is simply the difference between the received and expected signals, and the additional step of calculating its energy as the MSE was simple and obvious arithmetic. In fact, a POSITA would have also recognized that the MSE determination was commonly applied to error signals in the context of equalizer adaptation in characterizing channel performance. This basic method was taught in textbooks. *See, e.g.*, Ex-1022, 000033 ("There are several algorithms that can be used for adapting the equalizer. The most popular one for compact hardware implementation is mean-square error (MSE) calculated between the recovered signal and the training data in the time domain at sampling points[.] ... The LMS algorithm optimizes the filter coefficients based on minimizing the mean-squared error."). Accordingly, no additional data collection or measurement would have been required to determine error signal energy, and its calculation would have been known and obvious to a POSITA.

150. A POSITA would have been motivated, with a reasonable expectation of success, to determine the error signal energy as an indicator of channel performance on top of those enumerated by Berke. The consideration of MSE metrics would have been consistent with Berke's teaching of evaluating channel performance based on the height of eye diagram opening. Ex-1004, ¶22. In a eye diagram, the vertical opening (eye height) representing voltage margin. A POSITA would have recognized that the MSE of the error signal quantifies vertical eye closure. Given that Berke already teaches using the eye diagram as performance characteristics, it would have been an obvious and natural extension to quantify that performance by computing the MSE of the underlying error signals, providing precise metrics that correlate with the visual indicators in the eye diagram. *See generally* Ex-1028, cl.19 (teaching estimating performance metrics of a system with an equalizer based on "error signal power"); Ex-1029, 7:35-62 (using "error signal power" to compare performance of different "channel models"), Table 1. A POSITA would have been motivated to incorporate this additional performance characteristic alongside eye diagram and BER to obtain a more comprehensive understanding of channel behavior. This is consistent with and advances Berke's interest in using the training method to find coefficient settings that are "more robust." *See, e.g.*, Ex-1004, ¶31.

4. Dependent Claim 9

9. The communications link of claim 8, wherein at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.

151. As explained above, the combination of Berke and Cornelius (Ground 1) and the combination of Berke and Ran (Ground 2) each renders obvious the method of claim 8. *See* above §§V.A.8, V.B.3. Grounds 3 and 4 each further renders obvious the additional limitation of claim 9 for the same reasons as explained for claim 2. *See* above §V.C.2.

5. Dependent Claim 12

12. The communications link of claim 11, wherein the performance characteristic is an error signal energy.

152. As explained above, the combination of Berke and Cornelius (Ground 1) and the combination of Berke and Ran (Ground 2) each renders obvious the communication link of claim 11. *See* above §§V.A.10, V.B.5. Grounds 3 and 4 each further renders obvious the additional limitation of claim 12 for the same reasons as explained for claim 5. *See* above §V.C.3.

6. Dependent Claim 18

18. The pluggable module transceiver of claim 17, wherein the performance characteristic is an error signal energy.
--

153. As explained above, the combination of Berke and Cornelius (Ground 1) and the combination of Berke and Ran (Ground 2) each renders obvious the

pluggable module transceiver of claim 17. *See above* §§V.A.15, V.B.8. Grounds 3 and 4 each further renders obvious the additional limitation of claim 18 for the same reasons as explained for claim 5. *See above* §V.C.3.

D. Ground 5

154. It is my opinion that a POSITA would have found claims 1-19 obvious based on the teachings of Ran, Mejia, and Stauffer. Ran teaches techniques for selecting among presets of equalization settings and updating the settings in the context of serial communications between a network interface controller and a pluggable optical module. Mejia teaches a preset search process for selecting a preset that is suitable for a channel. Stauffer teaches storing equalization coefficients in registers, among other teachings. The motivation to combine is explained below. *See below* §§V.D.2[1a], V.D.3, V.D.4.

1. U.S. Patent Application Publication No. 2014/0237301 (“Mejia”) (Ex-1009)

155. I have reviewed the Mejia reference. Mejia was filed on February 15, 2013 and published on August 21, 2014. I understand Mejia is prior art to the ’111 patent under §§102(a)(1)-(2). *See below* Appendix 4 (Understanding of the Law).

156. Mejia is directed to “techniques that utilize preset evaluation to improve input/output performance in high-speed serial interconnects.” Ex-1009, ¶1. Mejia recognizes a problem that “some high-speed serial I/O implementations

(such as PCIe) operate at frequency ranges at which the ISI becomes a major challenge for achieving a target high-speed data transfer.” Ex-1009, ¶2.

157. To address the problem, Mejia proposes a method for “evaluat[ing] the system performance at different recommended equalization preset values and select[ing] the best equalization preset for the given system.” Ex-1009, ¶10. Mejia teaches a preset search technique that is applicable to a scenario where various presets for equalization settings are available (such as under the PCIe standard) for a transmitter. Ex-1009, ¶12 (“A vast number of settings within an equalization map are valid for Gen3 (generation 3) transmitters, ranging from 0 db to –9.8 db of de-emphasis and 0 db to 8.0 db pre-shoot. Within this equalization map, the [PCIe] Base Specification defines a set of equalization ‘presets.’”). The preset search technique selects the presets that best fit characteristics of the channel, including the topology or configuration of the system and the specific channel loss associated with the system. Ex-1009, ¶¶29-32.

158. Mejia is in the same field of endeavor as the ’111 patent because Mejia, like the ’111 patent, discloses systems and methods for serial communications. *See* above §II; *e.g.*, Ex-1009, Abstract, ¶¶1-2, ¶¶10-14. Mejia is also reasonably pertinent to the problem with which the inventors of the ’111 patent were involved because both Mejia and the ’111 patent relate to ensuring reliable data transmission through a communication channel. *Compare* Ex-1001,

1:19-35 (“As demand continues for ever-higher data rates, extensions to the standard must account for increased channel attenuation and dispersion even as the equalizers are forced to operate at faster symbol rates.”) *with* Ex-1009, ¶2 (“And, some high-speed serial I/O implementations (such as PCIe) operate at frequency ranges at which the ISI becomes a major challenge for achieving a target high-speed data transfer.”).

2. Independent Claim 1

[1pre]. A SerDes communications method that comprises, in a transceiver:

159. Ran discloses a **SerDes communications method**. Ran is directed to “[s]erial communication between a transmitter and a receiver at very high data rates.” Ex-1006, 1:34-41, 4:23-30 (“The network interface controller 50 is communicatively coupled via an xGAUI-n (a high data rate serial link) electrical connection to an optical module 70.”). For example, Ran teaches “a high data rate serial link (referred to as xGAUI-n ...)” that connects an optical module on a cable to a Ethernet port on a network interface controller (NIC). Ex-1006, 1:23-33 (“The optical module typically has an electrical connection to a chip (for example, a host application specific integrated circuit (ASIC)). This electrical connection is a high data rate serial link (referred to as xGAUI-n where AUI stands for Attachment Unit Interface, x is the data rate per lane in Giga bits per second (Gbps), and n is the number of lanes).”); *see also* FIG. 2A.




FIG. 2A

160. As was well-known to a POSITA, xGAUI-n interfaces transmit high-speed serial data over multiple lanes, each of which typically uses a SerDes to convert between parallel data (inside the chip) and serial data (for transmission over a lane). Each electrical lane in an xGAUI interface corresponds to a high-speed serial link that relies on a SerDes to convert wide parallel data links into serial symbol streams for transmission, and to perform the inverse operation on the receiving end. Furthermore, Ran applies its teachings to standards with SerDes architectures, such as PCIe and SATA. Ex-1006, 13:31-40, 14:59-67. The PCIe and SATA standards include serialization and deserialization, which a POSITA would have understood given the widespread use of both standards. *See above* §V.A.3[pre]; *e.g.*, Ex-1010, 000192-000193; Ex-1011, 000621; Ex-1008, 000058-

61. Because Ran provides teachings for common SerDes architectures, Ran teaches a method for SerDes communications.

161. Ran teaches a **transceiver** configured on the optical module or the NIC. Ex-1006, 1:23-33 (“An optical module is a hot-pluggable optical transceiver.”), 5:26-33 (“The optical module 120 also includes a module receiver 212 and module transmitter 210 for one lane of the xGAUI-1 electrical link between the network interface controller 100 and the optical module 120.”), 5:26-33 (“The network interface controller 100 also includes a host receiver 206 and a host transmitter 208.”), FIGS. 2A-B.

[1a] selecting one of multiple registers to specify initial pre-equalizer coefficient values, each of the multiple registers corresponding to a different channel model;
--

162. Ran teaches performing pre-equalization in a transceiver based on pre-equalizer coefficient values. Ran first notes the use of pre-equalizer in serial communication:

Serial communication between a transmitter and a receiver at very high data rates requires equalization to mitigate frequency dependent signal attenuation. Equalization can be applied at the transmitter (Tx) and/or at the receiver (Rx). Typically, a combination of transmitter and receiver equalization is used.

Ex-1006, 1:34-41. Ran teaches a pre-equalizer that is “a transmitter FFE of any length and step size” which has “settings” that could be adjusted “using requests

from a receiver.” Ex-1006, 3:14-28. Ran’s FFE performs pre-equalization because it is configured on the transmitter side and it addresses signal attenuation before the signal is transmitted to the receiver. *Id.* A POSITA would have understood that these adjustable “settings” include equalization coefficients of the transmitter FFE, which are pre-equalizer coefficient values. *Id.* Ran adjusts the equalization coefficients for the transmitter side based on requests from the receiver side. Ex-1006, 3:14-28 (“Various embodiments provide for a receiver to adjust settings of a transmitter FFE of any length and step size using requests from a receiver.”).

163. Ran teaches using a register to **specify initial pre-equalizer coefficient values** that are selected from multiple sets of pre-equalizer coefficient values, called “presets.” Ex-1006, 10:6-14, FIG. 3A. Ran teaches that the transceiver on each side of its serial link includes “management registers” for controlling equalization for the channel including selection of a preset. Ex-1006, 5:4-14 (“FIG. 2A illustrates a set of management registers in a network interface controller 100 and an optical module 120 for one lane of a multi-lane Attachment Unit Interface (AUI) Chip-to-Module (C2M) to enable control of equalization in the receiver in the optical module 120 by the network interface controller 100.”), 5:15-25 (“The host management registers 202 and module management registers 204 are used at least in connection with two-way communications to perform timing recovery and equalization for a single-lane xGAUI-1 electrical link between

the network interface controller 100 and the optical module 120 or other electrical interconnect module device.”), FIG. 2A:




FIG. 2A

The management registers on each side include a request register and a response register. Ex-1006, 5:15-25 (“Host management registers 202 in the network interface controller 100 include a host request register 220 and a host response register 222. Module management registers 224 in the optical module 120 include a module request register 224 and a module response register 226.”). Figure 3A shows a control field structure that is applicable to the request registers at both the host and the module. Ex-1006, 9:45-59 (“FIG. 3A depicts an embodiment of a

control field structure in a request register 300. Register 300 can be used as one of two registers in a two-way communication channel on at least one lane of an AUI Chip-to-Chip (C2C) or Chip-to-Module (C2M). Any field or parameter of the control field structure in the request register 300 may be written in host request register 220 and module request register 224 shown in FIG. 2A.”), FIG. 3A:




FIG. 3A

164. With the management registers described above, Ran teaches selecting preset values for the transmitter FFE by setting the module request

register. Ran teaches that each request register has a two-bit “initial condition request field 304,” used “to select one of three predefined transmitter equalizer configurations (‘presets’).” Ex-1006, 10:6-14. Figure 3B shows how bits 12 and 13 (corresponding to the initial condition request field) are set differently to select one of Preset1-Preset3 for transmitter equalizer configurations.

Bit 13	Bit 12	Description
1	1	Preset3
1	0	Preset2
0	1	Preset1
0	0	Co-efficient Control

FIG. 3B

Ex-1006, FIG. 3B. A predefined transmitter equalizer configuration includes settings for coefficients of the pre-equalizer. Because these coefficients are specified as an “initial condition,” they are used as **initial pre-equalizer coefficient values**. Ex-1006, 10:6-14. Therefore, the two bits in the initial condition request field of the request register specifies initial pre-equalizer coefficient values by **selecting** a set of such values from three possible sets respectively corresponding to Preset1, Preset2, and Preset3.

165. Mejia teaches a “preset search” technique for **selecting** the initial pre-equalizer coefficient values from multiple sets of initial pre-equalizer coefficient values. Ex-1009, ¶10.

166. Mejia is directed to “utiliz[ing] preset evaluation to improve input/output (I/O or IO) performance in high-speed serial interconnects such as PCIe.” Ex-1009, ¶10, ¶1. Mejia’s preset search technique is applicable to a scenario where various presets for equalization settings are available (such as under the PCIe standard) for a transmitter. Ex-1009, ¶12 (“A vast number of settings within an equalization map are valid for Gen3 (generation 3) transmitters, ranging from 0 db to –9.8 db of de-emphasis and 0 db to 8.0 db pre-shoot. Within this equalization map, the [PCIe] Base Specification defines a set of equalization ‘presets.’”). A preset is used as the starting point of the equalization training process where initial equalization coefficients are refined into final coefficients for use in transmitting data. Ex-1009, ¶12 (“At the beginning of the equalization training sequence, the link partners communicate starting point for the equalization via preset exchange.”), ¶13 (“PCIe 3.0 introduces adaptive backchannel equalization, where both upstream and downstream partners interact between the TX (Transmit) and RX (Receive) on each given lane to establish optimum equalization settings. ... A preset generally refers to starting setting(s) (which may be later refined into final settings).”). Here, the equalization settings for the transmitter include pre-equalizer coefficient values. Ex-1009, ¶11 (“As discussed herein, an equalization preset value (e.g., as supported for 8.0 GT/s transmitter) generally refers to a preset value defined in terms of voltage ratio(s), relating the

pre-cursor and post-cursor equalization voltages and a margin value generally refers to the capability of the link to guarantee compliance and interoperability in production or real world conditions.”). Mejia’s selected preset corresponds to initial pre-equalizer coefficient values because it is intended to be used as the starting point of equalization training.

167. Mejia teaches selecting or specifying a set of pre-equalizer coefficients from among multiple sets. Mejia’s preset search technique is “provided to evaluate the system performance at different recommended equalization preset values and select the best equalization preset for the given system.” Ex-1009, ¶¶31, 10. The preset search is “designed to calculate/determine and compare the margin(s) associated with different equalization presets” to select the best preset or “recommend the best equalization preset value(s) based on information such [as] gain, jitter and equalization controls, which can be obtained from the RX (Receiver) circuits.” Ex-1009, ¶10. Preset search also evaluates the presets for their bit error ratio (BER). Ex-1009, ¶32 (“Preset search flow will guarantee target BER by selecting the best preset on those boards; (b) the flow reduces the risk of having OEM (Original Equipment Manufacturer) configuring the system to a wrong preset.”), ¶12 (“Generally, the presets should at least guarantee a Bit Error Rate (BER) of 10^{-4} . The subsequent equalization process aims to take the link to a $BER < 10^{-12}$.”).

