

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 IPR2025-01220
U.S. Patent No. 11,032,111

PATENT OWNER'S RESPONSE

Mail Stop "Patent Board"
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Contents

I.	Introduction.....	1
II.	The '111 Patent	4
III.	Petitioner's Challenge to the '111 Patent	7
	A. Berke (EX-1004)	7
	B. Cornelius (EX-1005)	10
	C. Ran (EX-1006)	11
	D. Stauffer (EX-1008).....	15
	E. Mejia (EX-1009)	16
IV.	Legal Standards	18
V.	Level of Skill in the Art	19
VI.	Claim Construction.....	19
VII.	The Challenged Claims are Not Unpatentable	20
	A. Ground 1: Berke in view of Cornelius (Claims 1, 3, 4, 6-8, 10- 11, 13-17, 19)	21
	1. Claim 1	21
	2. Claims 8 and 11	32
	3. Claim 16.....	37
	B. Ground 2: Berke in view of Ran (Claims 1, 3, 4, 6-8, 10-11, 13-17, 19)	37
	1. Claim 1	37
	2. Claim 8.....	37

3.	Claims 11.....	41
4.	Claim 16.....	44
C.	Ground 3: Berke in view of Cornelius and Stauffer (Claims 2, 5, 9, 12, 18) and Ground 4: Berk in view of Ran and Stauffer (Claims 2, 5, 9, 12, 18).....	44
1.	Claims 2 and 9.....	45
2.	Claims 5, 12, and 18	46
D.	Ground 5: Ran in view of Mejia and Stauffer (Claims 1- 19).....	50
1.	Claim 1	50
2.	Claim 8.....	55
3.	Claim 11	55
4.	Claim 16	58
E.	Grounds 1-5: Secondary Indicia of Non-obviousness Further Undermine Petitioner’s Challenges to the ’111 Patent	58
VIII.	Conclusion	59

TABLE OF AUTHORITIES

Cases

<i>Apple Inc. v. Samsung Elecs. Co., Ltd.</i> , 839 F.3d 1034 (Fed. Cir. 2016).....	58
<i>Graham v. John Deere Co.</i> , 383 U.S. 1 (1966)	19
<i>Institut Pasteur & Universite Pierre Et Marie Curie v. Focarino</i> , 738 F.3d 1337 (Fed. Cir. 2013).....	58
<i>Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.</i> , 821 F.3d 1359 (Fed. Cir. 2016).....	20
<i>KSR Int’l Co. v. Teleflex Inc.</i> , 550 U.S. 398 (2007)	18, 20
<i>Metalcraft of Mayville, Inc. v. Toro Co.</i> , 848 F.3d 1358 (Fed. Cir. 2017).....	29
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	19

Regulations

37 C.F.R. § 42.100(b).....	19
----------------------------	----

TABLE OF EXHIBITS

Exhibit	Description
2001	In re Certain Active Electrical Cables and Components Thereof, Inv. No. 337-TA-3814, Complaint (Mar. 13, 2025).
2002	“Credo Introduces PAM4 DSP for High Performance Data Centers and Enterprise Networks,” March 17, 2021, https://credosemi.com/credo-introduces-pam4-dsp-for-high-performance-data-centers-and-enterprise-networks/
2003	“AEC Applications,” https://credosemi.com/products/hiwire-aec/applications/
2004	“Credo Introduces 800G HiWire ZeroFlap AECs to Support AI Backend Networks,” October 10, 2024, available at https://s205.q4cdn.com/511065572/files/doc_news/2024/10/credo-introduces-800g-hiwire-zeroflap-aecs-to-support-ai-backend-networks.pdf
2005	“CREDO Receives Cabling Installation & Maintenance Innovators 2019 Gold Award for HiWire™ Active Electrical Cables (AEC),” October 1, 2019, https://credosemi.com/news/credo-receives-cabling-installation-maintenance-innovators-2019-gold-award-for-hiwire-active-electrical-cables-aec/
2006	“Credo Files AEC Patent Infringement Complaint Against Amphenol, Molex, TE Connectivity, and Volex with United States International Trade Commission,” March 13, 2025, https://investors.credosemi.com/news-events/news/news-details/2025/Credo-Files-AEC-Patent-Infringement-Complaint-Against-Amphenol-MoLEX-TE-Connectivity-and-VoLEX-with-United-States-International-Trade-Commission/default.aspx
2007	Order No. 12 Granting Marvell’s Motion to Disqualify Fish, 337-TA-1446, public version
2008	Bert Reiser & Ruthie Wu, Why the International Trade Commission is such an appealing forum for patent disputes,” Reuters (June 11, 2025), https://www.reuters.com/legal/legalindustry/why-international-trade-commission-is-such-an-appealing-forum-patent-disputes-2025-06-11/
2009	Order No. 19 Granting Motion to Amend Procedural Schedule, 337-TA-1446

Exhibit	Description
2010	Respondents' Initial Invalidity Contentions, In the Matter of Certain Active Electrical Cables and Components Thereof, Inv. No. 337-TA-1446, served June 26, 2025
2011	Exhibit C-12, 111 Patent Invalidity Chart to the ITC Respondents' Initial Invalidity Contentions, In the Matter of Certain Active Electrical Cables and Components Thereof, Inv. No. 337-TA-1446, served June 26, 2025
2012	Exhibit C-3, 111 Patent Invalidity Chart to the ITC Respondents' Initial Invalidity Contentions, In the Matter of Certain Active Electrical Cables and Components Thereof, Inv. No. 337-TA-1446, served June 26, 2025
2013	Exhibit C-11, 111 Patent Invalidity Chart to the ITC Respondents' Initial Invalidity Contentions, In the Matter of Certain Active Electrical Cables and Components Thereof, Inv. No. 337-TA-1446, served June 26, 2025
2014	USPTO, "FAQs for Interim Processes for PTAB Workload Management," available at https://www.uspto.gov/patents/ptab/faqs/interim-processes-workload-management
2015	Respondents' Final Invalidity Contentions, Oct. 24, 2025, Inv. No. 337-TA-1446
2016	Exhibit C-21 to the ITC Respondents' Final Invalidity Contentions, Inv. No. 337-TA-1446
2017	Declaration of Tim A. Williams, May 12, 2026
2018	Transcript of the May 8, 2026 Deposition of Dr. Michael Chen, Ph.D.
2019	Press Release: Credo and 3M Enter Into Patent License Agreement, Jan. 22, 2026
2020	Press Release: Credo and Amphenol Reach Settlement in Active Electrical Cable Patent Infringement Disputes, Aug. 14, 2025
2021	Press Release: Credo and Molex Reach Settlement in Active Electrical Cable Patent Infringement Disputes, Mar. 26, 2026
2022	Press Release: Credo Licenses Its Active Electrical Cable Patents To The Siemon Company, Nov. 24, 2025
2023	Press Release: Credo and TE Connectivity Reach Settlement in Active Electrical Cable Disputes, Mar. 27, 2026
2024	Press Release: Credo and Volex Reach Settlement in Active Electrical Cable Patent Infringement Dispute, Aug. 26, 2025

Exhibit	Description
2025	“Credo Announces Production Availability of HiWire Active Electrical Cables,” Press Release dated September 4, 2019, available at https://investors.credosemi.com/news-events/news/news-details/2019/Credo-Announces-Production-Availability-of-HiWire-Active-Electrical-Cables-09-04-2019/default.aspx
2026	Equity Research: Credo Technology Group Holding Ltd: Initiation of Research Coverage, William Blair, dated September 18, 2025, available at https://www.williamblair.com/News/Credo-Technology-Group-Holding-Ltd-Initiation
2027	“This Company Supplying AI Datacenters With Cables Just Minted Two New Billionaires,” Forbes December 1, 2025, available at https://www.forbes.com/sites/emilygarcia/2025/12/01/this-company-supplying-ai-datacenters-with-cables-just-minted-two-new-billionaires/

I. INTRODUCTION

Petitioner Marvell Semiconductor, Inc., challenges the patentability of U.S. Patent No. 11,032,111 (the “’111 patent”). The invention described and claimed in the ’111 patent, assigned to Patent Owner Credo Technology Group Limited, arose from Credo’s creation of a brand new category of Ethernet cables now known as Active Electrical Cables (“AECs”). Credo coined the term Active Electrical Cable when it released its first AEC in 2019. EX-2025. Until then, data centers typically used optical cables for high-speed connections between components in a server rack. Traditional Ethernet cables could not support the increased demand for data speed, leaving optical cables as the only choice. Pioneering an entirely new market, Credo introduced AECs that are less expensive to buy, require less power to operate, and perform more reliably than optical cables. EX-2026. Within months, Credo’s AECs won the 2019 Gold Award from the Cabling Installation & Maintenance Innovators. EX-2005. Since then, Forbes magazine reports that Credo’s distinctive purple AECs, pictured below, “have become indispensable in the AI infrastructure buildout.” EX-2027.



