

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


Reply Declaration of Michael S. Chen

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

1. I previously submitted a declaration in this proceeding (Ex-1003). I have been asked to respond to certain opinions of Dr. Tim A. Williams. I have considered the additional materials listed in the attached Appendix 1. I make this declaration based upon my own personal knowledge and, if called upon to testify, would testify competently to the matters contained herein.

2. Dr. Williams argues that Berke does not perform all of the claimed steps of claim 1 “in a transceiver.” Ex-2017, ¶¶87-93. As part of his argument, Dr. Williams asserts that I “map[] the ‘selecting’ limitation [1a] to Berke’s control logic module 140” (Ex-2017, ¶89) and “rel[y] on Berke’s control logic module 140, which resides at receiver 130 on the second component, to supply the ‘updating’ limitation [1b]” (*id.*, ¶90). This mischaracterizes my testimony.

3. As explained in my original declaration, Berke teaches that “transmitter 110 and receiver 130 communicate with each other to optimize and adjust various compensation values [e.g., pre-equalizer coefficients] within the transmitter and the receiver to compensate for the insertion loss and other signal degradations of transmission channel 120.” Ex-1003, ¶47; Ex-1004, ¶13, Fig. 1:



100

FIG. 1

The communication includes transmitting “adaptation instructions” through a “back channel” or “side-band communication channel” between the channel management module 116 on the transmitter 110 and the channel management module 142 on the receiver 130. Ex-1004, ¶13, ¶19. I explained that both the selecting and updating steps, as performed by Berke, involve both the transmitter and receiver.

4. For the “selecting” step, I explained that “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.” Ex-1003, ¶52. Berke’s whitelisted pre-equalizer coefficients are stored in the whitelist registers 118 on the transmitter and used to configure the pre-equalizer on the transmitter (e.g., FFE module 114). Ex-1004, ¶14, ¶30, ¶32. The transmitter uses pre-equalizer coefficients from a whitelist register to configure the serial channel. *Id.* The receiver, via the control logic module 140, monitors the performance of the transmitted signal and communicates adaptation instructions to the transmitter. Ex-1004, ¶19, ¶33. Based on the adaptation instructions, the transmitter may select another whitelist register and use the set of pre-equalizer coefficients stored therein to configure the serial channel. Ex-1004, ¶33. Therefore, the operations for “selecting one of multiple registers” are carried out by both the transmitter and the receiver.

5. For the “updating” step, I explained that “[Berke’s] training process adjusts compensation values within the transmitter, including the pre-equalizer coefficients, via communications between the transmitter and receiver.” Ex-1003, ¶58. The transmitter runs variations of pre-equalizer coefficients by using them to program the pre-equalizer and the receiver records performance characteristics of the transmitted signals. Ex-1004, ¶20. Based on the performance characteristics, the control logic module 140 on the receiver uses adaptation instructions to direct the transmitter to update the pre-equalizer coefficients for further training. *See, e.g., id.* (“Control logic module 140 continues the training by directing transmitter

110 to select a next tuning setting of channel output module 112 and repeating the process of recording the associated eye height or BER for each of the various combinations of settings.”). Here, the pre-equalizer coefficients reside on the transmitter and are updated by the transmitter based on adaptation instructions from the receiver.

6. In light of Berke’s teachings regarding back channel adaptation, a POSITA would have understood Berke’s selection and training of pre-equalizer coefficients are performed by both the transmitter and the receiver. This is because, had these adaptation steps been performed solely on the receiver side, no back channel communication between the transmitter and the receiver would have been needed.

7. Dr. Williams argues that Berke’s “dedicated adaptation channel must carry real-time performance feedback” requiring “high-bandwidth, low-latency.” Ex-2017, ¶¶69-72. I disagree with Dr. Williams. I observe that Dr. Williams’s declaration does not cite any disclosure of Berke that supports a requirement for the adaptation channel to be “real-time,” “high-bandwidth,” or “low-latency.” Ex-2017, ¶¶69-72. I reviewed Berke and found no disclosure that imposes a specific performance requirement on the back channel. *See* Ex-1004, ¶13, ¶19. This is consistent with a POSITA’s understanding of back channels for equalization adaptation. Back channel communication for equalization adaptation generally

does not require high bandwidth because it carries only a modest volume of data.

The adaptation instructions typically convey identification information for equalization settings, such as a register address, or directives to modify equalization coefficients. The resulting data payload is orders of magnitude smaller than the data transmitted over the serial channel. In addition, back channel communication for equalization adaptation does not require low latency.