168. Mejia’s different equalization presets **correspond to different channel models**. Mejia’s different presets contain pre-equalizer coefficient values that are defined to compensate for “signal distortion” of a communication channel (Ex-1009, ¶2), including “counteract[ing] Inter-Symbol Interference” (Ex-1009, ¶12). Different presets are designed to be suitable for the communication channels of different systems. Ex-1009, ¶10 (“In an embodiment, a flow (that may be automated) is provided to evaluate the system performance at different recommended equalization preset values and select the best equalization preset for the given system.”), ¶31. As a POSITA would have understood, variations in channel loss or signal distortion arise from differences in channel characteristics—such as design, settings, and environmental conditions—where different channel characteristics correspond to different channel models.

169. Pre-equalizer coefficients of Mejia’s different presets correspond to and deliver high performance for different channel characteristics that impact channel loss, including channel characteristics related to the design (e.g., channel topology, material) and settings (e.g., channel configurations) of the systems:

In some embodiments: (a) the flow ensures that the system is going to be adjusted to use the best preset on any topology/configuration and with any specific Add-In Card (AIC). For example, some board manufactures may use different type materials that might not aligned directly with recommendation(s). On such systems, the channel loss

will not be directly matching with the expected loss per inch average numbers. Preset search flow will guarantee target BER by selecting the best preset on those boards.

Ex-1009, ¶32. The channel characteristics further include the channel length and insertion loss, which also impact channel loss:

Furthermore, in some embodiments, optimal preset that maximizes electrical margins may be a function of channel insertion loss as follows: (1) shortest channels (e.g., less than or equal to 3 inch or equivalent channel insertion loss), where best point is found at or near to P5 (or Preset 5); (2) medium channel (e.g., between 3-8 inch or equivalent channel insertion loss), where the best point is found at or near to P3 (or Preset 3); and (3) long channel (e.g., more than 8 inch or equivalent channel insertion loss), where the best point is found at or near to P7 (or Preset 7).

Ex-1009, ¶36. In addition, Mejia further teaches recommending presets that provide high performance based on information from the receive circuits, such as “equalization controls.” Ex-1009, ¶10 (“Some embodiments may also recommend the best equalization preset value(s) based on ... equalization controls, which can be obtained from the RX (Receiver) circuits.). A POSITA would have understood that equalization controls include the types and configurations of equalization, and it would have been obvious for such information obtained from the receiver side to indicate the equalization type used at the receiver and/or a resulting control

influenced by the equalization type used at the receiver. This teaches or suggests that the channel characteristics further includes receiver equalization types.

170. Each of the aforementioned channel characteristics impact the channel loss/distortion and frequency response characteristics. Ex-1009, ¶¶31-32, ¶36. As Mejia's pre-equalizer coefficients for each preset provide high performance for certain channel characteristics (Ex-1009, ¶¶31-32, ¶36), this preset corresponds to a channel model that represents or assumes the channel characteristics. Mejia's multiple presets of pre-equalizer coefficient values (Ex-1009, ¶10, ¶12, ¶31) are therefore designed to compensate for different channel loss/distortion properties resulting from channel characteristics that correspond to different channel models. This includes, for example, channel models with different amounts of insertion loss and related factors such as different channel length, topology, and materials. Ex-1009, ¶¶31-32, ¶36. To account for these different channel models, Mejia performs preset search to select a preset with suitable pre-equalizer coefficients from multiple options (e.g., P3, P5, P7). *See, e.g.*, Ex-1009, ¶36. Accordingly, Mejia teaches selecting from different presets corresponding to different channel models.

171. **Motivation to Combine.** Ran teaches selecting one preset from multiple presets but does not limit how the presets are structured or how the selection is performed. Ex-1006, 10:6-14, FIG. 3B. Mejia teaches a preset search technique that enables identifying a preset that is suitable for the characteristics of

a communication channel. Ex-1009, ¶¶10, ¶31. Mejia's preset search technique precisely provides a way of implementing the functionality of Ran. A POSITA would have found it obvious and been motivated to apply Mejia's preset search technique to implement Ran's selection of a preset, because it simply applies a known technique (e.g., the preset search technique) to a known method/system ready for improvement (e.g., Ran's equalization training method/serial channel) in the same way to obtain a predictable result (e.g., selecting a preset set of pre-equalizer coefficient values with optimum performance).

172. By evaluating performance characteristics of different presets on a channel and selecting the best preset, Mejia allows identifying and selecting pre-equalizer coefficients that best fit the characteristics of a channel and provide optimum and satisfactory performance. Ex-1009, ¶¶10, ¶¶31-32, ¶36. A POSITA thus would have been motivated to apply Mejia's preset search teachings to Ran to bring about the benefits.

173. A POSITA would have understood that the performance of pre-equalization was impacted by characteristics of the channel such as its length, material, layout, and loss profile. Both Ran and Mejia recognize variabilities in the channel, for example, due to different manufacturers, use cases, configurations, settings, and even environmental conditions. Ex-1006, 2:26-39 ("Transmit equalization is implementation specific and may be tuned by the host board

manufacturer, based on the board characteristics, to optimize the eye opening.”), 15:45-50 (“[F]or example, temperature and humidity [] may cause an increase in bit errors over the communications link.”), Abstract; Ex-1009, ¶32 (recognizing considerations of “topology/configuration and with any specific Add-In Card” in preset selection), ¶36 (recognizing considerations of channel length and insertion loss in preset selection). Having presets corresponding to different sets of characteristics or channel models would have provided sets of initial pre-equalizer coefficients that are suitable for different channels. A POSITA would have recognized that this is beneficial as it allows suitable initial pre-equalizer coefficient values to be applied to a variety of channel configurations and use cases and would have been motivated by such benefits to configure different presets to correspond to different channel models.

174. Furthermore, a POSITA would have found Mejia’s preset search technique to be an effective way to select from multiple sets of pre-equalizer coefficient values because directly evaluating performance characteristics of different pre-equalizer coefficient values would have allowed one to quantify the performance of different options and identify one that has optimum and satisfactory performance. For example, Mejia teaches that its selected preset provides the best performance and meets the expected BER for initial equalization settings. *See, e.g.*, Ex-1009, ¶31 (“[Preset search] may also recommend the best

equalization preset value(s) based on information from the RX (Receive) circuits.”) ¶32 (“Preset search flow will guarantee target BER by selecting the best preset on those boards”). Given the effectiveness of Mejia’s preset search technique and its consistency with Ran’s goal of picking a starting point for equalization training, a POSITA would have been motivated to use this technique to implement Ran’s selection of a preset from multiple options.

175. Additionally, the teachings of Ran and Mejia were readily compatible with each other. Ran and Mejia are in the same field of serial communications. *See* above §§V.B.1, V.D.1. Both Ran and Mejia support serial communications compliant with the PCIe standard. *See, e.g.*, Ex-1006, 13:31-40 (“The I/O adapters 716 may include a Peripheral Component Interconnect Express (PCIe) adapter that is communicatively coupled over bus 744 to a network interface controller 100.”); Ex-1009, ¶10 (“Some embodiments utilize preset evaluation to improve input/output (I/O or IO) performance in high-speed serial interconnects such as PCIe.”). Both Ran and Mejia teach selecting a preset from a plurality of presets to provide initial pre-equalizer coefficient values for equalization training. *See, e.g.*, Ex-1006, 10:6-14, FIG. 3B; Ex-1009, ¶¶12-14. A POSITA would have a reasonable expectation of success in implementing methods or systems according to the combination because it would not have changed the principle of operation of either Ran or Mejia. This is because the combination uses Mejia’s preset search

technique for exactly the purpose taught by Mejia: to select a preset that is suitable in light of channel characteristics.

176. It would have been obvious to store pre-equalizer coefficient values in registers. Ran teaches using a value stored in register to select a preset and specify initial pre-equalizer coefficient values; Ran does not limit how the coefficient values themselves are stored. A POSITA would have understood that, by 2018, there were a finite number of options for storing equalization coefficients, e.g., registers and memories. *See, e.g.*, Ex-1004, ¶30. A POSITA would have had good reason to use registers, especially given that Ran already teaches that both its NIC and optical module are configured with registers (Ex-1006, 5:4-25), making the additional registers for storing equalization coefficients simple and obvious to implement. Furthermore, a POSITA would have been motivated to load the equalization coefficients into registers to take advantage of faster access times—for example, the single-clock cycle access offered by register implementations, explained below. Under this implementation, it would have been obvious to use each of multiple registers to store one set of pre-equalizer coefficients. This was a simple and reliable way to organize the sets of coefficients. As explained with respect to Ground 1, in view of the typical width of multi-bit registers in transceiver circuits, storing a set of three or more FFE tap coefficients per register was a straightforward and efficient design choice that would have been readily

apparent to a POSITA with experience in serial communications. *See above* §V.A.3[1a]. This design was obvious to implement under well-known protocols such as I²C and SPI⁸ and would have been advantages because it allowed a set of pre-equalizer coefficients to be loaded in one clock cycle, with a single register read, as compared to more complicated schemes. *See above* §V.A.3[1a]. Moreover, the use of registers for this purpose—with each register storing a set of coefficients—was a long-standing practice, and the prior art provides express motivation to do this. Ex-1030, 1:10-25, 2:23-51. This design was particularly suitable for Ran’s system, which uses a two-bit “initial condition request field 304” “to select one of three predefined transmitter equalizer configurations (‘presets’).” Ex-1006, 10:6-14. When each of multiple registers stored one set of equalizer configurations, the selection of an equalizer configuration could be achieved by simply correlating values of the two bits in the initial condition request field to different register addresses. Accordingly, using multiple registers—each holding a set of pre-equalizer coefficients—not only falls squarely within the teachings of Ran but also represents an obvious implementation that offers advantages that would have motivated a POSITA to use the design. As such, each register would

⁸ Ran contemplates implementing these known protocols. Ex-1006, 8:43-50, 19:64-67, 20:59-62.

store one set of pre-equalizer coefficients corresponding to a preset. Here, the register for a certain preset corresponds to the channel model of that preset. By selecting one of the presets whose coefficients are stored in a register, Ran renders obvious **selecting one of multiple registers** to specify initial pre-equalizer coefficient values, **each of the multiple registers** corresponding to a different channel model.

177. To the extent Ran does not expressly teach storing pre-equalizer coefficients in registers, Stauffer teaches using registers to store pre-equalizer coefficient values. Ex-1008, 000037-000042. Stauffer teaches using an FFE module for an HSS (“high-speed SerDes”) channel. Stauffer teaches storing coefficients of the FFE module (e.g., FFE coefficients C0-C2 for Tap0-Tap2 of the FFE) in registers (e.g., registers 0x4, 0x5, and 0x6). Ex-1008, 000050 (“The HSS EX10 transmitter includes a feed forward equalization (FFE) to reduce intersymbol interference (ISI) at the receiver.”), *Id.* 000041 (“Table 2.6 HSS EX IO transmitter slice register definitions”):

0x4	16	R/W	0x0000	Transmit Tap0 Coefficient Register
	3:0	R/W	0000	FFE Tap 0 Coefficient This register's value is applied to the analog logic after "Apply Load" (Transmit Coefficient Control Register 0x02 bit 0) is pulsed. The value read from this register is the actual value being driven to the analog logic. Value is unsigned magnitude. See the Transmit Polarity Register 0x08 for sign values
	15:4	R		Unused
0x5	16	R/W	0x003F	Transmit Tap1 Coefficient Register
	5:0	R/W	111111	FFE Tap 1 Coefficient See the description for the Transmit Tap0 Coefficient Register (0x04)
	15:6	R		Unused
0x6	16	R/W	0x0000	Transmit Tap2 Coefficient Register
	4:0	R/W	00000	FFE Tap 2 Coefficient See the description for the Transmit Tap0 Coefficient Register (0x04)
	15:5	R		Unused

178. Stauffer's teachings shows that storing pre-equalizer coefficient values in registers was well known in the art in the context of serial communications.⁹ This further confirms that it would have been obvious to store initial pre-equalizer coefficient values in registers.

⁹ While Stauffer teaches storing coefficients of different FFE taps in registers 0x4, 0x5, and 0x6, etc., Stauffer states that this is a tutorial example that does not limit actual implementations. Ex-1008, 000049 ("HSS EX10 is only a tutorial example... implementations may vary..."). Moreover, a POSITA would have understood that designation of registers was often a logical or organizational construct, rather than a rigid physical constraint. Registers are typically defined by their addressability and usage within a system, and it was well known in the art

[1b] updating the initial pre-equalizer coefficient values during a training phase; and
--

179. Ran teaches a training process that updates the pre-equalizer coefficient values. As explained above, Ran teaches selecting a set of pre-equalizer coefficient values corresponding to a preset for use as initial condition for a communication channel. *See* above §V.D.2[1a].

180. After the initial condition is set, Ran teaches that “[c]ontinuous modification of the selected equalization coefficients is enabled on the communications link after the communications link is established.” Ex-1006, 13:41-48, 3:14-28 (“Various embodiments provide for a receiver to adjust settings of a transmitter FFE of any length and step size using requests from a receiver. ... Request and response registers are accessible to a management entity to enable

that multiple adjacent or related registers can be grouped together and treated as a single larger register for purposes of data storage, access, or processing. Thus, a set of data stored across multiple registers can be equivalently regarded as stored in a single aggregated register. As explained for Ground 1, a POSITA would have found it advantageous to store a set of pre-equalizer coefficients in one register rather than multiple separate registers.

changing one equalization coefficient at a time, with a directed change request and arbitrarily fine steps.”).

181. Ran illustrates example processes for updating equalization coefficient values in Figures 6A and 6B. *See generally* Ex-1006, 11:65-13:30, FIGS. 6A-6B. For example, Ran teaches “[a] management entity may continue monitoring the host management registers 202 and module management registers 204 as described in conjunction with the flowgraph shown in FIG. 6A, to enable continuous or periodic adaptation of equalizer settings used to transmit signals.” Ex-1006, 6:23-44.




FIG. 6A

Ex-1006, FIG. 6A. In the training process illustrated by Figure 6A, a request from the optical module's receiver to change a pre-equalizer coefficient of the NIC's (i.e., host's) transmitter is communicated by a management entity that reads a value in the module request register and writes the value in the host request register. Ex-1006, 11:65-12:14, 6:37-44. The management entity communicates a confirmation of the requested change by reading the host response register and writing in the module response register. Ex-1006, 12:10-31, 6:37-44. The process is repeated and the receiver continues to "tune any transmitter equalizer settings"

until, for example, a sufficiently low bit error rate (BER) is detected by the receiver. Ex-1006, 12:41-50. When the training process is complete, a register ready bit in the host request register will be set to 1 to trigger data transmission. *Id.* Accordingly, Ran teaches updating the initial pre-equalizer coefficient values during a training phase.