Recognizing Credo's success, competitors have since entered the AEC market with their own AECs, many of which are now licensed to Credo's patents. Petitioner designs and supplies components that may be used in AECs and is now challenging the '111 patent as well as two other Credo patents related to AECs. For multiple reasons, Petitioner has failed to meet its burden of showing that any claim of the '111 patent is unpatentable.

Petitioner challenges all nineteen claims of the '111 Patent, including independent claims 1, 8, 11, and 16. The Petition asserts five prior art references in various combinations across five grounds, and each ground fails for at least two reasons.

First, every challenged claim requires at least one feature that is totally absent from the asserted prior art. For independent claims 1, 8, and 16, every asserted prior art reference lacks at least the requirement of separate sets of equalization coefficients that correspond to different channel models. For independent claim 11, every asserted reference lacks at least the claimed performance characteristic determined by a port transceiver. Some of Petitioner's references are missing even more claimed features, as explained in detail below. Petitioner contorts the prior art in its effort to find these features, but they are nowhere to be found.

Second, Petitioner fails to explain *how* or *why* a POSITA would have been motivated to combine the prior art teaching within each ground. For example, Petitioner asserts that the prior art references come from similar fields of endeavor, but that is not enough. In another example, Petitioner proposes combinations that do not make sense. For example, Petitioner asserts that a POSITA would have been motivated to graft Mejia's preset selection procedures into the teachings of Ran. But Ran discloses its own preset selection procedures, so it has no need for Mejia's procedures. Petitioner makes similar arguments for the combination of Berke and Cornelius. In each of these combinations, Petitioner uses hindsight to reconstruct the claimed invention.

As explained in more detail below, Patent Owner respectfully submits that the Board should affirm the patentability of all nineteen claims of the '111 patent.

II. THE '111 PATENT

The '111 Patent claims improvements to serializer-deserializer (SerDes) communications, for example, using an “efficient adaptation interface for transmitter pre-equalizers” that enables “standards-suitable” data transmission over channels with varying “attenuation and dispersion.” EX-1001, 1:19-35; 1:39-2:16.

The '111 Patent explains that “as demand continues for ever-higher data rates [for data communications], 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, 1:19-35. Earlier techniques to address channel attenuation and dispersion “introduce[d] undesirable hardware complexity or cost.” EX-1001, 1:31-35. To provide a low cost and effective solution to, in part, address channel attenuation and dispersion, the '111 Patent discloses “SerDes pre-equalizer having adaptable preset coefficient registers.” EX-1001, title. The systems and methods described in the '111 Patent provide, among other benefits, improvement in signal quality while also reducing or eliminating the need for receiver-side equalization, thereby enabling digital filtering with reduced bit width, and reducing power consumption in high-speed communications links, including chip-to-module (C2M) channels. EX-1001, 6:40-54, 5:24-30.

For instance, the inventors discovered that desirable performance characteristics, such as improved efficiency and signal integrity, could be achieved by (i) providing multiple registers in a transceiver controller, each storing a set of preset initial pre-equalizer coefficient values corresponding to different channel models, (ii) allowing for selection of optimal starting point set of coefficients, (iii) adaptive updating of such coefficients during a training phase, and (iv) application of the updated coefficients to pre-equalize transmitted data streams, while minimizing power consumption, hardware complexity, and noise enhancement. *See, e.g.*, EX-1001, 11:53-14:11; 14:12-15:3; 6:40-54.

Further, the above-described adaptive pre-equalization process can be implemented in a pluggable module transceiver. *See, e.g.*, EX-1001, Abstract, 1:49-57, claims 8, 11, 16. Such pre-equalization can be applied to counteract expected loss and dispersion, which “generally requires knowledge of the channel.” EX-1001, 6:40-54. In an example, the above-noted selection of optimal starting point set of coefficients requires communication between a transmitter and a receiver, since testing each set of preset coefficients involves (i) the transmitter transmitting signals to the receiver and (ii) the receiver demodulating the received signals to “estimate or measure error signals.” EX-1001, 14:12-38, FIG. 8. Once the optimal starting point is determined, (i) the transmitter can “program[] the preequalizer with the preset coefficients from the selected register and transmit[] a training

pattern and/or data [to a receiver],” (ii) the receiver can “demodulate[] the receive signals, again obtaining error signals,” and (iii) a transceiver controller can “derive adaptation information” from the error signals and communicate remote adaptation information to the transmitter, e.g., via backchannel information fields. EX-1001, 14:36-67, FIG. 8. Such remote adaptation information can be used to adjust the transmit pre-equalizer. EX-1001, 14:57-15:3, FIG. 8.

During prosecution of the application that led to the ’111 Patent, the examiner issued a non-final Office Action dated July 9, 2020, which rejected original claims 1, 6-10, and 17-19 as obvious under §103 in view of Kolze (EX-1017), Thirumoorthy (EX-1018), and Binder (EX-1019). EX-1002, 000331-000345. The Examiner relied on Kolze for selecting pre-equalizer settings and updating coefficients, and Thirumoorthy for storing pre-equalizer coefficients in registers, *i.e.*, two references that collectively taught storing multiple coefficient sets in registers and using them during equalization training. EX-1002, 000334–000340. Yet in the same Office Action, the Examiner identified at least claims 2, 14, and 20, which recited that “each of the multiple registers corresponds to a different channel model,” as containing allowable subject matter. EX-1002, 000340.

III. PETITIONER’S CHALLENGE TO THE ’111 PATENT

Ground	Claims	Combination
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. Berke (EX-1004)

Berke is directed to a system and method for optimizing equalization coefficients in high-speed serial interfaces such as PCIe, SATA, and SAS channels. EX-1004, ¶¶ 3, 12. Berke identifies the problem that “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.” EX-1004, ¶ 3.

As shown below in FIG. 1, Berke’s high speed serial channel 100 comprises a transmitter 110 and a receiver 130 connected by a transmission channel 120, representing one half of a bidirectional serial data link between components within an information handling system. EX-1004, ¶ 12. The serial channel “provides 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.




FIG. 1

The transmitter includes a channel output module 112 with impedance/tuning settings, a feed-forward equalization (FFE) module 114, a channel management module 116, and an equalization coefficient whitelist 118. EX-1004, ¶¶ 14-15. The receiver includes a CTLE module 134, an AGC module 136, a DFE module 138, a control logic module 140, a channel management module 142, and an equalization coefficient whitelist 144. EX-1004, ¶ 16.

Berke’s system implements a training process to determine a “best set of tuning coefficients” for transmitter 110 and receiver 130 based on acceptable bit error rate (BER) or receiver eye height and width in serial channel 100. EX-1004,

¶¶ 13, 20-21. The training process is orchestrated by control logic module 140, which “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.” EX-1004, ¶¶ 20-21. The control logic module communicates with the transmitter via a backchannel established between channel management modules 116 and 142, which may be “a separate side-band communication channel for communicating adaptation instructions between transmitter 110 and receiver 130.” EX-1004, ¶ 19.

When training variability exceeds a threshold, the control logic module performs a sensitivity analysis to identify “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. Training ranges are narrowed for these sensitive coefficients, and the algorithm is re-run with the narrowed settings. EX-1004, ¶¶ 25-28; FIG. 3.

Berke’s equalization coefficient whitelists 116 and 144 store “certain equalization coefficient settings or combinations of equalization coefficient settings,” which are used “as a starting point for running the training algorithm” by the control logic module 140. EX-1004, ¶ 30. Whitelisted equalization coefficient

settings are populated through two pathways: (1) “knowledge gained during the design, development, or manufacturing of serial channel 100”, where the coefficient settings are known to produce “one or more of a particularly robust eye height or eye width” (EX-1004, ¶ 30), and (2) “knowledge gained during normal operation of serial channel 100,” where the coefficient settings exhibit “a particularly low BER” (EX-1004, ¶ 31). Berke does not disclose determining whitelisted coefficients based on any “channel models.”

Berke describes whitelist storage as “a register or group of registers of transmitter 110 and receiver 130, a memory element of the transmitter and the receiver, another data storage element of the transmitter and receiver, or a combination thereof.” EX-1004, ¶ 30.

When performance degrades, the control logic module “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. Berke never uses the term “channel model” and does not describe indexing its whitelist entries to predefined channel models.