Adaptation instructions need not be transmitted continuously at every data symbol; rather, they are sent periodically after the receiver has evaluated the signal performance of a preceding set of training data during training. Because each set of training data includes a substantial amount of data symbols transmitted through the serial channel, a POSITA would therefore have expected the data rate required for adaptation instructions to be far lower than the symbol rate of the serial channel. Similarly, the latency requirement of the back channel is not high during training because a significant amount of training time is consumed by transmitting and evaluating the training data, which is orders of magnitude longer than data symbol time.

8. Dr. Williams further argues that Cornelius's "low-speed management channels" cannot support Berke's "real-time performance feedback," which requires "different bandwidth, latency, and protocol characteristics." Ex-2017, ¶71. I first observe that Dr. Williams does not explain what these different

characteristics are and how they are not supported by Cornelius. Cornelius teaches communication between cable microcontrollers and port microcontrollers. Ex-1005, 7:7-17. Dr. Williams argues this communication is conducted through a “LSx sideband bus,” which is a “low-speed management channel.” Ex-2017, ¶71. I disagree with this characterization. Cornelius describes the LSx bus as one “specific embodiment” and expressly notes that “other signal pins may be used.” Ex-1005, 6:63-67. Cornelius discloses such other signal pins to include ones that Cornelius characterizes to be high-speed. Ex-1005, 6:16-20. Moreover, regardless of whether the communication channel in Cornelius operates at high or low speed, a POSITA would have understood its capacity was sufficient to carry adaptation instructions for equalization adaptation. As I explained above, such instructions involve only a modest volume of data, with a required data rate far below the symbol rate of the serial channel, and are only sent periodically.

9. Indeed, Cornelius teaches that its channel is used for communicating configuration information and “operational parameters” allowing “cables to adapt to new hosts and devices as the cable is used in various system applications.” Ex-1005, 6:50-57. A POSITA would have understood that a prominent aspect of adapting to a new host is adjusting equalization settings to suit the channel between the cable and the host. Accordingly, a POSITA would have understood that Cornelius’s operational parameters encompass equalization settings and related

parameters. Berke's adaptation instructions convey information for selecting and updating pre-equalizer coefficients—operational parameters that enable a transceiver to adapt to channel conditions. *See, e.g.*, Ex-1004, ¶13, ¶19. This is precisely the type of information Cornelius contemplates transmitting through its channel. Furthermore, a POSITA would have understood that adaptation instructions for equalization adaptation do not impose any heightened requirements beyond those of typical operational parameters and configuration information. Because Cornelius's channel is designed to carry operational parameters and configuration information, a POSITA would have understood that the same channel is capable of transmitting adaptation instructions, such as that taught by Berke.

10. Dr. Williams argues that “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 that “[s]plitting 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.” Ex-2017, ¶97. I disagree. Although Berke describes a transmitter and a receiver within an information handling system, nothing in Berke requires the transmitter and receiver to be “co-located.” To the contrary, Berke contemplates a variety of architectures for the information handling system:

Here, the components can be understood to include elements within an information handling system, such as components that are attached to one or more printed circuit board of the information handling system, where transmission channel 120 can represent one or more circuit traces on the printed circuit board, and can include one or more connectors. 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.

Ex-1004, ¶12. A POSITA reading this disclosure would have understood that Berke contemplates architectures in which the transmitter and receiver communicate across the boundary between a host (e.g., “one or more printed circuit board[s]”) and a cable (e.g., “one or more transmission cables”) via “one or more connectors.” *Id.*

11. Dr. Williams argues that “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” and “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-2017, ¶113. Based on this characterization, Dr. Williams concluded that “[t]hese are fundamentally different adaptation architectures, and combining

them does not produce predictable results.” *Id.* I disagree. Berke expressly teaches that its back channel communication “can include a third party device or mechanism that provides system management ... or another embedded controller.” Ex-1004, ¶19. Ran’s management entity is precisely such a mechanism. It facilitates communication between the transmitter and the receiver by, for example, relaying a request from an optical module’s receiver to a host’s transmitter to adjust operational parameters such as equalizer settings. Ex-1006, 6:37-44, 11:65-12:31. Berke therefore expressly contemplates an architecture that uses a mechanism like Ran’s management entity to mediate back channel communication between the transmitter and the receiver.

12. 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 18_____, 2026.

By:

shuowei chen

Michael S. Chen

**Appendix 1: Additional Materials Considered
in the Preparation of this Declaration**

Exhibit No.	Description
Ex. 2017	Declaration of Tim A. Williams