182. Notably, Ran teaches using the same training process for the opposite direction of data transmission—updating pre-equalizer coefficients used by the transmitter of the optical module based on requests from the receiver of the NIC:

Selection of equalization coefficients to configure a communications link between a receiver in a host system and a transmitter in an optical or electrical communication module is performed by a management entity with access to management registers in the receiver and transmitter. Continuous modification of the selected equalization coefficients is enabled on the communications link after the communications link is established to handle varying operating conditions such as temperature and humidity.

Ex-1006, Abstract, 6:52-62 (“In an embodiment, there may be four additional registers, two module management registers in the optical module 120 and two host management registers in the network interface controller 100 to enable control of transmitter equalization in the optical module 120 by host or NIC 100.”), 13:24-30.

[1c] using the updated pre-equalizer coefficient values to convey a transmit data stream.

183. After the training process (*see* above §V.D.2[1b]), Ran teaches using the updated pre-equalizer coefficient values to convey a transmit data stream. Ex-1006, 12:41-50 (“The management entity may use the indication of ‘receiver ready’ bit set to ‘1’ to permit transmission of real data (as opposed to training data) from the host transmitter to the module for transmission to a remote station (link partner).”), 10:60-67 (“A “receiver ready” field 402 can be used to signal the local receiver state to the link partner or any entity (e.g., management entity). When this bit is logical ‘1’, it indicates that one or more lanes in the local receiver has completed training and is prepared to receive data.”), 5:50-54, 6:23-44, 9:60-64, 12:64-13:23.

3. Dependent Claim 2

2. The method of claim 1, wherein at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.

184. Ground 5 renders obvious the method of claim 1. *See* above §V.D.2. As explained below, Ground 5 further renders obvious the additional limitation of claim 2.

185. As explained above, Ran teaches chip-to-module (C2M) channels. *See* above §V.B.3[8pre]. Ran teaches selecting a preset in the context of C2M

channels. Ex-1006, 9:45-59 (“Register 300 can be used as one of two registers in a two-way communication channel on at least one lane of an AUI Chip-to-Chip (C2C) or Chip-to-Module (C2M).”), 10:6-14 (“A two-bit initial condition request field 304 is used to select one of three predefined transmitter equalizer configurations (‘presets’).”). The combination of Ran, Mejia, and Stauffer (*see* above §V.D.2[1a] (motivation subsection explaining how and why Ran, Mejia, and Stauffer would have been combined)) teaches storing presets with different sets of initial pre-equalizer coefficients that respectively provide high performance for different channel models, each having a set of characteristics of the channel. *See* above §V.D.2[1a]. In light of Ran’s teaching of using presets in the context of C2M channels, the combination contemplates having presets that provide high performance for a C2M channel. The **channel models** corresponding to these presets **are for chip-to-module (C2M) channels**.

186. Mejia teaches using channel **insertion loss** as a characteristic of a channel model. *See* above §V.D.2[1a]; Ex-1009, ¶36. In addition, Stauffer teaches that **package loss** as a cause of signal degradation that should be considered in configuring equalization coefficients. *See* above §V.C.2; Ex-1008, 000029-000031. In light of Stauffer’s teachings, a POSITA would have recognized that package loss impacts the performance of a particular pre-equalizer configuration. Given that Mejia’s preset search technique is “provided to evaluate the system

performance at different recommended equalization preset values and select the best equalization preset for the given system” (Ex-1009, ¶¶31, 10), a POSITA would have found it obvious to consider package loss in defining presets because doing so would have provided high performance presets for different package loss characteristics. Accordingly, the combination of Ran, Mejia, and Stauffer teaches implementing presets corresponding to channel models that are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.

187. **Motivation to Combine.** The combination of Ran and Mejia teaches selecting a preset of pre-equalizer coefficients that best compensate for channel loss (e.g., signal distortion) of a particular channel. Ex-1009, ¶2, ¶12; *see above* §V.D.2[1a]. Stauffer teaches that package loss is a cause of signal degradation that is separate from the insertion loss of the channel (*see above* §V.C.2; Ex-1008, 000029-000031). Signal degradation is exactly why Mejia uses preset search: to “select the best equalization preset for [a] given system” that best compensates for known signal degradation of the system. Ex-1009, ¶10.

188. For high-speed scenarios where robust performance is needed—for example, “high-bandwidth data communications applications” (Ex-1006, 1:23-41) that require high data throughput or extended cable lengths—a POSITA would have been motivated to maximize signal quality. While a POSITA might also be motivated to address other factors beyond package loss, that would not preclude

them from specifically addressing package loss, which was a known source of signal degradation. Therefore, a POSITA would have found it obvious and been motivated to select sets of pre-equalizer coefficients to address the package loss characteristics that Stauffer teaches, as part of the channel model. *See* Ex-1008, 000031.

189. A POSITA would have understood that the teachings of Ran, Mejia, and Stauffer are compatible with each other. Ran, Mejia, and Stauffer are in the same field of serial communications. The compatibility of these references is further confirmed because the teachings of all of them are applicable to, for example, the PCIe standard. Ex-1006, 13:31-40; Ex-1009, ¶10; Ex-1008, 000061. Applying Stauffer's teachings would have merely involved applying well-known understanding of a POSITA and would not have changed the principle of operation of Ran, Mejia, or Stauffer.

4. Dependent Claim 3

3. The method of claim 1, wherein at least some of the different channel models presume different types of receiver equalization.
--

190. Ground 5 renders obvious the method of claim 1. *See* above §V.D.2. Ground 5 further renders obvious the additional limitation of claim 3.

191. Stauffer teaches applying **different types of receiver equalization**. Stauffer teaches a Decision Feedback Equalizer (DFE) at the receiver and that the receiver operates in “one of three equalization modes: non-DFE, DFE3, or DFE5.”

Ex-1008, 000052-000053. Among the three equalization modes, at least DFE3 and DFE5 refer to different types of receiver equalization because they include different filter taps:

In non-DFE mode, the received input signal is processed through the variable gain amplifier (VGA) then captured in the front end logic and deserialized. In DFE3 and DFE5 modes, the deserialized samples are processed using DFE algorithms that automatically adjust the receiver threshold to better compensate for severe pattern-dependent distortions in the channel. ... DFE5 and DFE3 modes use 5 and 3 filter taps, respectively for feedback. ... DFE5 provides maximum performance in terms of BER, while DFE 3 may be used in cases where signal degradation is not as severe. The non-DFE mode uses a fixed threshold receiver for cases where signal degradation is small enough not to require further equalization.

Id., 000053-000054. Mejia further teaches another type of receiver equalization, CTLE. Ex-1009, ¶32.

192. As explained above, the combination of Ran, Mejia, and Stauffer teaches or suggests multiple registers that store sets of pre-equalizer coefficients, each of which corresponds to a channel model. *See* above §V.D.2[1a]. Both Ran and Mejia teach that the transmitter communicates with the receiver to determine the equalization coefficients to use for the communication channel. Ex-1006, 3:14-28 (“Various embodiments provide for a receiver to adjust settings of a transmitter FFE of any length and step size using requests from a receiver.”); Ex-1009, ¶31

(“Some embodiments may also recommend the best equalization preset value(s) based on information from the RX (Receive) circuits.”). Mejia teaches selecting a set of equalization coefficients based on “equalization controls” from the receiver side, which a POSITA would have understood to include the receiver equalization type, where it would have been obvious for the information received from the receiver side to indicate the equalization type used at the receiver and/or a resulting control influenced by the equalization type used at the receiver. Ex-1009, ¶10; *see* above §V.D.2[1a]. The receiver equalization type is among channel characteristics that can be represented or assumed by channel models. *See* above §V.D.2[1a]. Given that various types of receiver equalization were available, a POSITA would have found it obvious for **at least some of the different channel models to presume** different types of receiver equalization.

193. Having different channel models to presume different types of receiver equalization would have been obvious further because the type of receiver equalization used impacts the channel performance. Ex-1008, 000054 (explaining different DFE modes should be used to compensate different extent of channel distortions). A POSITA would have understood that, for a set of pre-equalization settings to achieve high performance (e.g., low BER) during operation of a serial channel, they must be suitable for the receiver equalization type. Ex-1009, ¶12 (“Generally, the presets should at least guarantee a Bit Error Rate (BER) of 10^{-4} .”).

This is because transmitter pre-equalization and receiver equalization were conventionally jointly co-designed to compensate for the overall channel loss (*see* above §I.A), the receiver equalization settings impact the overall channel performance and inform the decision as to what pre-equalizer coefficients to use. A POSITA would have found it obvious to consider receiver equalization type in framing the presets in order to have presets that achieve satisfactory performance. Doing so would have provided the benefit of improving the performance of presets available and offering a better initial condition for the equalization training process. A POSITA would have had a reasonable expectation of success in considering this characteristic as the relationship between receiver equalization and channel performance was well known. Ex-1008, 000054. Therefore, Ground 5 renders obvious that at least some of the different channel models presume different types of receiver equalization.

5. Dependent Claim 4

4. The method of claim 1, wherein said selecting includes using the initial pre-equalizer coefficient values to determine a performance characteristic for each of the multiple registers.

194. Ground 5 renders obvious the method of claim 1. *See* above §V.D.2.

Ground 5 further renders obvious the additional limitation of claim 4.

195. As explained above, the combination of Ran, Mejia, and Stauffer teaches multiple registers that store sets of pre-equalizer coefficients, each of

which corresponds to a channel model. *See* above §V.D.2[1a] (motivation subsection explaining how and why Ran, Mejia, and Stauffer would have been combined). Mejia’s preset search technique evaluates the performance characteristics of the pre-equalizer coefficient values corresponding to each preset in order to select the best performing preset. *See* above §V.D.2[1a]; Ex-1009, ¶31, ¶10. The performance characteristics evaluated include margin, BER, gain, jitter, etc. *See, e.g.*, Ex-1009, ¶31 (“[A] preset search may be performed ... to calculate/determine and compare the margin(s) associated with different equalization presets ... Some embodiments may also recommend the best equalization preset value(s) based on information from the RX (Receive) circuits.”), ¶32 (“Preset search flow will guarantee target BER by selecting the best preset.”), ¶10 (“Some embodiments may also recommend the best equalization preset value(s) based on information such gain, jitter and equalization controls, which can be obtained from the RX (Receiver) circuits.”). The combination thus teaches that the selecting includes using the initial pre-equalizer coefficient values to determine a performance characteristic for each of the multiple registers.

6. Dependent Claim 5

5. The method of claim 4, wherein the performance characteristic is an error signal energy.
--

196. Ground 5 renders obvious the method of claim 4. *See* above §V.D.5. Ground 5 further renders obvious the additional limitation of claim 5.

197. As explained above, Ground 5 teaches using the initial pre-equalizer coefficient values to determine a performance characteristic for each of the multiple registers. *See* above §V.A.5. In particular, Mejia teaches using BER, margins, gain, and/or jitter as example performance characteristics. *Id.*; Ex-1009, ¶¶31-34. Mejia does not limit what performance characteristic to determine, but provides examples.

198. As explained above, Stauffer teaches or suggests using the error signal energy, as measured by MSE, as a performance characteristic. *See* above §V.C.3; Ex-1008, 000031-32; Ex-1027, 000002.

199. A POSITA would have found it obvious to apply Stauffer's teachings to Ran and Mejia and use error signal energy as measured by MSE as one of the metrics of channel performance. Determining error signal energy for the communications channel taught by Ran and Mejia would have been obvious because the same signal data used to generate Mejia's margin and BER would have been available for calculating error signal energy. Error signal was a known measurement that compared two signals: the received signal and the expected signal. *See generally* Ex-1024, cls.1-2, ¶17, FIGS. 6A-6B, 7. Mejia uses an expected training signal to calculate margin and BER. Mejia teaches evaluating performance metrics such as margin and BER during the preset search process. Ex-1009, ¶31 ("[A] preset search may be performed ... to calculate/determine and

compare the margin(s) associated with different equalization presets”), ¶32 (“Preset search flow will guarantee target BER by selecting the best preset on those boards.”). Bit error rate is calculated by comparing received bits versus expected bits; in order to determine which bits were erroneous, the transmitter sends a pre-defined sequence of bits, called a training signal, which the receiver looks for and compares with the received signal. Mejia measures margins using two methods, “jitter injection and voltage margining.” Ex-1009, ¶11. Voltage margining, for example, conventionally involved applying controlled voltage offsets to the receiver’s decision threshold and monitoring the system for BER. In such cases, both the transmitted and received signals are known, enabling accurate BER calculation. Therefore, Mejia teaches or suggests the use of training sequences and already compares the received signal to the expected training signal to calculate BER and margin. Ex-1009, ¶¶31-32. The error signal is simply the difference between the received and expected signals, and the additional step of calculating its energy as the MSE was simple and obvious arithmetic. In fact, a POSITA would have also recognized that the MSE determination was commonly applied to error signals in the context of equalizer adaptation in characterizing channel performance. This basic method was taught in textbooks. *See, e.g.*, Ex-1022, 000033 (“There are several algorithms that can be used for adapting the equalizer. The most popular one for compact hardware implementation is mean-square error

(MSE) calculated between the recovered signal and the training data in the time domain at sampling points[.] ... The LMS algorithm optimizes the filter coefficients based on minimizing the mean-squared error.”). Accordingly, no additional data collection or measurement would have been required to determine error signal energy, and its calculation would have been known and obvious to a POSITA.

200. A POSITA would have been motivated, with a reasonable expectation of success, to determine the error signal energy as an indicator of channel performance on top of those enumerated by Mejia. Mejia’s voltage margining reveals the amplitude range over which the channel remains operational. Ex-1009, ¶11. A POSITA would have recognized that the MSE of the amplitude error signal affects such a range and it would have been an obvious and natural extension to consider a performance characteristic by computing the MSE of the underlying error signals. *See generally* Ex-1028, cl.19; Ex-1029, 7: 35-62 (using “error signal power” to compare performance of different “channel models”), Table 1. A POSITA would have been motivated to incorporate this additional performance characteristic alongside margin and BER to obtain a more comprehensive understanding of channel behavior. This is consistent with, and advances, Mejia’s goal of “select[ing] the best equalization preset for the given system.” *See, e.g.*, Ex-1009, ¶31.

7. Dependent Claim 6

6. The method of claim 4, wherein the performance characteristic is a bit error rate.
--

201. Ground 5 renders obvious the method of claim 4. *See* above §V.D.5.

Ground 5 further renders obvious the additional limitation of claim 6. As explained with respect to claim 4, Mejia teaches using the bit error rate (BER) as a “performance characteristic.” *See* above §V.D.5; Ex-1009, ¶¶12, 32.

8. Dependent Claim 7

7. The method of claim 4, further comprising updating the selected register with the updated pre-equalizer coefficient values.