B. Cornelius (EX-1005)

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. The reference addresses the problem of accommodating a legacy standard (DisplayPort) and newer standards (such as “HSIO” or “T29”) on the same physical connector. EX-1005, 3:54-4:9, 5:33-39. Cornelius discloses active cables in Figures 5 and 6 that include “two active plugs...one on each end of cable,” each with “dual clock and data recovery circuits for retiming data” and a cable microcontroller to configure those circuits. *Id.*, 6:6-24, Figs. 5-6.

The core teaching of Cornelius is a connector pinout arrangement and associated circuitry that enables signals from two different standards to share common connector pins without interfering with one another. EX-1005 at 2:20-26, 3:65-5:59, Figs. 3-4. Cornelius describes how pins for a newer high-speed standard may be arranged to reduce crosstalk and interference among themselves, and how additional circuitry—such as resistors, PiN diodes, multiplexers, coupling capacitors, and inductors—may be used to isolate the signal paths of the two standards from each other. *Id.*

C. Ran (EX-1006)

Ran is directed to a system and method for dynamically configuring and adjusting transmit equalization coefficients in high-speed communication links, such as the chip-to-module (C2M) link between a network interface controller (NIC) and a hot-pluggable optical module. EX-1006, 1:15-25, 3:33-40. According

to Ran, “[a]t current rates of 50 Gb/s per lane, equalization is typically implemented in a receiver in the optical module and in a transmitter in the host.” EX-1006, 1:39-41. Ran recognizes the problem that “[f]ixed [transmit] 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. Ran aims to address this problem by having a receiver adjust equalization settings for a transmitter using a request-response procedure performed by a management entity. EX-1006, 3:14-61.

Ran’s system includes a NIC 100 with a host transmitter 208 and host receiver 206, communicatively coupled via an XGAUI-1 electrical connection to an optical module 120 with a module transmitter 210 and module receiver 212. EX-1006, 5:3-30; FIG. 2A. The NIC and optical module each include management registers to enable control of equalization in the receiver in the optical module by the network interface controller: host request register 220 and host response register 222 on the NIC side, and module request register 224 and module response register 226 on the module side. EX-1006, 5:14-18; FIG. 2A. Ran also teaches that four additional registers may be used “to enable control of transmitter equalization in the optical module 120 by host or NIC 100,” enabling bidirectional equalization control. EX-1006, 6:52-62.




FIG. 2A

A management entity mediates equalization adaptation by copying register contents between the host and module sides. EX-1006, 5:36-60, 8:22-9:24, 11:65-13:30; FIG. 6A. The management entity reads the module request register and “writes the value read from module request register to host request register,” conveying the equalization change request from the module to the host. EX-1006, 5:46-50. The management entity then reads the host response register and copies its content to the module response register, completing the transaction. EX-1006, 6:1-18. The management entity “may be in a device driver in an operating system, logic in a network interface controller, or logic in a media access controller.” EX-1006, 5:39-42.

Each request register includes a “two-bit initial condition request field 304,” which is “used to select one of three predefined transmitter equalizer configurations (‘presets’). If one of the presets is not selected, that is, the binary value of the initial condition request field 304 is ‘00’, coefficient select field 312 and coefficient request field 314 are used to update the transmitter equalizer configuration.” EX-1006, 10:6-14, FIGS. 3A-3B.

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

FIG. 3B

A coefficient change request uses the three-bit “coefficient select” field 312 to identify a specific target coefficient, e.g., $c(0)$, $c(-2)$, $c(-1)$, and $c(1)$, and the two-bit “coefficient request” field 314 to specify incrementing, decrementing or holding its value or setting it to “No equalization.” EX-1006, 10:25-48; FIGS. 3D-3E.

Bit 4	Bit 3	Bit 2	Description
1	1	0	c(-2)
1	1	1	c(-1)
0	0	0	c(0)
0	0	1	c(1)

FIG. 3D

Bit 1	Bit 0	Description
1	1	No Equalization
1	0	Decrement
0	1	Increment
0	0	Hold

FIG. 3E

Ran's three presets—Preset1, Preset2, and Preset3—are selectable by setting bits 12 and 13 of the request register. EX-1006, FIG. 3B. Ran does not disclose what these presets contain, how they differ, or what channel characteristics each is designed to address. Ran does not use the term “channel model.”

D. Stauffer (EX-1008)

Stauffer is a textbook that provides “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, at 000011. Stauffer covers SerDes features, functions, design, and protocol logic for standards including PCIe, SATA, and Ethernet. EX-1008, at 000058–000061.

E. Mejia (EX-1009)

Mejia is directed to techniques for using preset evaluation to improve I/O performance in high-speed serial interconnects, specifically PCIe. EX-1009, ¶¶ 1, 10. Mejia recognizes the 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, ¶ 6.

Mejia addresses the problem by proposing a “preset search” technique: “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.” EX-1009, ¶ 10. The preset search may be “implemented as a software-based characterization tool (e.g., implemented as part of BIOS (Basic IO System) reference code)” that “calculate[s]/determine[s] and compare[s] the margin(s) associated with different equalization presets for PCIe, e.g., at 8.0 GT/s or higher.” EX-1009, ¶ 10. The search may also “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.

Mejia’s presets are standardized configurations defined by the PCIe Base Specification. Mejia explains: “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.’” EX-1009, ¶ 12. These presets are protocol-defined configurations communicated between link partners during the PCIe equalization training process: “At the beginning of the equalization training sequence, the link partners communicate starting point for the equalization via preset exchange.” EX-1009, ¶ 12.

Mejia’s preset search process (shown in Figure 3) evaluates different presets by restarting the link with each preset, measuring performance (including jitter margin), and repeating “until all presets are evaluated.” EX-1009, ¶ 31, FIG. 3. “Once done, ... the preset with the highest margins is selected and the link is retained using that preset.” EX-1009, ¶ 31. The selected preset may be “stored (e.g., in BIOS flash or other memory devices discussed herein) for future boots.” EX-1009, ¶ 31.

Mejia observes that “optimal preset that maximizes electrical margins may be a function of channel insertion loss” and identifies three channel-length categories: “(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.

Mejia acknowledges that static preset mapping may be unreliable: “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.

Mejia does not use the term “channel model” and does not describe storing presets in hardware registers indexed to channel models. Mejia does not disclose a controller that selects among registers based on channel model identification.

IV. LEGAL STANDARDS

A patent claim is unpatentable for obviousness if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). In *Graham v. John Deere Co.*, 383 U.S. 1 (1966), the Supreme Court set out a framework for assessing obviousness that requires consideration of four factors: (1) the “level of ordinary skill in the pertinent art,” (2) the “scope and content of the prior art,” (3) the “differences between the prior art and the claims at issue,” and (4) “secondary considerations”

of non-obviousness such as “commercial success, long felt but unsolved needs, failure of others, etc.” 383 U.S. at 17-18.

V. LEVEL OF SKILL IN THE ART

Petitioner defines a person of ordinary skill in the art (“POSITA”) as an individual “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.” Pet., 13-14. Because the ordinary level of skill in the art does not affect the ultimate analysis, neither Patent Owner nor Dr. Williams takes a position in general with respect to Petitioner’s proposed level of ordinary skill in the art. Solely for the purposes of this proceeding, however, Patent Owner and Dr. Williams employ the same level of skill in the art that Petitioner has proposed. *See also* EX-2017, ¶¶ 17-18.

VI. CLAIM CONSTRUCTION

In an *inter partes* review filed on or after November 13, 2018, claim terms are construed based on their ordinary and customary meaning. 37 C.F.R.

§ 42.100(b). Further, the Board should always construe claims to sustain their validity, if possible. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1329 (Fed. Cir. 2005).

Patent Owner requests that the Board adopt the ordinary and customary meaning of the claim terms as understood by one of ordinary skill in the art.

VII. THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE

A patent claim “is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art.” *KSR*, 550 U.S. at 418.

An obviousness determination based on a combination of references also requires finding “both ‘that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.’”

Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd., 821 F.3d 1359, 1367-68 (Fed. Cir. 2016) (citation omitted); *see KSR*, 550 U.S. at 418.

Petitioner has failed to show that the cited prior art renders obvious the Challenged Claims and has not articulated any meaningful motivation to combine or a reasonable expectation of success. Petitioner also fails to account for secondary indicia of non-obviousness, which further demonstrate that the original claims of the '111 Patent were patentable. For these and the reasons stated below, the Petition fails.

**A. Ground 1: Berke in view of Cornelius
(Claims 1, 3, 4, 6-8, 10- 11, 13-17, 19)**

1. Claim 1

a) Berke does not teach multiple registers for pre-equalizer coefficients, each corresponding to a different channel model

Limitation [1a] requires “selecting one of multiple registers to specify initial pre-equalizer coefficient values, each of the multiple registers corresponding to a different channel model.” Petitioner relies solely on Berke to supply the “each of the multiple registers corresponding to a different channel model” feature of limitation [1a]. Pet., 24-26. But in its analysis Petitioner conflates general background awareness of channel conditions with the structured, model-based register organization the claims require.

Berke does not teach that whitelisted coefficients are determined based on knowledge of respective channel characteristics. Instead, Berke’s whitelisted coefficients are selected based on the measured performance of the transmission channel in terms of robust receiver eye height and width or BER. EX-1004, ¶¶ 30-31. While different channel characteristics may impact performance of the transmission channel, Berke does not teach that any specific channel characteristics are known and relied upon to determine whitelisted coefficients. Highlighting this distinction, Berke actually teaches that its training algorithm *may need to be*

repeated multiple times to derive a final converged set of tuning coefficients since “variables in the environment, such as the temperature of serial channel 100, the operating voltages of transmitter 110 and receiver 130, voltage transients, cross talk from other nearby serial channels, random noise, and other effects that are not deterministic and not reproducible.” EX-1004, ¶ 21; EX-2017, ¶ 76. Berke’s whitelist is organized around observed performance – coefficient settings that are “known to produce one or more of a particularly robust eye height or eye width.” EX-1004, ¶¶ 30-31. The organizing principle is output quality, not input characterization. Berke merely remembers what produced good results – it does not define, characterize, or index to discrete models. EX-2017, ¶ 77.

Berke discloses two pathways for populating the equalization coefficient whitelists: (1) knowledge gained during design, development, or manufacturing (EX-1004, ¶ 30); and (2) knowledge gained during normal operation (EX-1004, ¶ 31). Neither pathway involves identifying discrete channel characteristics corresponding to a “channel model” and then deriving corresponding coefficients for it. EX-2017, ¶ 78. In both cases, Berke simply saves coefficients known to produce “particularly robust eye heights or eye widths.” EX-1004, ¶¶ 30-31. Again, Berke merely remembers what produced good results – it does not define, characterize, or index to discrete models. EX-2017, ¶ 78.

Furthermore, Berke's whitelists are a collection of empirically derived "robust" coefficient settings—not a classification of predefined channel representations. Recognizing that a channel *has* characteristics is fundamentally different from defining discrete "channel models" that allegedly represent different configurations of those characteristics. EX-2017, ¶ 79. The organizing principle in Berke is output quality (robust eye performance), not an input model of the channel. Each whitelist entry is simply a coefficient set known to produce robust performance—it "corresponds to" good eye openings, not to any characteristics of the channel. There is no disclosure in Berke that each whitelisted set of coefficients is organized around, indexed by, or defined to match a defined model of the channel or its characteristics as alleged by Petition. EX-2017, ¶ 79.

Berke's teachings regarding selection among whitelisted coefficients reinforces the lack of channel-model-indexed coefficient sets. For example, Berke teaches that after using a first whitelist set for training, "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." EX-1004, ¶ 33. This is a try-one-then-fallback mechanism where the selection criterion is real-time performance feedback, not a determination of which predefined "channel model" applies. In contrast, the claims imply a model-based selection where the

system identifies or assumes a channel model and selects the register that corresponds to it. Berke provides no mechanism for identifying a channel model or mapping an identified model to a particular register. EX-2017, ¶ 80.

It is also telling that Berke never uses the term “channel model,” nor does it discuss any equivalent concept. Petitioner attempts to bridge this gap by asserting that “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.” Pet., 62; EX-1003, ¶ 53. But by this logic, every stored coefficient set in any system would “correspond to a channel model.” All coefficient sets are influenced by some set of real-world conditions. EX-2017, ¶ 81. Such an application would render the “channel model” limitation meaningless.

Furthermore, despite Petitioner’s listing of specific channel characteristics, there is no disclosure of defining discrete “channel models” and indexing equalization coefficient registers to those models. EX-2017, ¶ 82. Indeed, Berke contradicts the notion of a one-to-one mapping between different sets of whitelisted coefficients and discrete channel models representing or assuming different channel conditions. EX-2017, ¶ 83. Berke acknowledges that “multiple different sets of compensation values” can “result in acceptable BER or receiver eye characteristics in serial channel 100.” EX-1004, ¶ 13. As such, different sets of

whitelisted coefficients might work acceptably for the same channel conditions in Berke. The whitelist is thus a pool of known, good starting points for Berke's training algorithm, not a mapping where each entry uniquely covers a distinct channel condition. EX-2017, ¶ 83. Berke explains that "even on a particular information handling system, operating at different times, the back channel adaptation mechanism may operate to provide different sets of compensation values based upon minute variations in the operating condition." (EX-1004, ¶ 13), confirming this lack of correspondence between channel conditions and coefficients.

The prosecution history further distinguishes the claimed invention from Berke's generic practice of storing empirically-validated coefficient sets in a whitelist. As discussed in Section II, the examiner, who had before him references allegedly teaching multiple stored coefficient sets in registers (*i.e.*, Kolze in combination with Thirumoorthy), nonetheless found the "channel model" limitation to recite non-obvious subject matter. Thus, if "channel model" merely meant "any stored coefficient set that was influenced by channel conditions," the limitation presumably would have been met by the combination of Kolze and Thirumoorthy, and the Examiner would not have identified it as allowable. The "channel model" feature as claimed requires more than merely having coefficients informed by channel conditions stored in registers. Accordingly, Berke fails to

teach multiple registers for pre-equalizer coefficients, each corresponding to a different channel model.

b) Berke does not perform all of the claimed steps “in a transceiver”

The preamble [1pre] of Claim 1 expressly states that the method comprises, “in a transceiver,” the steps of [1a], [1b], and [1c]. The use of the singular article “a transceiver” followed by all three limitations indicates that a single transceiver performs all three steps. EX-2017, ¶ 87. Thus, a POSITA reading this claim language would understand that the selecting [1a], updating [1b], and using [1c] all occur within the same transceiver—not distributed across two different devices at opposite ends of a serial link. EX-2017, ¶ 87.

Berke teaches “a high speed serial channel 100 of an information handling system” that “includes a transmitter 110, a transmission channel 120, and a receiver 130.” EX-1004, ¶ 12, FIG. 1. Berke explains that “[s]erial 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.” EX-1004, ¶ 12 (emphasis added). Berke further explains that “[t]he 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.

Under Berke’s bidirectional structure, Petitioner points to different components with respect to the claimed functionality. For example, Petitioner points to Berke’s control logic module 140 as satisfying the “selecting” limitation [1a], Pet., 24-25), and the “updating” limitation [1b], (Pet., 26-27), but the “using” limitation [1c] is mapped to Berke’s FFE module 114 (Pet., 28). FFE module 114 is part of transmitter 110 at the first component. *See* EX-1004, ¶ 14. Whereas, the control logic module 140 resides at the receiver 130 on the second component. *See* EX-1004, ¶¶ 12, 16, FIG. 1.

The first component and the second component are different physical devices—they are the two endpoints of the serial channel, connected by transmission channel 120. EX-1004, ¶ 12. Even though each component may be called a “transceiver,” the control logic module 140 that selects the whitelist set is at the second component’s receiver, while the FFE module 114 that uses the coefficients to transmit data is at the first component’s transmitter. These are not the same transceiver; they are at opposite ends of the link. EX-2017, ¶¶ 88-93.

The claim requires all three steps “in a transceiver”—a single transceiver. Petitioner’s mapping places [1a] in one transceiver (the second component) and [1c] in a different transceiver (the first component), in contravention of the claim’s

structural requirement. Accordingly, Berke does not disclose the “using” [1c] limitation performed by the same transceiver as the “selecting” [1a] and “updating” [1b] limitations.

c) A POSITA would not have been motivated to combine Berke and Cornelius

Claim 1 is directed to a “SerDes communication method.” EX-1001, Claim 1. The Petition suggests that a POSITA would have been motivated “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. Pet., 19. While at best, these two references are in the same field, Petitioner has failed to explain *why* or *how* a POSITA would combine the references.

As an initial matter, Petitioner’s SerDes-context rationale for combining the two references is circular because Berke purportedly already operates in a SerDes context (*e.g.*, PCIe/SATA). This generic rationale is inadequate and does not explain why or how a POSITA would have combined these references to relocate, in a SerDes context, Berke’s transmitter-side whitelist/register/FFE/controller architecture into Cornelius’ cable plug and repurpose Cornelius’s management-oriented communication channels for iterative equalization training as taught by Berke. EX-2017, ¶ 66.