202. Ground 5 renders obvious the method of claim 4. *See* above §V.D.5.

Ground 5 further renders obvious the additional limitation of claim 7.

203. As explained with respect to limitation [1b], Ran teaches updating the initial pre-equalizer coefficient values to updated pre-equalizer coefficient values during a training phase. *See* above §V.D.2[1b]. The combination of Ran, Mejia, and Stauffer teaches or suggests storing pre-equalizer coefficients in registers. *Id.*; Ex-1006, 12:10-31, 6:37-44. A POSITA would have recognized that the updated pre-equalizer coefficients obtained during the training phase represent a better fit with the characteristics of the channel, including topology, configuration, lengths, etc., and provide better performance than the starting point of the training. Using the updated values would have improved the efficiency of future training processes

by having a better starting point. This would have motivated a POSITA to replace the initial pre-equalizer coefficients with the updated pre-equalizer coefficients. At the same time, Ran teaches having a limited number of available presets. Ex-1006, 10:6-14 (“A two-bit initial condition request field 304 is used to select one of three predefined transmitter equalizer configurations (‘presets’).”), FIG. 3B. Given the limited space on the registers, it would have been obvious to update the selected register with the updated pre-equalizer coefficient values. Using existing register space to store updated pre-equalizer coefficient values would have provided the benefit of saving silicon space and cost.

9. Independent Claim 8

[8pre] A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

204. For the same reasons explained above for Ground 2, Ran teaches this limitation. *See* above §V.B.3[8pre].

[8a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

205. As explained above with respect to limitation [1a], Ran teaches a transmit FFE that pre-equalizes data to be transmitted on the chip-to-module communications link. *See* above §V.D.2[1a]. As serial channels transmit data symbols sequentially, the data going through a serial channel is a serial symbol

stream. A POSITA would have understood that the transmitter FFE is a type of FIR filter commonly used in serial communications. *See, e.g.,* Ex-1008, 000020 (referring to a three-tap FFE as “the filter” and each tap’s tap weight value as “a filter coefficient”). Therefore, Berke teaches **one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted.**

[8b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.
--

206. As explained above with respect to limitation [1a], the combination of Ran, Mejia, and Stauffer teaches **multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model.** *See* above §V.D.2[1a]. In addition, the combination teaches selecting one of the multiple registers to specify initial pre-equalizer coefficient values. *Id.* In this way, the selected register is **used to set initial coefficient values for the one or more transmit filters.** *See* above §V.D.9.[8a] (explaining that the transmitter FFE is a transmit filter).

207. Ran further teaches **a controller** performing the function of using a register to set initial coefficient values. As explained above, Ran teaches using management registers located on both sides of the serial link to control equalization, where the management registers include request and response

registers. *See* above §V.D.2[1a]; Ex-1006, 5:4-25, FIG. 2A. Each request register includes an “initial condition request field” that is used to select a preset from multiple options. *See* above §V.D.2[1a]; Ex-1006, 10:6-14, FIG. 3B. Ran further teaches using a management entity to control the value of the request register. Ex-1006, 9:45-59 (“FIG. 3A depicts an embodiment of a control field structure in a request register 300. ... [A] management entity can access a register that stores one or more of the parameters and copy one or more of the parameters to a destination register.”), 11:65-12:31, FIGS. 3, 6.

208. Ran teaches or suggests implementing the management entity on the transmitter side. For example, when setting pre-equalizer coefficients of a host’s transmitter based on requests from an optical module’s receiver, Ran teaches implementing the management entity on the host device (where the host transmitter is located). Ex-1006, 5:50-61 (“The host response register 222 and/or module request register 224 may also include a one-bit receiver ready field to indicate to a host management entity that the equalization tuning is complete and that data may be reliably transmitted from the NIC 100 to the optical module 120.”), 5:39-49 (“Module request register 224 is a read-only register that can only be read by a host management entity.”). This would have been a reasonable design choice because the management entity responds to requests by editing the host’s request register;

being located locally on the host would have allowed the management entity to easily edit that register.

209. Ran also teaches the same training process in the opposite data transmission direction, for setting pre-equalizer coefficients used by the optical module's transmitter based on requests from the host's receiver. *See* above §V.D.2[1b]; Ex-1006, Abstract, 6:52-62, 13:24-30. The management entity responds to requests—this time from the host—to edit the module's request register. *Id.*, Ex-1006, 11:65-12:14, 6:37-44. Because this request register resides on the optical module, a POSITA would naturally have interpreted Ran as teaching or at least suggesting that the management entity is likewise located on the optical module. Like the host example, where the management entity was local to the transmitter side (there, the host), it would also have made sense to have the management entity local to the transmitter (here, the optical module). This would have allowed the management entity to be local to the optical module, to locally edit the request register in that module. A POSITA would have recognized the benefits of doing so. For example, it would have been more efficient and secure to use a local management entity to access and edit the registers at the optical module, as compared with using a management entity on another device. In addition, such an in-module controller configuration would have simplified dependency management of the circuits. This symmetric approach is consistent with Ran's

teaching of using the same training process for both data transmission directions and would have simplified the control logic for implementing the management entity. Ex-1006, Abstract, 6:52-62, 13:24-30.

210. Components of Ran's optical module that carry out the functionalities of the management entity and the registers that store initial pre-equalizer coefficients, which the combination of Ran, Mejia, and Stauffer teaches or suggests (*see above* §V.D.2[1a]), constitute a controller at the pluggable module transceiver. The combination thus teaches or suggests a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.

10. Dependent Claim 9

9. The communications link of claim 8, wherein at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.

211. Ground 5 renders obvious the communications link of claim 8. *See above* §V.D.9. Ground 5 further renders obvious the additional limitation of claim 9 for the same reasons as explained for claim 2. *See above* §V.D.3.

11. Dependent Claim 10

10. The communications link of claim 8, wherein at least some of the different channel models presume different types of receiver equalization.

212. Ground 5 renders obvious the communications link of claim 8. *See* above §V.D.9. Ground 5 further renders obvious the additional limitation of claim 10 for the same reasons as explained for claim 3. *See* above §V.D.4.

12. Independent Claim 11

[11pre] A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

[11a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

213. Limitations [11pre]-[11a] are identical to [8pre]-[8a]. Ground 5 renders obvious [11pre]-[11a] for the same reason explained for [8pre]-[8a]. *See* above §§V.D.9[8pre]-[8b].

[11b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values for which the port transceiver determines a performance characteristic, the controller using one of the registers selected by the port transceiver to specify the initial coefficient values for the one or more transmit filters.

214. As explained with respect to limitation [8b], the combination of Ran, Mejia, and Stauffer teaches **a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values, the controller**

using one of the registers to specify the initial coefficient values for the one or more transmit filters. *See* above §V.D.9[8b].

215. As discussed for Ground 5 with respect to claim 4, the combination also teaches **determining a performance characteristic** for each set of initial coefficient values contained in the multiple registers. *See* above §V.D.5; Ex-1009, ¶¶11-12, 31-34, 50-51, 61-62, 67-68, FIG. 3.

216. Furthermore, Mejia teaches that the **determining** and **selecting** are performed on the receiver side. As explained above, Mejia teaches evaluating presets for performance characteristics such as margin and BER and selecting the best preset accordingly. *See* above §V.D.2[1a]; Ex-1009, ¶¶31-32. These performance characteristics are determined at the receiver side of a communications channel based on the received signal. *See, e.g.*, Ex-1009, ¶31 (“Some embodiments may also recommend the best equalization preset value(s) based on information from the RX (Receive) circuits.”). Furthermore, Ran teaches that the management entity sets pre-equalizer coefficient values on the transmitter side by copying bit values from the receiver’s request register to the transmitter’s request register, where the request register includes the initial condition request field that selects a set of initial pre-equalizer coefficient values. Ex-1006, 11:65-12:14, 10:6-14; *see* above §V.D.2[1a-1b]. This teaches performing the selecting step on the receiver side.

217. As explained above, Ran's optical module includes a pluggable module transceiver that includes transmit filters. *See* above §V.D.9[8pre-8a], §V.B.3[8pre]. When the optical module is on the transmitter side, Ran's port transceiver in the NIC (*see* above §V.B.3[8pre]) is on the receiver side. The combination thus teaches each of the multiple registers containing a set of initial coefficient values **for which the port transceiver determines a performance characteristic**, the controller using one of the registers **selected by the port transceiver** to specify the initial coefficient values for the one or more transmit filters.

13. Dependent Claim 12

12. The communications link of claim 11, wherein the performance characteristic is an error signal energy.
--

218. Ground 5 renders obvious the communications link of claim 11. *See* above §V.D.12. Ground 5 further renders obvious the additional limitation of claim 12 for the same reasons as explained for claim 5. *See* above §V.D.6.

14. Dependent Claim 13

13. The communications link of claim 11, wherein the performance characteristic is a bit error rate.
--

219. Ground 5 renders obvious the communications link of claim 11. *See* above §V.D.12. Ground 5 further teaches the additional limitation of claim 13 for the same reasons as explained for claim 6. *See* above §V.D.7.

15. Dependent Claim 14

14. The communications link of claim 11, wherein the port transceiver generates updates for the coefficient values of the one or more transmit filters.

220. Ground 5 teaches the communications link of claim 11. *See* above §V.D.12. Ground 5 further renders obvious the additional limitation of claim 14.

221. As explained with respect to limitation [1b], Ran teaches a training process that updates the equalization coefficient values for the transmitter FFE. *See* above §V.D.2[1b]. The transmitter FFE teaches a transmit filter. *See* above §V.D.9[8a].

222. Ran's training process "provide[s] for a receiver to adjust settings of a transmitter FFE of any length and step size using requests from a receiver." Ex-1006, 3:14-28. This teaches or suggests the receiver generating updates for the coefficient values for the transmitter FFE to adjust its settings.

223. As noted above with respect to limitation [11b], as Ran's optical module includes a pluggable module transceiver that includes transmit filters, the corresponding receiver is located in the port transceiver. *See* above §V.D.12[11b]. The combination thus teaches or suggests that the port transceiver generates updates for the coefficient values of the one or more transmit filters.

16. Dependent Claim 15

15. The communications link of claim 14, wherein the controller saves updated coefficient values in one of the registers.

224. Ground 5 teaches the communications link of claim 14. *See* above §V.D.15. Ground 5 further renders obvious the additional limitation of claim 15.

225. As explained above, Ran teaches a management entity carrying out a training process that updates the initial pre-equalizer coefficient values (*see* above §V.D.2[1b]) and suggests implementing the management entity in the optical module (*see* above §V.D.9[8b]). As further explained with respect to claim 7, the combination of Ran, Mejia, and Stauffer renders obvious updating the selected register with the updated pre-equalizer coefficient values. *See* above §V.D.8. Because components of Ran's optical module that carry out the functionalities of the management entity and the registers that store initial pre-equalizer coefficients constitute a controller at the pluggable module transceiver (*see* above §V.D.9[8b]), the combination teaches or suggests that the controller saves updated coefficient values in one of the registers.

17. Independent Claim 16

[16pre] A pluggable module transceiver including:

[16a] one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to a port transceiver; and

[16b] a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.

226. For the same reasons explained with respect to limitations [8pre]-[8b], Ground 5 renders obvious [16pre]-[16b]. *See* above §§V.D.9[8pre]-[8b].

18. Dependent Claim 17

17. The pluggable module transceiver of claim 16, wherein the port transceiver determines a performance characteristic for each register.

227. Ground 5 renders obvious the pluggable module of claim 16. *See* above §V.D.17. Ground 5 teaches the additional limitation of claim 19 for the same reasons explained for limitation [11b]. *See* above §V.D.12[11b].

19. Dependent Claim 18

18. The pluggable module transceiver of claim 17, wherein the performance characteristic is an error signal energy.

228. Ground 5 teaches the pluggable module of claim 17. *See* above V.D.18. Ground 5 discloses the additional limitation of claim 18 for the same reasons explained for claim 5. *See* above §V.D.6.

20. Dependent Claim 19

19. The pluggable module transceiver of claim 17, wherein the performance characteristic is a bit error rate.

229. Ground 5 teaches the pluggable module of claim 17. *See above* §V.D.18. Ground 5 discloses the additional limitation of claim 19 for the same reasons explained for claim 6. *See above* §V.D.7.

VI. Background and Qualifications

230. I am a tenured full professor of Electrical Engineering at the University of Southern California, where I have been a faculty member since 2011. As a professor, I research and teach electronic circuit and system design. My research focuses on analog mixed-signal and RF circuit architectures for data converters (both ADC and DAC), transceivers (both wireline and wireless), clock generation (such as PLL and DLL), digital VLSI, SoC, mostly-digital circuits, artificial intelligence (AI) and machine learning (ML) accelerators, analog CAD, and advanced signal processing techniques (e.g. adaptive filter, calibration, equalization, etc.) for analog circuit and systems. My research has been applied to wireline, wireless communications, computing, IoT systems, autonomous vehicles, etc.

231. I earned my Ph.D. and M.Sc. in Electrical Engineering and Computer Science from the University of California, Berkeley in 2006 and 2002 respectively.

I received my B.S. degree from National Taiwan University, where I graduated as the top student in the upper division.

232. My work has been recognized by both academic and commercial entities. I have published over 90 conference, journal, magazine articles, including multiple publications in the IEEE International Solid-State Circuits Conference (ISSCC), IEEE Symposium on VLSI Circuits (VLSIC), IEEE Custom Integrated Circuits Conference (CICC), IEEE European Conference on Solid-State Circuit (ESSCIRC), IEEE Radio Frequency Integrated Circuits Symposium (RFIC), IEEE Journal of Solid-State Circuits (JSSC), and Transactions on Circuit and System I (TCAS-I). My research has been funded by the National Science Foundation (NSF), ONR, IARPA, SRC, DARPA and various industry companies. I have also served as a technical consultant for companies including Qualcomm, Tetramem, BAE Systems and Samsung.

233. I have received recognition for my work and research. In 2024, I was honored with the grade of IEEE Fellow to recognize my technical contributions in solid-state circuit society. This grade is reserved for 0.1% of IEEE members. In 2023, I was the co-recipient of the ISSCC Jack Kilby Award for significant advances in analog mixed-signal circuits. In 2022, I was a co-recipient for the RFIC 2022 Best Student Paper Award for my work with students in my laboratory on a circuit prototype. From 2021 to 2023, I was the IEEE Solid-State Circuit

Society (SSCS) Distinguished Lecturer. I have been invited to offer seminars, workshops, keynote talks, and panel discussions in major IEEE conferences, IEEE SSCS chapters, and universities on a range of topics covering fundamental to advanced integrated circuit design and how to improve the performance and/or reduce cost of a circuit or system. In 2019, I was also the recipient of Qualcomm Faculty Award, and in 2014 both the NSF Faculty Early Career Development (CAREER) Award and DARPA Young Faculty Award (YFA). As a student, I was the recipient of a UC Regents' Fellowship at Berkeley in 2000 and the Analog Devices Outstanding Student Award for IC design in 2006.

234. I have extensive experience and familiarity with integrated circuit (IC) design for wireline communication as well as relevant signal processing techniques for such communication systems. As a principle investigator, I have conducted various research projects innovating and prototyping the most critical circuit blocks of wireline transceivers. I have also served in various review panels (e.g. NSF panels, IEEE conferences and journals), reviewing research papers related to wireline or serial link or SerDes circuit, system and signal processing techniques. Lastly, I have been teaching various circuit courses at USC. I teach my students a wide range of circuit and signal processing techniques related to this wireline area, so that they learn how those circuit components can be designed and implemented.

A. Compensation

235. For my efforts in connection with the preparation of this declaration I have been compensated at my standard rate for this type of consulting activity. My compensation is in no way contingent on the results of these or any other proceedings relating to the above-captioned patent.

B. Materials and Other Information Considered

236. I have considered information from various sources in forming my opinions. I have reviewed and considered each of the exhibits listed in the attached Appendix 2 (Materials Considered in the Preparation of This Declaration) in forming my opinions.

VII. Understanding of the Law

237. I am not an attorney. In forming my opinions in this Declaration, I applied the relevant legal principles provided to me by counsel, which are summarized in Appendix 4.

VIII. Reservation of Rights

238. My opinions are based upon the information that I have considered to date. I am unaware of any evidence of secondary considerations with respect to the '111 Patent that would render any of the challenged claims non-obvious. I reserve the right, however, to supplement my opinions in the future to respond to any arguments raised by the owner of the '111 Patent and to take into account new information that becomes available to me.

239. I declare that all statements made herein of my knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed on August 8, 2025.

By:



Michael S. Chen

APPENDIX 1: CURRICULUM VITAE OF MICHAEL S. CHEN

Mike Shuo-Wei Chen

3737 Watt Way, PHE 622, Los Angeles, CA 90089
Email: shuowei@gmail.com (Alt: swchen@usc.edu)
Phone: 510-517-0052

Curriculum Vitae

EDUCATION

University of California, Berkeley 2002 – 2006

- PhD in Electrical Engineering and Computer Science
- Dissertation: “High-Speed, Low-Power A/D Converter for a Low-Cost UWB Sub-sampling Radio”
- Advisor: Robert W. Brodersen

University of California, Berkeley 2000 – 2002

- Master of Science in Electrical Engineering and Computer Science
- Thesis: “Ultra Wideband Baseband Design and Implementation”

National Taiwan University 1994 – 1998

- Bachelor of Science in Electrical Engineering
- Top 1 student for upper division

ACADEMIC EXPERIENCE

- Full Professor, Department of Electrical Engineering, University of Southern California, 2021 – Present.
- Associate Professor, Department of Electrical Engineering, University of Southern California, 2017 – 2021.
- Colleen and Roberto Padovani Early Career Chair, USC, 2014 – 2021
- Assistant Professor, Department of Electrical Engineering, University of Southern California, 2011 – 2017.
- Graduate Student Researcher, Berkeley Wireless Research Center, Berkeley, CA, 2001 – 2006.

WORK EXPERIENCE

- Analog Senior Technical Staff, Atheros Communications, Santa Clara, CA, 2006–2010.
-Designed RF, analog mixed-signal circuits for various wireless products

CONSULTING EXPERIENCE

- Legal Consulting
(2019-2022): Xilinx-ADI litigation

Legal consultant, working with counsels to prepare/review legal documents, non-infringement arguments, deposition, source codes, PTAB, etc.

(2023-Present): Aptiv-Microchip

Legal expert, working with counsel on IPR expert declaration, source code review for litigation, etc.

(2023-2024): ARM-Qualcomm litigation

Legal expert, working on source code review, expert report, and deposition, testifying expert during trial

(2025-Present): Apple-Vampire Labs

Legal expert, IPR

(2025-Present): BOE – Samsung

Legal expert, IPR

(2025-Present): Marvell – Credo

This matter

(2025-Present): Renesas-Monterey Research

Legal expert, patent re-examination

- Technical Consulting:
Qualcomm, Samsung, BAE system, Tetramem → Provide technical advices on various electronic projects for the client.

HONORS and AWARDS

- IEEE Fellow 2024
- ISSCC 2022 Jack Kilby Award (co-recipient)
- RFIC 2022 Best Student Paper Award – First Place (co-recipient)
- IEEE Solid-State Circuit Society (SSCS) Distinguished Lecturer 2021-2023.
- Qualcomm Faculty Award, 2019.
- DARPA Young Faculty Award (YFA), 2014.
- NSF Faculty Early Career Development (CAREER) Award, 2014.
- Analog Devices Outstanding Student Award, 2006
- UC Regents Fellowship, UC Berkeley, 2000
- Lin's Foundation Award for Top 1 engineering college student over the academic year, 1996-1997
- Presidential Awards, National Taiwan University, 1996-1998
- Member of National Mathematics Team of Taiwan, awarded with Honourable Mention in Asian Pacific Mathematics Olympiad, 1994.

AREA OF INTEREST

I worked on various Analog Mixed-Signal, VLSI, SoC and RF circuit architectures/techniques, covering a wide range of topics/areas in academia and industry, including

- Data Converters (e.g. ADC and DAC)
- RF/mm-Wave Wireless Communications (e.g. WiFi, 5G/6G, Bluetooth, GPS, etc.)

- I/O, Wireline Communications (e.g. SerDes, CDR, Clock generation)
- Computing (e.g. AI/ML accelerators, neuromorphic processor)
- Design Methodology (AI-ML assisted CAD).
- Digital VLSI, SOC, RTL
- Analog Mixed-Signal Circuit
- Power management circuit

PUBLICATIONS

Conferences

1. S. Mahapatra, M. Ayes, C. Yang, M. Palaria, S. Su, A. Zhang, **M. S.-W. Chen**, "A 24GHz Direct Digital Transmitter Using Multiphase Subharmonic Switching PA Achieving 3.2Gbps Data Rate and -30.8dB EVM in 65nm CMOS", IEEE International Solid-State Circuits Conference (ISSCC), Feb. 2025.
2. C. Yang, S. Su, M. Ayes, S. Mahapatra, M. Hamada, V. Chenna, H. Hashemi, **M. S.-W. Chen**, "A Blocker Tolerant Receiver with VCO-based Non-Uniform Multi-Level Time-Approximation Filter With -36dB EVM in 28nm CMOS", IEEE International Solid-State Circuits Conference (ISSCC), Feb. 2025. (to be presented)
3. J. Liu, A. Shabra, S. Ho, G. Manganaro, **M. S.-W. Chen**, "A 16GS/s 10b Time-domain ADC using Pipelined-SAR TDC with Delay Variability Compensation and Background Calibration Achieving 153.8dB FoM in 4nm CMOS" in *IEEE Symposium on VLSI Technology and Circuits (VLSI Technol. Circuits)*, June 2024.
4. Q. Zhang*, S. Su*, B. R. Biswas, S. Gupta, **M. S.-W. Chen**, "Synthesizable 10-bit Stochastic TDC Using Common-Mode Time Dithering and Passive Approximate Adder With 0.012mm² Active Area in 12nm FinFET," in *IEEE Symposium on VLSI Technology and Circuits (VLSI Technol. Circuits)*, June 2024. (* equal contribution)
5. Q. Zhang*, S. Su*, Z. Liu*, H.-C. Cheng, Z. Qiu, M. Palaria, J. Ye, D. Meng, B. Chen, S. Hossain, W. Wu, **M. S.-W. Chen**, "A Stochastic Analog SAT Solver in 65nm CMOS Achieving 6.6 μ s Average Solution Time with 100% Solvability for Hard 3-SAT Problems," in *IEEE Symposium on VLSI Technology and Circuits (VLSI Technol. Circuits)*, June 2024. (* equal contribution)
6. M. Ayes, S. Mahapatra, C. Yang and **M. S.-W. Chen**, "A 0.072-mm² 18-21GHz Non-Uniform Sub-Sampling Receiver with a Non-Uniform Discrete-Time FIR Filter Achieving 42dB Blocker Rejection in 28nm CMOS", *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2024.
7. M. Palaria, S. Su, H.-C. Cheng, R. Rasul, Q. Zhang, S. Mahapatra, C.-F. Law, S. Hossain, R. Bena, W. Wu, Q. Nguyen, M. S.-W. Chen, "Analog Kalman Filter with Integration and Digitization via a Shared Thyristor-Based VCO for Sensor Fusion in 65 nm CMOS", in *IEEE European Solid-State Circuits Conference (ESSCIRC)*, Sep. 2023.
8. H.-C. Cheng, S. Su, M. Palaria, Q. Zhang, C. Yang, S. Hossain, R. Bena, B. Chen, Z. Liu, J. Liu, R. Rasul, Q. Nguyen, W. Wu, **M. S.-W. Chen**, "A

- Memristor-Based Analog Accelerator for Solving Quadratic Programming problems,” in IEEE Custom Integrated Circuits Conference (CICC), April 2023.
9. S. Su*, Q. Zhang*, and **M. S.-W. Chen**, “A 2GS/s 8.5-Bit Time-Based ADC Using a Segmented Stochastic Flash TDC,” in 2023 IEEE Custom Integrated Circuits Conference (CICC), Apr. 2023 (* equal contribution)
 10. Q. Zhang, H.-C. Cheng, S. Su, and **M. S.-W. Chen**, “A Fractional-N Digital MDLL with Injection Error Scrambling and Background Third-Order DTC Delay Equalizer Achieving -67dBc Fractional Spur,” in IEEE International Solid-State Circuits Conference (ISSCC), Feb. 2023.
 11. Q. Zhang*, S. Su*, and **M. S.-W. Chen**, “A Cost-Efficient Fully Synthesizable Stochastic Time-to-Digital Converter Design Based on Integral Nonlinearity Scrambling,” in 2022 59th ACM/EDAC/IEEE Design Automation Conference (DAC), July 2022. (*contributed equally to this work)
 12. C. Yang, S. Su and **Mike Chen**, “A Millimeter-Wave Mixer-First Receiver with Non-Uniform Time-Approximation Filter Achieving $>45\text{-dB}$ Blocker Rejection,” in IEEE Radio Frequency Integrated Circuits Symposium (RFIC), June 2022. (Best Student Paper Award – First Place)
 13. S. Su, and **M. S.-W. Chen**, “High-Speed Digital-to-Analog Converter Design Towards High Dynamic Range,” (invited paper) *IEEE Custom Integrated Circuits Conference (CICC)*, Apr. 2022.
 14. J. Liu, M. Hassanpourghadi, and **M. S.-W. Chen**, “A 10GS/s 8b 25fJ/c-s 2850um² Two-Step Time-domain ADC Using Delay-Tracking Pipelined-SAR TDC with 500fs Time Step in 14nm CMOS Technology“, *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2022.
 15. S. Su, Q. Zhang, M. Hassanpourghadi, J. Liu, R.A. Rasul, and **M. S.-W. Chen**, “Analog/Mixed-Signal Circuit Synthesis Enabled by the Advancements of Circuit Architectures and Machine Learning Algorithms,” (invited paper) *Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2022.
 16. S. Su, Q. Zhang, J. Liu, M. Hassanpourghadi, R.A. Rasul, and **M. S.-W. Chen**, “TAFA: Design Automation of Analog Mixed-Signal FIR Filters Using Time Approximation Architecture,” *Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2022.
 17. J. Liu, S. Su, M. Madhusudan, M. Hassanpourghadi, S. Saunders, Q. Zhang, R. Rasul, Y. Li, J. Hu, A. Kumar, S. S. Sapatnekar, R. Harjani, A. Levi, S. Gupta and **M. S.-W. Chen**, “From Specification to Silicon: Towards Analog/Mixed-Signal Design Automation using Surrogate NN Models with Transfer Learning“, *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, November 2021.
 18. M. Hassanpourghadi, S. Su, R.A. Rasul, J. Liu, Q. Zhang, and **M. S.-W. Chen**, “Circuit Connectivity Inspired Neural Network for Analog Mixed-Signal Functional Modeling,” *58th ACM/EDAC/IEEE Design Automation Conference (DAC)*, Dec. 2021.(to appear)
 19. R.A. Rasul, and **M. S.-W. Chen**, “A 128×128 SRAM Macro with Embedded Matrix-Vector Multiplication Exploiting Passive Gain via MOS Capacitor for Machine Learning Application,” *IEEE Custom Integrated Circuits Conference (CICC)*, Apr. 2021.
 20. A. Zhang, M. Ayesh, S. Mahapatra, and **M. S.-W. Chen**, “A 24-28 GHz Concurrent Harmonic and Subharmonic Tuning Class E/F_{2,2/3} Subharmonic

- Switching Power Amplifier Achieving Peak/PBO Efficiency Enhancement,” *IEEE Custom Integrated Circuits Conference (CICC)*, Apr. 2021.
21. Q. Zhang, S. Su, C.-R. Ho, and **M. S.-W. Chen**, “A Fractional-N Digital MDLL with Background Two-Point DTC Calibration Achieving -60dBc Fractional Spur,” *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2021.
 22. A. Zhang, C. Yang, M. Ayes, and **M. S.-W. Chen**, “A 5-6 GHz Current-Mode Subharmonic Switching Digital Power Amplifier for Enhancing Power Back-off Efficiency,” *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2021.
 23. Q. Zhang, S. Su, J. Liu and **M. S.-W. Chen**, “CEPA: CNN-based Early Performance Assertion Scheme for Analog and Mixed-Signal Circuit Simulation,” in *2020 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, November 2020.
 24. J. Liu, M. Hassanpourghadi, Q. Zhang, S. Su and **M.S.-W. Chen**, “Transfer Learning with Bayesian Optimization-Aided Sampling for Efficient AMS Circuit Modeling,” in *2020 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, November 2020.
 25. C. Yang, M. Ayes, A Zhang, T.F. Wu, **M. S.-W. Chen**, “A 29-mW 26.88-GHz Non-Uniform Sub-Sampling Receiver Front-End Enabling Spectral Alias Spreading,” *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, June, 2020.
 26. S. Su and **M. S.-W. Chen**, “A SAW-Less Direct-Digital RF Modulator with Tri-Level Time-Approximation Filter and Reconfigurable Dual-Band Delta-Sigma Modulation,” *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2020.
 27. T.-F. Wu and **M. S.-W. Chen**, “A 40MHz-BW 76.2dB/78.0dB SNDR/DR noise-shaping nonuniform sampling ADC with single phase-domain level crossing and embedded nonuniform digital signal processor in 28nm CMOS,” *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2020.
 28. S. Su and **M. S-W Chen**, “A 1–5GHz Direct-Digital RF Modulator with an Embedded Time-Approximation Filter Achieving -43dB EVM at 1024 QAM,” *IEEE Symposia on VLSI Technology and Circuits (VLSIC)*, June 2019.
 29. J.W. Nam, **M.S-W Chen**, “A 12.8-Gbaud ADC-based NRZ/PAM4 Receiver with Embedded Tunable IIR Equalization Filter Achieving 2.43-pJ/b in 65nm CMOS,” *IEEE Custom Integrated Circuits Conference (CICC)*, April 2019.
 30. M. Hassanpourghadi, **M.S-W Chen**, “A 2-way 7.3-bit 10 GS/s Time-based Folding ADC with Passive Pulse-Shrinking Cells,” *IEEE Custom Integrated Circuits Conference (CICC)*, April 2019.
 31. A. Zhang and **M. S-W Chen**, “A Watt-Level Phase-Interleaved Multi-Subharmonic Switching Digital Power Amplifier Achieving 31.4% Average Drain Efficiency,” *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2019.
 32. A. Zhang, **M. S-W Chen**, “A Sub-Harmonic Switching Digital Power Amplifier with Hybrid Class-G Operation for Enhancing Power Back-off Efficiency,” *IEEE Symposia on VLSI Technology and Circuits (VLSIC)*, 2018.