The fact that both Berke and Cornelius relate to SerDes communications establishes field-of-end endeavor overlap, but says nothing about any reasonable expectation of success for the proposed combination in this case. By 2018, SerDes was a dominant paradigm for high-speed digital communication – PCIe, SATA, USB, Ethernet, Fibre Channel, InfiniBand, and dozens of other standards all used SerDes architectures despite significant differences and lack of interoperability among them. EX-2017, ¶ 67. A POSITA would have understood that merely sharing a SerDes context is not sufficient motivation to combine any SerDes-related references without more. *Id.*

The proposed combination effectively requires transplanting into Cornelius's cable plug: (a) Berke's FFE module 114; (b) Berke's channel management module 116; (c) Berke's equalization coefficient whitelist 118; and (d) the entire register-based selection and training architecture described in Berke's paragraphs 19–33. None of these elements exists in Cornelius's cable plug. EX-2017, ¶ 68. The result is not a combination of two references – it is Berke's system re-housed in Cornelius's physical form factor, without any teaching from Cornelius suggesting such a configuration. This is the hallmark of impermissible hindsight reconstruction. *See Metalcraft of Mayville, Inc. v. Toro Co.*, 848 F.3d 1358, 1367 (Fed. Cir. 2017) (“[W]e cannot allow hindsight bias to be the thread that stitches together prior art patches into something that is the claimed invention.”).

Berke's equalization teachings are directed to a fundamentally different cable architecture than that of Cornelius. Berke's training relies on a backchannel between the transmitter's channel management module 116 and the receiver's control logic module 140, through which "adaptation instructions" flow. EX-1004, ¶ 19. Berke teaches that "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. This dedicated adaptation channel must carry real-time performance feedback and coefficient adjustment instructions in an iterative, convergence-driven process.

In contrast, Cornelius discloses communication pathways between host devices and cable plugs—specifically, an LSx sideband bus for "configuration and identification information" and messaging between port microcontrollers and cable microcontrollers for sharing "cable feature information," "diagnostics and fault isolation," and "power management optimization." EX-1005, 6:43–57, 7:7–17. None of these pathways is described as carrying iterative equalization coefficient training instructions.

The LSx bus is a low-speed management channel designed for reading and writing cable configuration and vendor identification data. The microcontroller messaging is for cable feature sharing, diagnostics, fault isolation, and power

optimization. Berke's adaptation channel is expressly for real-time performance feedback and coefficient adjustment – a fundamentally different function requiring different bandwidth, latency, and protocol characteristics. Neither the Petition nor Dr. Chen provide any explanation of how Berke's iterative coefficient training would be established over Cornelius's management pathways. EX-2017, ¶ 71.

In particular, Berke's training involves “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.” EX-1004, ¶ 20. This requires rapid, iterative exchange of coefficient settings and performance metrics. Cornelius's low-speed management channels are not designed for, and are not disclosed as capable of supporting, this kind of high-bandwidth, low-latency adaptation protocol. EX-2017, ¶ 72.

Furthermore, Cornelius describes calibration, but it is fundamentally different from Berke's receiver-driven iterative training. Cornelius's Figure 11 illustrates a loopback-based procedure in which a near end of the cable is placed in loopback mode, a signal is transmitted and received via the loopback path, and “transmit and receive parameters for the near-end circuits may be optimized,” followed by the same procedure for the far end. EX-1005, 11:45–62. This loopback calibration operates locally—the host transmits a signal, it loops back at the near-

end cable plug, and the host receives its own signal back. EX-2017, ¶ 73. There is no disclosed mechanism for the host receiver to send coefficient adjustment instructions to the cable plug's transmitter. The calibration optimizes circuits in isolation as a batch process, not through Berke's iterative convergence based on receiver-side performance monitoring across the full link. EX-2017, ¶ 73.

2. Claims 8 and 11

a) **Berke does not teach multiple registers for pre-equalizer coefficients, each corresponding to a different channel model**

Limitation [8b] requires “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.” Petitioner relies solely on Berke to supply the “each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model” feature of limitation [8b]. *See* Pet., 36; EX-1003, ¶¶ 84. As discussed with respect to limitation [1a], Berke does not disclose multiple registers each containing a set of initial coefficient values corresponding to a different channel model. *See* Section VII.A.1.a).

b) Berke And Cornelius Do Not Disclose “The Port Transceiver Determines a Performance Characteristic”

Limitation [11b] requires “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.” Petitioner relies on Cornelius to supply the “port transceiver” feature and Berke to supply the “performance characteristic” feature of limitation [11b].

Specifically, Claim 11 requires that “the port transceiver determines a performance characteristic” for each register and that “the controller [uses] one of the registers selected by the port transceiver to specify the initial coefficient values for the one or more transmit filters.” This creates a specific functional architecture: the host-side port transceiver evaluates multiple register-stored coefficient sets, determines which performs best, and the module-side controller implements the host’s selection. EX-2017, ¶ 106.

Petitioner maps Berke’s receiver-based control logic module 140 to the functionality of the “port transceiver” and argues it selects whitelisted coefficients in a register by determining their performance. Under this theory, the control logic module 140, configured as part a host-based port transceiver in Ran’s chip-to-

module serial communication link, must: (a) determine performance characteristics for multiple coefficient sets stored in registers in the optical module; (b) select the best-performing register; and (c) instruct the module-side controller to use that register. EX-2017, ¶ 106.

Cornelius does not teach this process. Instead, Cornelius teaches communication pathways between host devices and cable plugs – it does not teach the host evaluating multiple stored coefficient sets in the module for performance and selecting among them. EX-1005, 6:6-24, Figs. 5-6; EX-2017, ¶ 106.

Berke does teach evaluating performance metrics (eye height/width, BER) and selecting among whitelist sets based on performance. EX-1004, ¶33. But Berke performs this evaluation and selection within a single endpoint system where the receiver directly observes the transmitter's signal. Under Petitioner's mapping for claim 11, this process must occur *across* the chip-to-module boundary: the port transceiver (host receiver) must evaluate coefficient sets stored *in the module's controller* and select among them. EX-2017, ¶ 106. Neither Berke nor Cornelius teaches this cross-boundary performance evaluation and selection. EX-2017, ¶ 106.

c) A POSITA would not have been motivated to combine Berke and Cornelius

Claims 8 and 11 are directed to a “chip-to-module communication link.”

EX-1001, Claims 8, 11. Petitioner claims that “a POSITA would have been motivated to apply Berke’s teachings to the serial communication paths in Cornelius’s active cable [in the context of chip-to-module communications link] 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.” Pet., 34. But this generic statement does not address *why* or *how* a POSITA would have combined these references. In particular, it does not address why a POSITA would relocate Berke’s transmitter-side whitelist/register/FFE/controller architecture into Cornelius’s cable plug and split Berke’s tightly-coupled transmitter-receiver training system across a host to cable boundary and repurpose Cornelius’s management-oriented communication channels for iterative equalization training as taught by Berke.

Petitioner’s various attempts to rationalize this combination fall flat. For example, the Petition states that “Cornelius teaches pre-equalization” for its communications link and “Berke’s best-effort link training techniques” would help Cornelius achieve its goal of providing the desired equalization for its communications link. Pet., 34. But Cornelius’s mention that CDR circuits “may

employ equalizer circuits, buffers, emphasis, and de-emphasis circuits as appropriate” (EX-1005, 6:30–42) contains no detail about FFE filters, coefficient values, registers, or training processes. It provides no suggestion to implement Berke’s elaborate coefficient whitelist, register selection, and training architecture.

The Petition also alleges that “[a] POSITA would have had a reasonable expectation of success because the combination simply applies Berke’s serial communications techniques to Cornelius’s serial links.” Pet., 35. But Berke’s training system is designed for endpoint-to-endpoint operation where both the transmitter and receiver are co-located within a single information handling system and the receiver has full observability and control of the channel’s effect on the transmitted signal. EX-1004, ¶¶ 13, 19–20. Splitting Berke’s transmitter components into a cable plug and keeping the receiver in a “port transceiver” at the host introduces a new cable-to-host boundary between the transmitter and receiver, which would require wholly re-engineering both devices without any suggestion to do so. EX-2017, ¶¶ 94-103.

Finally, the lack of motivation with respect to claim 1 applies equally to claims 8 and 11. As explained with respect to claim 1 (Ground 1), there is no motivation to combine Berke’s equalization teachings, which are directed to a fundamentally different cable architecture than that of Cornelius. *See* Section VII.A.1.c).