33. T.F. Wu, **M. S-W Chen**, "A 200MHz-BW 0.13mm² 62dB-DR VCO-Based Non-Uniform Sampling ADC with Phase-Domain Level Crossing in 65nm CMOS," *IEEE Custom Integrated Circuits Conference (CICC)*, 2018.
34. S. Su and **M. S-W Chen**, "A 16-bit 12GS/s Single/Dual-Rate DAC with Successive Bandpass Delta-Sigma Modulator Achieving <-67dBc IM3 within DC to 6GHz Tunable Passbands," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2018.
35. C.R. Ho, **M. S-W Chen**, "A digital frequency synthesizer with dither-assisted pulling mitigation for simultaneous DCO and reference path coupling," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2018.
36. C.R. Ho, **M. S-W Chen**, "A fractional-N digital PLL with background dither noise cancellation loop achieving <-62.5dBc worst-case near-carrier fractional spur in 65nm CMOS," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2018.
37. R. Rasul, P. Teimouri, and **M. S-W Chen**, "A Time Multiplexed Network Architecture for Large-Scale Neuromorphic Computing," *IEEE MWSCAS*, August 2017.
38. C.R. Ho, **M. S-W Chen**, "Interference-Induced DCO Spur Mitigation for Digital Phase Locked Loop in 65-nm CMOS," *IEEE European Solid-State Circuits Conference (ESSCIRC)*, Sep. 2016.
39. J.W. Nam, M. Hassanpourghadi, A. Zhang, **M.S-W Chen**, "A 12-bit 1.6 GS/s Interleaved SAR ADC with Dual Reference Shifting and Interpolation Achieving 17.8 fJ/conv-step in 65nm CMOS," *IEEE Symposia on VLSI Technology and Circuits (VLSIC)*, June 2016.
40. C.R. Ho, **M. S-W Chen**, "A Digital PLL with Feedforward Multi-Tone Spur Cancellation Loop Achieving <-73dBc Fractional Spur and <-110dBc Reference Spur in 65nm CMOS," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2016.
41. S. Su, **M. S-W. Chen**, "A 12b 2GS/s Dual-Rate Hybrid DAC with Pulsed Timing-Error Pre-Distortion and In-Band Noise Cancellation Achieving >74dBc SFDR up to 1GHz in 65nm CMOS," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2016.
42. T.F. Wu, C.R. Ho, **M. S.-W. Chen**, "A Flash-Based Non-Uniform Sampling ADC Enabling Digital Anti-Aliasing Filter in 65nm CMOS," *IEEE Custom Integrated Circuits Conference (CICC)*, Sep 2015.
43. S. Su, T. Tsai, P. Sharma, **M. S.-W. Chen**, "A 12-bit Hybrid DAC with 8GS/s Unrolled Pipeline Delta-Sigma Modulator achieving >75dB SFDR over 500MHz in 65nm CMOS," *IEEE Symposia on VLSI Technology and Circuits (VLSIC)*, June 2014.
44. C.R. Ho, **M. S.-W. Chen**, "A Fractional-N DPLL with Adaptive Spur Cancellation and Calibration-Free Injection-Locked TDC in 65nm CMOS," to be presented at *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, June, 2014.
45. J.W. Nam, D. Chiong, and **M. S.W. Chen**, "A 95-MS/s 11-bit 1.36-mW Asynchronous SAR ADC with Embedded Passive Gain in 65nm CMOS," *IEEE Custom Integrated Circuits Conference (CICC)*, Sep. 2013

46. P.K. Sharma, and **M. S.W. Chen**, "A 6b 800MS/s 3.62mW Nyquist AC-coupled VCO Based ADC in 65nm CMOS," *IEEE Custom Integrated Circuits Conference (CICC)*, Sep. 2013
47. **M. S.W. Chen**, "Trend of High-Speed SAR ADC towards RF Sampling," *10th International Conference on Sampling Theory & Applications (SampTA)* (Invited), July. 2013.
48. D. Hand, and **M. S.W. Chen**, "A Non-Uniform Sampling ADC Architecture with Embedded Alias-Free Asynchronous Filter," *Global Telecommunications Conference (GLOBECOM)*, Dec. 2012.
49. **M. S.W. Chen**, "Overhead Minimization Techniques for Digital Phase-Locked Loop Frequency Synthesizer," *IEEE MWSCAS (Invited Session)*, Aug. 2012.
50. **M.S.W. Chen**, **D. Su**, S. Mehta, "A Calibration-Free 800MHz Fractional-N Digital PLL with Embedded TDC," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2010.
51. Nathawad, L.; Zargari, M.; Samavati, H.; Mehta, S.; Kheirkhahi, A.; Chen, P.; Gong, K.; Vakili-Amini, B.; Hwang, J.; **Chen, M.S.W.**; Terrovitis, M.; Kaczynski, B.; Limotyrakis, S.; Mack, M.; Gan, H.; Lee, M.; Abdollahi-Alibeik, S.; Baytekin, B.; Onodera, K.; Mendis, S.; Chang, A.; Jen, S.; Su, D.; **Wooley, B.**, "A Dual-Band CMOS MIMO Radio SoC for IEEE 802.11n Wireless LAN," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2008.
52. A. Fort, **M. S.-W. Chen**, R.W. Brodersen, C. Desset, P. Wambacq, **L. Van Biesen**, "Impact of Sampling Jitter on Mostly-Digital Architectures for UWB Bio-Medical Applications," *IEEE International Conference on Communications (ICC)*, June 2007.
53. A. Fort, **M. S.-W. Chen**, C. Desset, P. Wambacq, **L. Van Biesen**, "Clock offset tracking for subsampling UWB architectures in a body area network," *IEEE International Conference on Ultra-Wideband (ICUWB)*, Sep. 2007.
54. **M. S.-W. Chen** and R. W. Brodersen, "Digital Complex Signal Processing Techniques for Impulse Radio," *Global Telecommunications Conference (GLOBECOM)*, Nov. 2006.
55. **M. S.-W. Chen** and R. W. Brodersen, "Implementation Considerations for a Sub-sampling Impulse Radio," *IEEE International Conference on Ultra-Wideband (ICUWB)*, Sep. 2006.
56. **M. S.-W. Chen** and R. W. Brodersen, "A 6b 600MS/s 5.3mW Asynchronous ADC in 0.13 μ m CMOS," *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2006.
57. D. Cabric, **M. S.-W. Chen**, D. Sobel, J. Yang, and R. Brodersen, "Future Wireless Systems: UWB, 60 GHz, and Cognitive Radios," (Invited paper) at *Custom Integrated Circuits Conference (CICC)*, Sep. 2005
58. **M. S.-W. Chen**, R. W. Brodersen, "The Impact of a Wideband Channel on UWB System Design," *Military Communication Conference (MILCOM)*, Nov. 2004.
59. **M. S.-W. Chen**, R. W. Brodersen, "A Subsampling UWB Radio Architecture by Analytic Signaling," *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, May 2004.
60. I. O'Donnell, **M. S.-W. Chen**, S. Wang, R. W. Brodersen, "An Integrated, Low-Power, Ultra-Wideband Transceiver Architecture for Low-Rate, Indoor Wireless Systems," *IEEE CAS Workshop on Wireless Communications and Networking*, Sep. 2002.

Journals

1. S. Su, Q. Zhang, B. R. Biswas, S. K. Gupta, **M. S.-W. Chen**, "Stochastic TDC Using Common-Mode Time Dithering and Passive Approximate Adders," *IEEE J. Solid-State Circuits*, April 2025.
2. M. Ayes, S. Mahapatra, C. Yang and **M. S.-W. Chen**, "A Blocker-Tolerant Non-Uniform Sub-Sampling Receiver With a Non-Uniform Discrete-Time FIR Filter," *IEEE J. Solid-State Circuits*, 2024.
3. W. Song, M. Rao, Y. Li, C. Li, Y. Zhuo, F. Cai, M. Wu, W. Yin, Z. Li, Q. Wei, S. Lee, H. Zhu, L. Gong, M. Barnell, Q. Wu, P. Beerel, **M. S.-W. Chen**, N. Ge, M. Hu, Q. Xia, J. Yang, "Programming memristor arrays with arbitrarily high precision for analog computing," *Science*, 2024.
4. Q. Zhang, H.-C. Cheng, S. Su, and **M. S.-W. Chen**, "Fractional-N Digital MDLL with Injection-Error Scrambling and Calibration," *IEEE J. Solid-State Circuits*, 2023.
5. C. Yang, S. Su, and **M. S.-W. Chen**, "Millimeter-Wave Receiver with Non-Uniform Time-Approximation Filter," *IEEE J. Solid-State Circuits*, 2023 (in press).
6. J. Liu, M. Hassanpourghadi, and **M. S.-W. Chen**, "A 10-GS/s 8-bit 2850- μm^2 Two-Step Time-Domain ADC With Speed and Efficiency Enhanced by the Delay-Tracking Pipelined-SAR TDC," *IEEE J. Solid-State Circuits*, 2022.
7. S. Su and **M. S.-W. Chen**, "SAW-Less Direct RF Transmitter with Multi-Mode Noise Shaping and Tri-Level Time-Approximation Filter," *IEEE J. Solid-State Circuits*, Mar. 2022.
8. Q. Zhang, S. Su, C.-R. Ho, and **M. S.-W. Chen**, "A Fractional-N Digital MDLL With Background Two-Point DTC Calibration," (Invited) *IEEE J. Solid-State Circuits (JSSC)*, Jan. 2022.
9. M. Hassanpourghadi, R. A. Rasul, and **M. S. W. Chen**, "A Module-Linking Graph Assisted Hybrid Optimization Framework for Custom Analog and Mixed-Signal Circuit Parameter Synthesis" *ACM Transactions on Design Automation of Electronic Systems*, June 2021.
10. X. Yan, J. Ma, T. Wu, A. Zhang, J. -B. Wu, M. Chin, Z. Zhang, M. Dubey, W. Wu, **M. S.-W. Chen**, J. Guo, H. Wang "Reconfigurable Stochastic Neurons Based on Tin Oxide/MoS₂ Hetero-memristors for Simulated Annealing and the Boltzmann Machine" *Nature Communications*, 12, Article number: 5710 (2021)
11. S. Su, **M. S.-W. Chen**, "A Time-approximation Filter for Direct RF Transmitter," *IEEE J. Solid-State Circuits (JSSC)* 2021.
12. J.W. Nam, **M. S.-W. Chen**, "A 12.8-Gbaud ADC-based Wireline Receiver with Embedded IIR Equalizer," (Invited) *IEEE J. Solid-State Circuits (JSSC)*, Mar. 2020.
13. A. Zhang, **M. S.-W. Chen**, "A Watt-Level Phase-Interleaved Multi-Subharmonic Switching Digital Power Amplifier," (Invited) *IEEE J. Solid-State Circuits (JSSC)* Nov., 2019.
14. A. Zhang, **M. S.-W. Chen**, "A Subharmonic Switching Digital Power Amplifier for Power Back-Off Efficiency Enhancement," (Invited) *IEEE J. Solid-State Circuits (JSSC)* Feb., 2019.

15. T.F. Wu, **M. S.-W. Chen**, "A VCO-Based Nonuniform Sampling ADC with Phase-Domain Level Crossing," (Invited) *IEEE J. Solid-State Circuits (JSSC)* Mar., 2019.
16. S. Su, **M. S.-W. Chen**, "A 16-bit 12GS/s Single/Dual-Rate DAC with a Successive Bandpass Delta-Sigma Modulator Achieving $<-67\text{dBc}$ IM3 within DC to 6GHz Tunable Passbands," (Invited) *IEEE J. Solid-State Circuits (JSSC)* Dec., 2018.
17. T.F. Wu, **M. S.-W. Chen**, "A Subranging-Based Nonuniform Sampling ADC With Sampling Event Filtering," *IEEE Solid-State Circuits Letters (SSC-L)* April, 2018.
18. J.W. Nam, M. Hassanpourghadi, A. Zhang, **M.S.-W. Chen**, "A 12-bit 1.6/3.2/6.4 GS/s 4 b/cycle Time-interleaved SAR ADC with Dual Reference Shifting and Interpolation," *IEEE J. Solid-State Circuits (JSSC)*, June, 2018.
19. T.F. Wu, C.R. Ho, **M. S.-W. Chen**, "A Flash-Based Non-Uniform Sampling ADC With Hybrid Quantization Enabling Digital Anti-Aliasing Filter," *IEEE J. Solid-State Circuits (JSSC)* Sep., 2017.
20. M. Hassanpourghadi, P.K. Sharma, and **M. S.-W. Chen**, "A 6-bit Nyquist AC-coupled VCO Based ADC at 800MS/s," *IEEE Transactions on Circuits and Systems (TCAS-I)* 2017.
21. C.R. Ho, **M. S.-W. Chen**, "A Digital PLL with Feedforward Multi-Tone Spur Cancellation Loop Achieving $<-73\text{dBc}$ Fractional Spur and $<-110\text{dBc}$ Reference Spur in 65nm CMOS," (Invited) accepted for *IEEE J. Solid-State Circuits (JSSC)* Dec., 2016.
22. S. Su, **M. S.-W. Chen**, "A 12b 2GS/s Dual-Rate Hybrid DAC with Pulsed Timing-Error Pre-Distortion and In-Band Noise Cancellation Achieving $>74\text{dBc}$ SFDR up to 1GHz in 65nm CMOS," (Invited) accepted for *IEEE J. Solid-State Circuits (JSSC)* Dec., 2016.
23. J.W. Nam, **M. S.-W. Chen**, "An Embedded Passive Gain Technique for Asynchronous SAR ADC Achieving >10 ENOB, 1.36 mW at 95MS/s in 65nm CMOS," accepted by *IEEE Transactions on Circuits and Systems (TCAS-I)* 2016
24. T.F. Wu, S. Dey, **M. S.-W. Chen**, "A Non-Uniform Sampling ADC Architecture with Reconfigurable Digital Anti-aliasing Filter," accepted by *IEEE Transactions on Circuits and Systems (TCAS-I)* 2016
25. C.R. Ho, **M. S.-W. Chen**, "A Fractional-N DPLL with Calibration-free Multi-phase Injection-locked TDC and Adaptive Single-tone Spur Cancellation Scheme," accepted by *IEEE Transactions on Circuits and Systems (TCAS-I)* 2016
26. S. Su, T. Tsai, P. Sharma, **M. S.W. Chen**, "A 12-bit 1GS/s Dual-Rate Hybrid DAC with an 8GS/s Unrolled Pipeline Delta-Sigma Modulator Achieving $>75\text{dB}$ SFDR over the Nyquist Band," (Invited) *IEEE J. Solid-State Circuits*, April, 2015.
27. Y. Cao, J. Velamala, K. Sutaria, **M. S.W. Chen**, J. Ahlbin, I. Esqueda, M. Bajura, M. Fritze, "Cross-Layer Modeling and Simulation of Circuit Reliability", *IEEE Transactions on CAD of Integrated Circuits and Systems*, Jan. 2014.
28. **M.S.W. Chen**, D. Su, S. Mehta, "A Calibration-Free 800MHz Fractional-N Digital PLL with Embedded TDC," (Invited) *IEEE Journal of Solid-State Circuits*, Dec. 2010.
29. Zargari, M.; Nathawad, L.Y.; Samavati, H.; Mehta, S.S.; Kheirkhahi, A.; Chen, P.; Gong, K.; Vakili-Amini, B.; Hwang, J.A.; **Chen, S.-W.M.**; Terrovitis, M.;

- Kaczynski, B.J.; Limotyrakis, S.; Mack, M.P.; Gan, H.; MeeLan Lee; Chang, R.T.; Dogan, H.; Abdollahi-Alibeik, S.; Baytekin, B.; Onodera, K.; Mendis, S.; Chang, A.; Rajavi, Y.; Jen, S.H.-M.; Su, D.K.; Wooley, B.A., "A Dual-Band CMOS MIMO Radio SoC for IEEE 802.11n Wireless LAN," *IEEE J. Solid-State Circuits*, Dec. 2008.
30. **M.S.W. Chen**, R. W. Brodersen, "A Subsampling Radio Architecture for Ultrawideband Communications," *IEEE Transactions on Signal Processing*, Oct. 2007.
 31. **M.S.W. Chen**, R. W. Brodersen, "A 6-bit 600-MS/s 5.3-mW Asynchronous ADC in 0.13-um CMOS," *IEEE J. Solid-State Circuits*, Dec. 2006.
 32. **M.S.W. Chen**, R. W. Brodersen, "A Subsampling UWB Impulse Radio Architecture Utilizing Analytic Signaling," (Invited Paper) *IEICE Trans. on Electronics*, June 2005.
 33. D. Cabric, **M. S.-W. Chen**, D. Sobel, S. Wang, J. Yang, and R. W. Brodersen, "Novel Radio Architectures for UWB, 60GHz, and Cognitive Wireless Systems," *EURASIP J. on Wireless Communcations and Networking*, 2006.

Magazine

1. **M. S.-W. Chen**, "Trend and New Opportunities in Digital PLL Design", (Invited) *IEEE Solid State Circuit Magazine*, 2020 winter issue.
2. C.R. Ho, **M. S.-W. Chen**, "Clock Generation in the future with Digital Signal Processing Technique for Mitigating Spur and Interference," *IEEE Microwave Magazine* 2019.
3. Sankaran, S.G.; Zargari, M.; Nathawad, L.Y.; Samavati, H.; Mehta, S.S.; Kheirkhahi, A.; Chen, P.; Ke Gong; Vakili-Amini, B.; Hwang, J.; **Chen, S.-W.M.**; Terrovitis, M.; Kaczynski, B.J.; Limotyrakis, S.; Mack, M.P.; Gan, H.; Lee, M.; Chang, R.T.; Dogan, H.; Abdollahi-Alibeik, S.; Baytekin, B.; Onodera, K.; Mendis, S.; Chang, A.; Rajavi, Y.; Jen, S.H.-M.; Su, D.K.; Wooley, B., "Design and implementation of a CMOS 802.11n soc –[integrated circuits for communications]," *Communications Magazine*, April 2009.
4. D. Cabric, I. O'Donnell, **M. S.-W. Chen**, and R.W. Brodersen, "Spectrum Sharing Radios," (Invited) *IEEE Circuits and Systems Magazine*, 2006.

Book Chapter

1. C.R. Ho, **M. S.-W. Chen**, "Fractional-N spur reduction techniques for DPLL," *Phase-Locked Frequency Generation and Clocking: Architectures and Circuits for Modern Wireless and Wireline Systems*, IET 2019 (under preparation).
2. **M. S.-W. Chen**, "Challenges and Emerging Trend of DSP Enabled Frequency Synthesizer," *Digitally-Assisted Analog and Analog-Assisted Digital IC Design*. Chapter4, 2015, Cambridge University Press.
3. **M. S.-W. Chen**, "Energy-Efficient ADC Topology Enabled with Asynchronous Techniques," *Circuits for Nanoscale: Communications, Imaging, and Sensing*. Chapter14, Sep. 2008.

Patent

1. **M.S.W. Chen**, C.R. Ho, "Adaptive spur cancellation techniques and multi-phase injection locked TDC for digital phase locked loop circuit," US patent #9941891, 2018.
2. **M. S.W. Chen**, "Non-uniform Sampling Analog to Digital Converter (ADC) Architecture with Digital Reconfigurable Anti-Aliasing Filter," filed for provisional patent application, No. 61/911,261, Dec., 2013.
3. **M. S.W. Chen**, D. Su, "Fractional and Integer PLL Architectures," US patent #8289086, 2012.

INVITED TALKS/TUTORIALS/WORKSHOPS

1. ITRI, Taiwan, "New Opportunities in Mixed-Signal ICs", June, 2011.
2. UMC, Taiwan, "New Opportunities in Mixed-Signal ICs", June, 2011.
3. HRL Laboratories, LLC: "Enhancing A-to-D Conversion Efficiency". July 2nd, 2012
4. "Analog-to-Digital Interface: A Time Approach" 2013 CMOS Emerging Technologies Research Symposium, Whistler, July 2013.
5. "Path towards High-Speed High-Resolution Data Converters with Diminishing Cost", Sep. 2013, Broadcom.
6. "Analog-Digital Interface Research, really?" Sep. 2013, Qualcomm-Atheros.
7. "Re-shape Future Mixed-Signal IC Design" Oct. 2013. Qualcomm.
8. "Path towards High-Speed High-Resolution Data Converters with Diminishing Cost", Nov. 2013, UT Dallas.
9. "Path towards High-Speed High-Resolution Data Converters with Diminishing Cost", Nov. 2013, Texas Instrument.
10. "Power efficient ADC Topologies towards RF Sampling" 2014 RFIC Tutorial.
11. "Path Towards Direct RF Synthesis: A Hybrid Digital-to-Analog Converter Architecture" 2014 IEEE MWSCAS Conference, Distinguished Speaker Series.
12. "Asynchronous SAR ADC: Past, Present and Beyond" 2014 IEEE MWSCAS Conference Tutorial
13. "Exploring Limits of Mixed-Signal ICs" Sep. 2014, MaxLinear Corp.
14. "Rethinking Analog-Digital Interface Circuit Architectures" Oct. 2014, IC Seminar, Columbia University, NY.
15. "Rethinking Analog-Digital Interface Circuit Architectures" Mar. 2015, IC Seminar, UC Berkeley, CA.
16. "Rethinking Analog-Digital Interface Circuit Architectures" Dec. 2015, IC Seminar, National Taiwan University(NTU)/IEEE Chapter, Taipei, Taiwan.
17. "Rethinking Analog-Digital Interface Circuit Architectures" Feb. 2016, IC Seminar, UT Austin, Texas.
18. "Asynchronous SAR ADC: Past, Present and Beyond", Feb. 2016, EE department Colloquium, UT Austin/IEEE Chapter.
19. "Rethinking Analog-Digital Interface Circuit Architectures" Feb. 2016, IC Seminar, University of Michigan, Ann Arbor.
20. "Rethinking Analog-Digital Boundary from Circuit to System Level towards Reconfigurability of Everything" Mar. 2016, SystemX Seminar, Stanford

21. "Rethinking Analog-Digital Interface Circuit Architectures" April. 2016, EE department Colloquium. Carnegie Mellon University (CMU).
22. "Rethinking Analog-Digital Interface Circuit Architectures" April. 2016, IC Seminar, University of California, Los Angeles (UCLA).
23. "IC Research overview," TSMC, 2017
24. "IC Research overview," Intel Lab, 2017
25. "IC Research overview," InPhi, 2017
26. "Advancing Low-Power High-Speed Analog-to-Digital Converters: An Asynchronous Design Approach", VLSI DAT tutorial, 2017
27. "Generic spur cancellation for digital PLL," Qualcomm 2017
28. "Evolutions of SAR ADC: from High Resolution to High Speed Regime," IEEE CICC 2018 tutorial.
29. "Emerging Opportunities in Analog Mixed-Signal Circuit Design Automation," ACM/IEEE ICCAD 2018 workshop.
30. "Analog-to-Digital Converter Architecture Opportunities in Emerging Wireless Systems," RFIC 2018 workshop
31. "How can hardware designers reclaim the spotlight?" moderator/co-organizer of ISSCC 2019 evening panel.
32. "Fundamentals of Analog-to-digital conversion," 2-day tutorial in Shanghai, China 2019.
33. "Digital Fractional-N Phase Locked Loop Design," tutorial, ISSCC Feb 2020
34. "Digital Fractional-N Phase Locked Loop Design," ISSCCedu Feb 2020
35. "Low Spur PLL architectures," MEAD tutorial in EPFL, 2021 (to appear)
36. "Low Spur PLL architectures and techniques," IEEE CICC 2020
37. "New Opportunities in Nonuniform Sampling," IEEE SSCS Webinar Oct 2020.
38. "High-Performance Digital-to-Analog Converter Design: A Path towards Digital Transmitter," IESD Keynote June 2021
39. "ADC Evolution via Architectural Rethinking: from Asynchronous SAR to Non-uniform Sampling ADC," IEEE SSCS Tainan Chapter, Oct 2021
40. "High-Performance Digital-to-Analog Converter Design: A Path towards Digital Transmitter," IEEE SSCS Atlanta Chapter, Nov. 2021
41. "Asynchronous SAR ADC: Past, Present and Beyond," IEEE Southern Alberta Chapter, Nov. 2021
42. "Non-Uniform Sampling Data Converters: A Journey to Uncharted Circuits and Systems" IEEE VLSI-DAT April 2022
43. "Asynchronous SAR ADC: Past, Present and Beyond," IEEE Swiss Chapter, May 2022
44. "New Opportunities in Nonuniform Sampling," IEEE Penang Chapter, July 2022
45. "Trend in Digital PLL Design and New Opportunities in Spur Cancellation," IEEE Southern Alberta Chapter, Sep. 2022
46. "High-Performance Digital-to-Analog Converter Design: A Path towards Digital Transmitter," IEEE Egypt Chapter Nov. 2022
47. "New Opportunities in Nonuniform Sampling," IEEE Taipei Chapter, Dec. 2022
48. "Trend in Digital PLL Design and New Opportunities in Spur Cancellation," IEEE Tainan Chapter, Dec. 2022

PROFESSIONAL SERVICES

Review Panel:

1. IEEE SSCS James D. Meindl Innovators Award
2. NSF Panelist
3. Samsung Research

Journal article review:

1. IEEE Journal of Solid-State Circuits
2. IEEE Solid-State Circuits Letters (SSC-L)
3. IEEE Transactions on Signal Processing
4. IEEE Transactions on Communications
5. IEEE Transactions on Circuits and Systems I
6. IEEE Transactions on Circuits and Systems II
7. IEEE Transactions on Vehicular Technology
8. IEEE Communications Letters
9. Journal of VLSI Signal Processing Systems

USC Department Services:

1. Munishian Series Committee
2. EFC
3. EE Festival reviewer

Conference TPC:

IEEE European Conference on Solid-State Circuits (ESSCIRC) (2022- present)
IEEE International Solid State Circuits Conference (ISSCC) (2018 - present)
IEEE Symposium on VLSI Circuits (VLSIC) (2017 - 2020)
IEEE Custom Integrated Circuits Conference (CICC) (2015 - 2019)
IEEE GlobalSIP (2014)

Organized/Participated Panel/Forum/Workshops for SSCS:

1. Forum: “Emerging Design Techniques for Data Converters” served as organizer at CICC 2017
2. Panel: “What can/should Circuit Designers do to Ride on the Wave of Machine Learning?” served as co-organizer and moderator at CICC 2018
3. Panel: “How can hardware designers reclaim the spotlight?” served as co-organizer and moderator at ISSCC 2019
4. Panel: “Favorite circuit design and testing mistakes of starting engineers?” served as co-organizer at ISSCC 2021
5. VLSI Symposia Mentoring Event June 2021
6. Panel: “How to choose career path, academia, industry, startups?” served as panelist at CICC 2022 (to appear)
7. Panel: “Open Source Systems, Circuits, and Design: Is It the Future?” served as panelist at CICC 2022 (to appear)

Associate Editor:

SSC-L, TCAS-II

Society Membership:
IEEE Fellow

APPENDIX 2: MATERIALS CONSIDERED IN THE PREPARATION OF THIS DECLARATION

Exhibit No.	Description
1001	U.S. Patent No. 11,032,111 to Sun et al. (the '111 patent)
1002	Certified Prosecution History for the '111 patent ("Prosecution History")
1004	U.S. Patent App. Pub. No. 2019/0288881 to Berke et al. ("Berke")
1005	U.S. Patent No. 8,516,238 to Cornelius, et al. ("Cornelius")
1006	U.S. Patent No. 10,715,357 to Ran ("Ran")
1007	U.S. Prov. App. Serial No. 62/722,517 ("Ran Provisional")
1008	Excerpts from Stauffer, High Speed Serdes Devices and Applications ("Stauffer")
1009	U.S. Patent App. Pub. No. 2014/0237301 to Mejia et al. ("Mejia")
1010	PCI Express Base Specification Revision 3.0 (Nov. 10, 2010)
1011	Serial ATA Revision Specification 3.1 (July 18, 2011)
1012	Excerpts of IEEE 802.3-2015 Standard
1013	Declaration of June Munford re: Stauffer
1014	Library of Congress Certification for Stauffer
1015	Internet Archive Declaration of Mina Ching re: Stauffer
1016	Declaration of Liliana Nunez re: Stauffer
1017	U.S. Patent App. Pub. No. 2013/0223506 to Kolze
1018	U.S. Patent App. Pub. No. 2010/0232493 to Thirumoorthy
1019	U.S. Patent App. Pub. No. 2013/0201316 to Binder et al.
1020	U.S. Patent No. 5,648,972 to Gharakhanian
1021	U.S. Prov. App. Serial No. 63467_62723701 ('111 Patent Provisional)
1022	Excerpts from Gasca, et al., CMOS Continuous-Time Adaptive Equalizers for High-Speed Serial Links
1023	U.S. Patent No. 7,233,617 to Gorecki
1024	U.S. Patent App. Pub. No. 2003/0035497 to Gorecki et al.
1025	TE Connectivity, SFP+ High Speed Copper Cable Assemblies
1026	Cisco, Cisco 40GBASE QSFP Modules Data Sheet
1027	Wang & Bovik, Mean Squared Error: Love It or Leave It?
1028	U.S. Patent No. 8,665,941 to Eliaz
1029	U.S. Patent No. 8,600,039 to Chen et al.
1030	U.S. Patent No. 5,392,315 to Laud

APPENDIX 3: CHALLENGED CLAIMS

[1pre]. A SerDes communications method that comprises, in a transceiver:

[1a]. selecting one of multiple registers to specify initial pre-equalizer coefficient values, each of the multiple registers corresponding to a different channel model;

[1b]. updating the initial pre-equalizer coefficient values during a training phase; and

[1c]. using the updated pre-equalizer coefficient values to convey a transmit data stream.