3. Claim 16

The Petition alleges that the combination of Berke and Cornelius renders obvious claim 16 for the same reasons set forth with respect to claim 8 under Ground 1. Pet., 42. For the same reasons discussed above with respect to claim 8 (*see* Section VII.A.2), Ground 1 fails to teach all of the elements of claim 16.

B. Ground 2: Berke in view of Ran (Claims 1, 3, 4, 6-8, 10-11, 13-17, 19)

1. Claim 1

Petitioner argues that Berke and Ran render obvious Claims 1, 3, 4, 6, and 7 for the same reasons as Berke and Cornelius under Ground 1. Pet., 43. This argument fails for the same reasons explained for Claims 1, 3, 4, 6, and 7 under Ground 1, which are incorporated by reference as if fully set forth herein. *See* Section VII.A.1.

2. Claim 8

a) Berke does not disclose multiple registers for pre-equalizer coefficients, each corresponding to a different channel model

Limitation [8b] requires “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.” Petitioner relies solely on Berke to supply the “each of the multiple registers containing a set of initial

coefficient values corresponding to a different channel model” feature of limitation [8b]. Pet., 47. As discussed in Section VII.A.1.a), Berke does not disclose multiple registers each containing a set of initial coefficient values corresponding to a different channel model.

b) A POSITA would not have been motivated to combine Berke and Ran

Claim 8 is directed to a “chip-to-module communication link.” EX-1001, Claim 8. Petitioner claims that “a POSITA would have been motivated to apply Berke’s teachings to Ran’s serial link 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.” Pet., 45. But this generic statement does not address *why* or *how* a POSITA would have combined these references. In particular, it does not address why a POSITA would graft Berke’s multi-register whitelist of empirically-determined coefficient sets onto Ran’s management entity, which is a relay mechanism with no native coefficient storage, or how Berke’s comprehensive multi-parameter training (which sweeps through coefficient combinations) would be implemented through Ran’s one-coefficient-at-a-time adjustment protocol and across a chip-to-module boundary, when Ran’s architecture provides no mechanism for such cross-boundary multi-register evaluation.

Again, Petitioner’s alleged motivations do not stand up to scrutiny. Petitioner argues that “Berke’s technique of configuring pre-equalization coefficient values” provided an efficient and effective process that addresses a problem Ran recognizes—namely that “fixed equalization as done in existing C2M AUIs may not work well in 100 Gb/s per lane.” Pet., 46. But Ran already addresses this problem through its own mechanism: a management-entity-driven adaptation protocol that enables “continuous modification of the selected equalization coefficients ... after the communications link is established to handle varying operating conditions such as temperature and humidity.” EX-1006, Abstract. A POSITA would not have been motivated to discard Ran’s purpose-built adaptation mechanism in favor of importing Berke’s fundamentally different whitelist-and-training approach, particularly when the two architectures are incompatible at the protocol level (comprehensive sweep vs. one-tap-at-a-time). EX-2017, ¶ 112.

Petitioner further argues that 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. Pet., 47. However, attempting to import Berke’s whitelist architecture into Ran’s management-entity framework would change the fundamental operation of both references. Berke’s system

operates within a single endpoint where the receiver has direct observability and control over the transmitter's coefficients through a dedicated backchannel. EX-1004, ¶¶ 13, 19–20. Ran's system operates through a multi-register relay protocol where the management entity copies requests and responses between host and module registers one coefficient at a time. EX-1006, 3:14–28, 11:65–12:31, FIG. 3E. These are fundamentally different adaptation architectures, and combining them does not produce predictable results. EX-2017, ¶ 113.

Ultimately, Berke and Ran are incompatible. Ran utilizes an on-tap-at-a-time adaptation which is fundamentally different from Berke's comprehensive training process. Berke's training process involves "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." EX-1004, ¶ 20. This comprehensive sweep requires the ability to configure multiple parameters simultaneously and evaluate the joint result. Berke's iterative refinement further narrows to "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. This sensitivity analysis requires configuring specific

coefficient combinations and measuring performance – an operation that requires more than one-tap-at-a-time adjustability. EX-2017, ¶¶ 114-15.

In contrast, Ran’s equalization adaptation protocol allows “changing one equalization coefficient at a time, with a directed change request and arbitrarily fine steps.” EX-1006, 3:14–28. Each request-response cycle involves: (1) the module receiver writes a coefficient select and coefficient request to the module request register; (2) the management entity copies this to the host request register; (3) the host transmitter adjusts the single identified coefficient by one increment or decrement; (4) the host writes a response to the host response register; and (5) the management entity copies this to the module response register. EX-2017, ¶ 116. A POSITA attempting to implement Berke’s whitelist-based selection and training through Ran’s management architecture would face a fundamental architectural mismatch: Berke’s algorithm assumes the ability to configure complete coefficient sets and evaluate their joint performance, while Ran’s protocol permits only incremental single-coefficient nudges. EX-2017, ¶ 117.

3. Claims 11

a) Berke does not disclose “the port transceiver determines a performance characteristic”

Limitation [11b] requires “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.” Petitioner relies on Ran to supply the “port transceiver” feature and Berke to supply the “performance characteristic” feature of limitation [11b]. Pet., 48-49.

Specifically, Claim 11 requires that “the port transceiver determines a performance characteristic” for each register and that “the controller [uses] one of the registers selected by the port transceiver to specify the initial coefficient values for the one or more transmit filters.” EX-1001, Claim 11. This creates a specific functional architecture: the host-side port transceiver evaluates multiple register-stored coefficient sets, determines which performs best, and the module-side controller implements the host’s selection. EX-2017, ¶ 122. The Petition attempts to map Berke’s receiver-based control logic module 140 to the functionality of the “port transceiver,” arguing that it selects whitelisted coefficients in a register by determining their performance. Pet., 49. Under this theory, the control logic module 140, configured as part a host-based port transceiver in Ran’s chip-to-module serial communication link, must: (a) determine performance characteristics for multiple coefficient sets stored in registers in the optical module; (b) select the best-performing register; and (c) instruct the module-side controller to use that register.

But Ran does not teach this process. Ran teaches a management entity that copies requests between registers – it does not teach the host evaluating multiple stored coefficient sets in the module for performance and selecting among them. EX-2017, ¶ 124. Ran’s initial condition request field “is used to select one of three predefined transmitter equalizer configurations (‘presets’)” (EX-1006, 10:6–14), but Ran does not disclose any process of determining a performance characteristic for each preset and selecting based on that determination. The selection mechanism for the initial condition in Ran is entirely undisclosed as it is simply a two-bit field that is set to some value through an unspecified process. EX-2017, ¶ 124.

Berke does teach evaluating performance metrics (eye height/width, BER) and selecting among whitelist sets based on performance. EX-1004, ¶ 33. But Berke performs this evaluation and selection within a single endpoint system where the receiver directly observes the transmitter’s signal. Under Petitioner’s mapping for claim 11, this process must occur *across* the chip-to-module boundary: the port transceiver (host receiver) must evaluate coefficient sets stored *in the module’s controller* and select among them. EX-2017, ¶ 125. Neither Berke nor Ran teaches this cross-boundary performance evaluation and selection.

b) A POSITA would not have been motivated to combine Berke and Ran

Like Claim 8, Claim 11 is also directed to a “chip-to-module communication link.” EX-1001, Claim 11. For the same reasons discussed in Section VII.B.2.b) with respect to Claim 8, a POSITA would not have sought to combine Berke and Ran to arrive at claim 11.

4. Claim 16

The Petition alleges that the combination of Berke and Ran renders obvious claim 16 for the same reasons set forth with respect to claim 8. Pet., 49. For the same reasons discussed in Section VII.B.2 with respect to Claim 8, claim 16 is also non-obvious.

C. Ground 3: Berke in view of Cornelius and Stauffer (Claims 2, 5, 9, 12, 18) and

Ground 4: Berk in view of Ran and Stauffer (Claims 2, 5, 9, 12, 18)

Grounds 3 and 4 challenge Claims 2, 5, 9, 12, and 18 based on the combination of Berke, Cornelius, and Stauffer or Berke, Ran, and Stauffer, respectively. Petitioner adds Stauffer to Ground 1 (Berke in combination with Cornelius) and Ground 2 (Berke in combination with Ran) to allegedly supply teachings for insertion loss/package loss and error signal energy.

1. Claims 2 and 9

Dependent Claims 2 and 9 require that “at least some of the different channel models are for chip-to-module (C2M) channels with different insertion loss and package loss characteristics.” EX-1001, claims 2, 9. Petitioner argues that Berke teaches channel models corresponding to different “insertion loss and other signal degradations” of transmission channels and asserts that “it would have been obvious for Berke’s ‘other signal degradations [of transmission channels] to include package loss.” Pet., 51. This argument fails for two reasons.

First, as explained with respect to claim 1 (Ground 1), Berke does not teach registers corresponding to different “channel models.” Indeed, Petitioner and Dr. Chen’s reasoning imposes two separate requirements on Berke’s whitelisted coefficients: “they must [i] yield high performance (e.g., robust eye heights or eye widths, low BER) and [ii] fit a set of channel characteristics of a channel model so as to compensate for channel loss/distortion.” EX-1003, ¶ 140; *see* Pet., 51. Berke, however, does not require coefficients to “fit a set of channel characteristics of a channel model.” Berke requires only that coefficients are known to produce “robust eye height or eye width.” EX-1004, ¶ 30; EX-2017, ¶ 130.

Second, Petitioner’s reliance on Stauffer falls short because Claim 2 does not merely require awareness that package loss exists, it requires that “channel models” account for “different insertion loss and package loss characteristics.” *See*

EX-1001, Claim 2; EX-2017, ¶ 131. Petitioner points to Stauffer’s discussion of “package models,” but these are transfer function representations of the chip package used in an offline channel analysis – the process of evaluating whether a channel design will support adequate BER performance using statistical simulation tools, channel transfer function, and transmitter/receiver transfer functions. EX-1008, 000031, 000033-34. Stauffer’s package models are not part of any channel model used for selecting equalization coefficients during system operation. EX-2017, ¶ 131.

2. Claims 5, 12, and 18

Dependent Claims 5, 12, and 18 each require that “the performance characteristic is an error signal energy.” Petitioner alleges that it would have been obvious to modify Berke based on Stauffer to use error signal energy as a metric of channel performance on top of Berke’s existing eye-based metrics for initial pre-equalizer coefficients. Again, this argument falls short for multiple reasons.

First, Petitioner’s reliance of Stauffer’s discussion of “root-mean-square value” is misplaced. Pet., 53-54. A careful reading of Stauffer reveals that the “root-mean-square” discussion occurs in the context of *random jitter* specification. EX-1008, 000031-33. Random jitter is distinguishable from other types of jitter, such as deterministic jitter (DJ). EX-1008, 000033 (explaining that RJ “does *not*

correlate with the data pattern being sent” whereas “several components of DJ are dependent on the data pattern being sent”). Specifically, Stauffer states:

The rationale for specifying different types of jitter is that some of the jitter can be corrected. For instance, the FFE in the transmitter device alters the transitions of the data edges based on the pattern of 1’s and 0’s being sent. In effect, the FFE is used to inject DJ on the transmitter output. . . . *On the other hand, RJ is not predictable, and therefore cannot be corrected by equalization.* It is important to limit the amount of RJ on the serial link.”

EX-1008, 000033 (emphasis added). This is confirmed by Dr. Chen’s own description of “equalization.” See EX-1003, ¶ 11 (“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’).”). Since the type of jitter that Stauffer specifies using RMS (random jitter) “cannot be corrected by equalization,” a POSITA reading Stauffer would understand that RMS jitter metrics are not useful for evaluating or selecting among equalization coefficient sets. EX-2017, ¶¶ 134-36. As such, Stauffer’s use of “root-mean-square” is not as a performance characteristic for evaluating or selecting among pre-equalizer coefficient sets. *Id.*

Second, the argument that “root-mean-square value” is equivalent to teaching error signal energy conflates two distinct concepts: RMS of jitter and MSE of Error Signals. EX-2017, ¶ 137. RMS of jitter is a statistical measure of the timing variation in bit transitions, used to specify random jitter in interface standards. *See* EX-1008, Page 000032. This measures *timing* variation, not *amplitude* error between received and expected signals. EX-2017, ¶ 137. In contrast, MSE of error signals is a measure of the energy content of the error signal, defined as the difference between the received signal and the expected signal. EX-1027, 000002; EX-1022, 000033. This measures amplitude discrepancy between what was received and what was expected. EX-2017, ¶ 137. These are fundamentally different quantities applied to different signal characteristics (timing jitter vs. amplitude error). The mathematical relationship between RMS and MSE as abstract statistical concepts does not mean that teaching RMS jitter specification is equivalent to teaching MSE-based error signal energy evaluation for equalization coefficient selection. EX-2017, ¶¶ 137-38.

Third, the suggestion that “no additional data collection or measurement would have been required to determine error signal energy” given Berke’s use eye diagram and BER (Pet., 54) oversimplifies the distinction between discrete bit-error counting and continuous error-signal energy computation. While Berke does compare received patterns to expected patterns during BER calculation (EX-1004,

¶¶ 24-26), calculating BER (counting bit errors) is a categorically different operation from computing error signal energy (calculating the mean square of the continuous-time amplitude difference between received and expected waveforms). EX-2017, ¶ 139. BER counts discrete errors, whereas MSE quantifies continuous-amplitude deviations. EX-1003, ¶ 149. The data used to compute BER (bit-level comparisons) is not the same data used to compute MSE (sample-level amplitude differences), and the computational infrastructure for the two is different. EX-2017, ¶ 139.

Fourth, Berke already discloses using eye height/width and BER as its performance metrics for evaluating and selecting among whitelist coefficient sets. EX-1004, ¶¶ 20, 24, 33. Even if one accepted that use of error signal energy as measured by MSE was known and could be applied on top of Berke's performance metrics, the mere availability of MSE metrics does not suffice to explain why a POSITA would have adopted an additional approach on top of Berke's existing solution. EX-2017, ¶ 140.

**D. Ground 5: Ran in view of Mejia and Stauffer
(Claims 1-19)**

1. Claim 1

**a) Neither Ran Nor Mejia Teaches Using Channel
Models**

Ran does not teach the use of channel models, nor does Petitioner allege as much. Instead, Petitioner asserts that Mejia teaches presets that “correspond to different channel models.” Pet., 62. However, the “preset search” that Petitioner identifies in Mejia actually teaches the opposite.

The “preset search” requires testing each preset to identify which one works best at that particular moment. Mejia explains the process as follows: “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.” EX-1009, ¶ 10. Mejia further explains that the preset search process (shown in Figure 3) evaluates different presets by restarting the link with each preset, measuring performance (including jitter margin), and repeating “until all presets are evaluated.” EX-1009, ¶ 31, FIG. 3. “Once done, ... the preset with the highest margins is selected and the link is retained using that preset.” EX-1009, ¶ 31. This process requires testing all of the presets to see which one works best because it knows nothing about the channel. EX-2017, ¶ 152. If the system stored different presets for different channels, then there would be no need to test every

single preset each time the system initializes, restarting the link with each preset to see how it performs compared to the other presets. EX-1009, ¶ 12; EX-2017, ¶ 152.

Mejia acknowledges that the “optimal preset... may be a function of channel insertion loss” and other characteristics of the channel. EX-1009, ¶ 36. However, that does not mean that each of Mejia’s presets is “corresponding to a different channel model,” as claim 1 requires. EX-2017, ¶ 153. If each of Mejia’s presets corresponded to a particular channel model, then there would be no need to test every preset to see how it performs. One could simply select the preset that corresponds to the current channel model, which would be much more efficient. Mejia’s selection process, which requires restarting the link to test each preset, teaches the opposite. EX-2017, ¶ 153.

b) A POSITA Would Not Have Been Motivated to Combine Ran, Mejia, and Stauffer

The Petition asserts that a POSITA would have included Mejia in combination with Ran because Mejia is “analogous art,” is “in the same field,” and “relates to ensuring reliable data transmission through a communication channel.” Pet., 56. Petitioner describes in detail the different selection procedures disclosed in Ran and Mejia, but Petitioner fails to explain *why* a POSITA would have used

Mejia's teachings when Ran already discloses a process for selecting equalizer presets.

In particular, the Petition does not address why a POSITA would graft Mejia's coefficient selection process onto Ran's management entity, or how Mejia's preset selection process would be implemented given that Ran already discloses an intricate request and response process for selecting coefficients. Petitioner asserts that "Ran teaches selecting one preset from multiple presets but does not limit how the selection is performed." Pet., 64. It is true that Ran does not expressly *limit* how the coefficient selection is performed. However, as Petitioner acknowledges, Ran teaches particular and detailed request and response procedures for performing the selection. Pet., 58-61. Petitioner also acknowledges that Mejia teaches its own particular and detailed coefficient selection procedures, which are different from Ran's procedures. Pet., 61-63. Petitioner identifies no teaching in Ran or Mejia that would invite or motivate a POSITA to replace or modify Ran's selection procedures to incorporate the different selection procedures described in Mejia. EX-2017, ¶ 144. Petitioner's argument is nothing more than impermissible hindsight reconstruction.