2. The method of claim 1, wherein at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.

3. The method of claim 1, wherein at least some of the different channel models presume different types of receiver equalization.

4. The method of claim 1, wherein said selecting includes using the initial pre-equalizer coefficient values to determine a performance characteristic for each of the multiple registers.

5. The method of claim 4, wherein the performance characteristic is an error signal energy.

6. The method of claim 4, wherein the performance characteristic is a bit error rate.

7. The method of claim 4, further comprising updating the selected register with the updated pre-equalizer coefficient values.

[8pre]. A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

[8a]. one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

[8b]. a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.

9. The communications link of claim 8, wherein at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.

10. The communications link of claim 8, wherein at least some of the different channel models presume different types of receiver equalization.

[11pre]. A chip-to-module communications link that comprises a port connector coupling a port transceiver to a pluggable module transceiver, the pluggable module transceiver including:

[11a]. one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to the port transceiver; and

[11b]. a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values for which the port transceiver determines a performance characteristic,

[11c]. the controller using one of the registers selected by the port transceiver to specify the initial coefficient values for the one or more transmit filters.

12. The communications link of claim 11, wherein the performance characteristic is an error signal energy.

13. The communications link of claim 11, wherein the performance characteristic is a bit error rate.

14. The communications link of claim 11, wherein the port transceiver generates updates for the coefficient values of the one or more transmit filters.

15. The communications link of claim 14, wherein the controller saves updated coefficient values in one of the registers.

[16pre]. A pluggable module transceiver including:

[16a]. one or more transmit filters to each pre-equalize a corresponding serial symbol stream being transmitted to a port transceiver; and

[16b]. a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model, the controller using one of the registers to set initial coefficient values for the one or more transmit filters.

17. The pluggable module transceiver of claim 16, wherein the port transceiver determines a performance characteristic for each register.

18. The pluggable module transceiver of claim 17, wherein the performance characteristic is an error signal energy.

19. The pluggable module transceiver of claim 17, wherein the performance characteristic is a bit error rate.

APPENDIX 4: UNDERSTANDING OF THE LAW

I have applied the following legal principles provided to me by counsel in arriving at the opinions set forth in this report.

Legal Standard for Prior Art

I am not an attorney. I have been informed by attorneys of the relevant legal principles and have applied them to arrive at the opinions set forth in this declaration.

I understand that the petitioner for inter partes review may request the cancelation of one or more claims of a patent based on grounds available under 35 U.S.C. § 102 and 35 U.S.C. § 103 using prior art that consists of patents and printed publications.

Anticipation and Prior Art

I understand that § 102 specifies when a challenged claim is invalid for lacking novelty over the prior art, and that this concept is also known as “anticipation.” I understand that a prior art reference anticipates a challenged claim, and thus renders it invalid by anticipation, if all elements of the challenged claim are disclosed in the prior art reference. I understand the disclosure in the prior art reference can be either explicit or inherent, meaning it is necessarily present or implied. I understand that the prior art reference does not have to use the same words as the challenged claim, but all of the requirements of the claim must

be disclosed so that a person of ordinary skill in the art could make and use the claimed subject-matter.

In addition, I understand that § 102 also defines what is available for use as a prior art reference to a challenged claim.

Under § 102(a)(1), a challenged claim is anticipated if it was patented, described in a printed publication, or in public use, on sale, or otherwise available to the public before the effective filing date of the challenged claimed invention.

Under § 102(a)(2), a challenged claim is anticipated if it was described in an issued patent, or in an application for patent that was published by the USPTO, in which the patent or application, as the case may be, names another inventor and was effectively filed before the effective filing date of the challenged claimed invention.

I understand that the filing date of a patent is generally the filing date of the application filed in the United States that issued as the patent. However, I understand that a patent may be granted an earlier effective filing date if the patent owner properly claimed priority to an earlier patent application.

I understand that when a challenged claim covers several structures, either generically or as alternatives, the claim is deemed anticipated if any of the structures within the scope of the claim is found in the prior art reference.

I understand that when a challenged claim requires selection of an element from a list of alternatives, the prior art teaches the element if one of the alternatives is taught by the prior art.

Legal Standard for Obviousness

I understand that even if a challenged claim is not anticipated, it is still invalid if the differences between the claimed subject matter and the prior art are such that the claimed subject matter would have been obvious to a person of ordinary skill in the pertinent art at the time the alleged invention.

I understand that obviousness must be determined with respect to the challenged claim as a whole.

I understand that one cannot rely on hindsight in deciding whether a claim is obvious.

I also understand that an obviousness analysis includes the consideration of factors such as (1) the scope and content of the prior art, (2) the differences between the prior art and the challenged claim, (3) the level of ordinary skill in the pertinent art, and (4) “secondary” or “objective” evidence of non-obviousness.

Secondary or objective evidence of non-obviousness includes evidence of: (1) a long felt but unmet need in the prior art that was satisfied by the claimed invention; (2) commercial success or the lack of commercial success of the claimed invention; (3) unexpected results achieved by the claimed invention; (4)

praise of the claimed invention by others skilled in the art; (5) taking of licenses under the patent by others; (6) deliberate copying of the claimed invention; and (7) contemporaneous and independent invention by others. However, I understand that there must be a relationship between any secondary evidence of non-obviousness and the claimed invention.

I understand that a challenged claim can be invalid for obviousness over a combination of prior art references if a reason existed (at the time of the alleged invention) that would have prompted a person of ordinary skill in the art to combine elements of the prior art in the manner required by the challenged claim. I understand that this requirement is also referred to as a “motivation to combine,” “suggestion to combine,” or “reason to combine,” and that there are several rationales that meet this requirement.

I understand that the prior art references themselves may provide a motivation to combine, but other times simple common sense can link two or more prior art references. I further understand that obviousness analysis recognizes that market demand, rather than scientific literature, often drives innovation, and that a motivation to combine references may come from market forces.

I understand obviousness to include, for instance, scenarios where known techniques are simply applied to other devices, systems, or processes to improve them in an expected or known way. I also understand that practical and common-

sense considerations should be applied in a proper obviousness analysis. For instance, familiar items may have obvious uses beyond their primary purposes.

I understand that the combination of familiar elements according to known methods is obvious when it yields predictable results. For instance, obviousness bars patentability of a predictable variation of a technique even if the technique originated in another field of endeavor. This is because design incentives and other market forces can prompt variations of it, and predictable variations are not the product of innovation, but rather ordinary skill and common sense.

I understand that a particular combination may be obvious if it was obvious to try the combination. For example, when there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. This would result in something obvious because the result is the product not of innovation but of ordinary skill and common sense. However, I understand that it may not be obvious to try a combination when it involves unpredictable technologies.

It is further my understanding that a proper obviousness analysis focuses on what was known or obvious to a person of ordinary skill in the art, not just the patentee. Accordingly, I understand that any need or problem known in the field of

endeavor at the time of invention and addressed by the patent can provide a reason for combining the elements in the manner claimed.

It is my understanding that the Manual of Patent Examining Procedure §2143 sets forth the following as exemplary rationales that support a conclusion of obviousness:

- Combining prior art elements according to known methods to yield predictable results;
- Simple substitution of one known element for another to obtain predictable results;
- Use of known technique to improve similar devices (methods, or products) in the same way;
- Applying a known technique to a known device (method, or product) ready for improvement to yield predictable results;
- Choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success;
- Known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art;

- Some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

A person of ordinary skill in the art looking to overcome a problem will often use the teachings of multiple publications together like pieces of a puzzle, even though the prior art does not necessarily fit perfectly together. Therefore, I understand that references for obviousness need not fit perfectly together like puzzle pieces. Instead, I understand that obviousness analysis takes into account inferences, creative steps, common sense, and practical logic and applications that a person of ordinary skill in the art would employ under the circumstances.

I understand that a claim can be obvious in light of a single reference, if the elements of the challenged claim that are not explicitly or inherently disclosed in the reference can be supplied by the common sense of one of skill in the art.

I understand that obviousness also bars the patentability of applying known or obvious design choices to the prior art. One cannot patent merely substituting one prior art element for another if the substitution can be made with predictable results. Likewise, combining prior art techniques that are interoperable with respect to one another is generally obvious and not patentable.

In order for a claim to be found invalid based upon a modification or combination of the prior art, there must be reasonable expectation that a person of ordinary skill would have successfully modified or combined the prior art to arrive at the claimed arrangement. This does not mean that it must be certain that a person of ordinary skill would have been successful – the law only requires that the person of ordinary skill in the art would have perceived a reasonable expectation of success in modifying or combining the prior art to arrive at the claimed invention.

In sum, my understanding is that obviousness invalidates claims that merely recite combinations of, or obvious variations of, prior art teachings using understanding and knowledge of one of skill in the art at the time and motivated by the general problem facing the inventor at the time. Under this analysis, the prior art references themselves, or any need or problem known in the field of endeavor at the time of the invention, can provide a reason for combining the elements of or attempting obvious variations on prior art references in the claimed manner.

Legal Standard for Claim Construction

I understand that before any invalidity analysis can be properly performed, the scope and meaning of the challenged claims must be determined by claim construction.

I understand that a patent may include two types of claims, independent claims and dependent claims. I understand that an independent claim stands alone

and includes only the limitations it recites. I understand that a dependent claim depends from an independent claim or another dependent claim. I understand that a dependent claim includes all the limitations that it recites in addition to the limitations recited in the claim (or claims) from which it depends.

In comparing the challenged claims to the prior art, I have carefully considered the patent and its file history in light of the understanding of a person of skill at the time of the alleged invention.

I understand that to determine how a person of ordinary skill would have understood a claim term, one should look to sources available at the time of the alleged invention that show what a person of skill in the art would have understood disputed claim language to mean. It is my understanding that this may include what is called “intrinsic” evidence as well as “extrinsic” evidence.

I understand that, in construing a claim term, one should primarily rely on intrinsic patent evidence, which includes the words of the claims themselves, the remainder of the patent specification, and the prosecution history. I understand that extrinsic evidence, which is evidence external to the patent and the prosecution history, may also be useful in interpreting patent claims when the intrinsic evidence itself is insufficient. I understand that extrinsic evidence may include principles, concepts, terms, and other resources available to those of skill in the art at the time of the invention.

I understand that words or terms should be given their ordinary and accepted meaning unless it appears that the inventors were using them to mean something else or something more specific. I understand that to determine whether a term has special meaning, the claims, the patent specification, and the prosecution history are particularly important, and may show that the inventor gave a term a particular definition or intentionally disclaimed, disavowed, or surrendered claim scope.

I understand that the claims of a patent define the scope of the rights conferred by the patent. I understand that because the claims point out and distinctly claim the subject matter which the inventors regard as their invention, claim construction analysis must begin with and is focused on the claim language itself. I understand that the context of the term within the claim as well as other claims of the patent can inform the meaning of a claim term. For example, because claim terms are normally used consistently throughout the patent, how a term is used in one claim can often inform the meaning of the same term in other claims. Differences among claims or claim terms can also be a useful guide in understanding the meaning of particular claim terms.

I understand that a claim term should be construed not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the entire specification. I understand that because the

specification is a primary basis for construing the claims, a correct construction must align with the specification.

I understand that the prosecution history of the patent as well as art incorporated by reference or otherwise cited during the prosecution history are also highly relevant in construing claim terms. For instance, art cited by or incorporated by reference may indicate how the inventor and others of skill in the art at the time of the invention understood certain terms and concepts. Additionally, the prosecution history may show that the inventors disclaimed or disavowed claim scope, or further explained the meaning of a claim term.

With regard to extrinsic evidence, I understand that all evidence external to the patent and prosecution history, including expert and inventor testimony, dictionaries, and learned treatises, can also be considered. For example, technical dictionaries may indicate how one of skill in the art used or understood the claim terms. However, I understand that extrinsic evidence is considered to be less reliable than intrinsic evidence, and for that reason is generally given less weight than intrinsic evidence.

I understand that in general, a term or phrase found in the introductory words or preamble of the claim, should be construed as a limitation if it recites essential structure or steps, or is necessary to give meaning to the claim. For instance, I understand preamble language may limit claim scope: (i) if dependence

on a preamble phrase for antecedent basis indicates a reliance on both the preamble and claim body to define the claimed invention; (ii) if reference to the preamble is necessary to understand limitations or terms in the claim body; or (iii) if the preamble recites additional structure or steps that the specification identifies as important.

On the other hand, I understand that a preamble term or phrase is not limiting where a challenged claim defines a structurally complete invention in the claim body and uses the preamble only to state a purpose or intended use for the invention. I understand that to make this determination, one should review the entire patent to gain an understanding of what the inventors claim they invented and intended to encompass in the claims.

I understand that 35 U.S.C. § 112(f) created an exception to the general rule of claim construction called a “means plus function” limitation. These types of terms and limitations should be interpreted to cover only the corresponding structure described in the specification, and equivalents thereof. I also understand that a limitation is presumed to be a means plus function limitation if (a) the claim limitation uses the phrase “means for”; (b) the “means for” is modified by functional language; and (c) the phrase “means for” is not modified by sufficient structure for achieving the specified function.

I understand that a structure is considered structurally equivalent to the corresponding structure identified in the specification only if the differences between them are insubstantial. For instance, if the structure performs the same function in substantially the same way to achieve substantially the same result. I further understand that a structural equivalent must have been available at the time of the issuance of the claim.