Petitioner further argues that a POSITA would have 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).” Pet., 64. But this argument falls short for two reasons.

First, Petitioner fails to sufficiently explain *how* a POSITA would “apply Mejia’s preset search technique to implement Ran’s selection of a preset.” Pet., 64. As Petitioner acknowledges, Ran and Mejia each teach their own detailed procedures for selecting coefficients, and Petitioner does not explain how a POSITA would have combined those different procedures. For example, Petitioner provides no explanation of how Mejia’s “preset search” procedures would fit into Ran’s five-step request-response selection procedures. Pet., 64-67; EX-1006, 3:14-61.

Petitioner generally asserts that “Ran and Mejia’s teachings were readily compatible; they are in the same field of serial communications... and both support serial communications for PCIe.” Pet., 65. Petitioner also asserts that both Ran and Mejia “select a preset to provide initial pre-equalizer coefficient values for equalization training.” Pet., 65. Petitioner then concludes that a POSITA “would have a reasonable expectation of success because *it* would not have changed the principle of operation of either reference.” Pet., 65 (emphasis added). Missing in Petitioner’s analysis is any explanation of what “it” is. Petitioner fails to explain

which part(s) of Mejia's selection procedures would be used or what part(s) of Ran's five-step selection procedures would require modification to accommodate the Mejia procedures. EX-2017, ¶¶ 147-48. The fact that two references both teach coefficient selection procedures in the same field of endeavor does not mean that a POSITA could simply combine those procedures without making necessary modifications. EX-2017, ¶¶ 147-48. Petitioner provides no explanation of the necessary modifications or why a POSITA would have been motivated to make them.

Second, Petitioner fails to sufficiently explain *why* a POSITA would have used Mejia's coefficient selection process to modify Ran's coefficient selection process. Ran already taught its own iterative, five-step selection process, which Ran proposed to solve a particular problem. The problem Ran solves is that at high speeds "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. Ran aims to address this problem by having a receiver adjust equalization settings for a transmitter using a detailed request-response procedure performed by a management entity. EX-1006, 3:14-61. Petitioner identifies no teaching in Ran that would have invited a POSITA to look elsewhere for alternative selection procedures. Given that Ran already includes a detailed

coefficient selection process, there would be no reason for a POSITA to do so. EX-2017, ¶ 149.

2. Claim 8

Largely mirroring its analysis from claim element [1a], Petitioner again relies on Ran and Mejia as the teaching for limitation [8b], which requires “a controller having multiple registers, each of the multiple registers containing a set of initial coefficient values corresponding to a different channel model.” Pet., 76-78. For the same reasons discussed in claim 1, neither Ran nor Mejia teaches storing or using different sets of coefficient values that correspond to different channel models. *See* Section VII.D.1.a).

Similarly, a POSITA would not have been motivated to combine Ran, Mejia, and Stauffer for claim 8 for the same reasons discussed above with respect to claim 1. *See* Section VII.D.1.b).

3. Claim 11

a) Neither Ran Nor Mejia Teaches a Port Transceiver Determining a Performance Characteristic

Petitioner relies on the general combination of Ran and Mejia as the teaching for [11b], which requires “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.” Pet., 79. As discussed above

(*see* Section VII.B.3.a)), Claim 11 and [11b] create a specific functional architecture: the host-side port transceiver evaluates multiple register-stored coefficient sets, determines which performs best, and the module-side controller implements the host's selection. EX-2017, ¶ 158.

Ran does not teach the right functionality. Ran teaches a management entity, separate from the pluggable module transceiver, that copies requests between registers. EX-2017, ¶ 160. Ran does not teach a transceiver evaluating multiple stored coefficient sets in the module for performance and selecting among them. Further, Ran's initial condition request field "is used to select one of three predefined transmitter equalizer configurations ('presets')" (EX-1006, 10:6–14), but Ran does not disclose any process of determining a performance characteristic for each preset and selecting based on that determination. The selection mechanism for the initial condition in Ran is entirely undisclosed as it is simply a two-bit field that is set to some value through an unspecified process. EX-2017, ¶ 160.

Petitioner also asserts that Mejia teaches the determination and selection under its analysis of [11b], but Mejia does not fill the gap in Ran's teaching. In particular, Petitioner cites paragraphs 31 and 32 of Mejia, including the sentence fragment: "recommend the best equalization preset values(s) based on information from the RX (Receive) circuits." Pet., 79; EX-1009, ¶ 31. This snippet from Mejia does not teach the structure or functionality of claim 11 or [11b] in particular.

First, Petitioner fails to identify which entity in Mejia corresponds to the claimed port transceiver. Second, by selectively quoting only a fragment of a sentence, Petitioner likewise fails to identify which entity “recommend(s) the best equalization preset value(s).” Pet., 79. Looking at the broader context of paragraph 31, Mejia itself fails to identify which entity provides the recommendation. EX-1009, ¶ 31.

To the extent Petitioner meant to assert that the “RX (Receive) circuits” constitute the claimed port transceiver, that assertion would not match the requirements of claim 11. Whatever entity provides the recommendation in Mejia, it clearly is *not* the RX (Receive) circuits because Mejia explains that the recommendation may be “based on information *from* the RX (Receive) circuits.” EX-1009, ¶ 31 (emphasis added). Therefore, Mejia does not teach a host-side port transceiver configured to perform the required selection of coefficients. EX-2017, ¶ 162.

b) A POSITA Would Not Have Been Motivated to Combine Ran, Mejia, and Stauffer

A POSITA would not have been motivated to combine Ran, Mejia, and Stauffer for claim 11 for the same reasons discussed in claim 1. *See* Section VII.D.1.b).

4. Claim 16

Petitioner argues that the combination of Ran, Mejia, and Stauffer renders obvious claim 16 for the same reasons set forth with respect to claim 8 under Ground 5. Pet., 81. For the same reasons as claim 8, Ground 5 does not teach claim 16. *See* Section VII.D.2.

E. Grounds 1-5: Secondary Indicia of Non-obviousness Further Undermine Petitioner’s Challenges to the ’111 Patent

Secondary indicia of non-obviousness further show that it would not have been obvious at the time of the invention to selectively combine particular teachings from the cited references, as Petitioner urges. “Objective indicia of nonobviousness must be considered in every case where present.” *Apple Inc. v. Samsung Elecs. Co., Ltd.*, 839 F.3d 1034, 1048 (Fed. Cir. 2016). Here, Patent Owner Credo Technology Group has licensed its AEC technology to multiple manufacturers of active cables. *See* EX-2019, EX-2020, EX-2021, EX-2022, EX-2023, EX-2024. Licensing activity provides “strong probative value” when evaluating secondary considerations of non-obviousness. *See Institut Pasteur & Universite Pierre Et Marie Curie v. Focarino*, 738 F.3d 1337, 1347 (Fed. Cir. 2013). Accordingly, Petitioner’s challenges to the claims of the ’111 Patent should be rejected.

VIII. CONCLUSION

For the reasons stated above, Patent Owner respectfully asks the Board to find that the challenged claims of the '111 Patent are not unpatentable.

Date: May 12, 2026

Respectfully submitted,

/Charles M. McMahon/

Charles M. McMahon
USPTO Reg. No. 44,926
Mircea A. Tipescu
USPTO Reg. No. 53,690
David H. Bluestone
USPTO Reg. No. 44,542
Thomas DaMario
USPTO Reg. No. 77,142
BENESCH FRIEDLANDER
COPLAN & ARONOFF LLP
71 South Wacker Drive, Suite 1600
Chicago, IL 60606
Tel: (312) 312-4949
Email: CredoIPR@beneschlaw.com

**CERTIFICATE OF COMPLIANCE WITH
TYPE-VOLUME LIMITATIONS**

I, the undersigned, do hereby certify that the foregoing Patent Owner Response, including footnotes, contains 11,580 words, as measured by the Word Count function of Microsoft Word as specified by 37 C.F.R. § 42.24(b)(2).

/Charles M. McMahon/

Charles M. McMahon

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. § 42.6(e)(4), I hereby certify that on May 12, 2026, a true and correct copy of the foregoing was served by electronic mail upon the following counsel of record for Petitioner:

Harper Batts (HBatts@sheppardmullin.com)
Chris Ponder (CPonder@sheppardmullin.com)
Jeffrey Liang (JLiang@sheppardmullin.com)
LegalTM-Marvell-Credo-IPRs@sheppardmullin.com

/Charles M. McMahon/

Charles M. McMahon