

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WESTERN DIGITAL TECHNOLOGIES, INC.,

Petitioner

v.

SHAMROCK INNOVATIONS, LLC,

Patent Owner

Case No.: IPR2023-00842

U.S. Patent No. 9,535,454

Title: COMPUTING MODULE WITH SERIAL DATA CONNECTIVITY

**DECLARATION OF TERRY R. LEE
IN SUPPORT OF PATENT OWNER'S RESPONSE**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. QUALIFICATIONS	1
III. COMPENSATION	5
IV. MATERIALS CONSIDERED.....	5
V. SUMMARY OF MY OPINIONS	6
VI. LEGAL PRINCIPLES.....	7
VII. LEVEL OF ORDINARY SKILL IN THE ART	11
VIII. CLAIM CONSTRUCTION	12
IX. SUMMARY OF THE ‘454 PATENT	12
A. ‘454 Patent	12
B. ‘214 Patent	16
C. ‘752 Patent	20
X. SUMMARY OF CITED REFERENCES	21
A. Vangemert	21
B. Copeland	24
C. Dobson	27
D. Reay	29
XI. THE PETITION FAILS TO DEMONSTRATE THAT THE CHALLENGED CLAIMS ARE UNPATENTABLE.....	30
A. The Petition Fails to Prove That Independent Claim 1 is Obvious Over Vangemert and Copeland	30
1. There is No Motivation to Combine Vangemert and Copeland	31
2. The proposed combination of Vangemert and Copeland, would render Vangemert inoperable for its intended purpose of providing a low latency algorithm for buffer credit calculation according to the Fibre Channel protocol.	39
3. Petition Fails to Demonstrate That Vangemert or Copeland Transmits the Flow Control Signal Through the Serial Link and Connector Without Passing the Flow Control Signal Through the Second Buffer.....	45

A.	The Petitioner Fails to Demonstrate That Vangemert Transmits the Flow Control Signal Without Passing the Flow Control Signal Through the Second Buffer.	45
B.	Petitioner Fails to Demonstrate That Copeland Transmits the Flow Control Signal Through the Serial Link as recited in the Challenged Claims	49
4.	Petitioner’s Stated Motivation to Combine Vangemert and Copeland and Minimize Latencies is Incorrect	53
B.	Petition Fails to Prove That Independent Claim 1 is Obvious Over Dobson and Copeland	56
1.	Petition Fails to Demonstrate That Dobson or Copeland teaches a module configured to transmit the flow control signal through the serial link and the connector without passing the flow control signal through the second buffer,” as recited in the challenged claims.....	57
2.	Petitioner’s proposed mapping of the recited “second buffer” is inconsistent and, therefore, does not show how the proposed combination of Dobson and Copeland teaches every element recited in the challenged claims.	60
XII.	CONCLUSIONS.....	63

I. INTRODUCTION

1. My name is Terry Lee. I have been retained as an expert witness by Patent Owner Shamrock Innovations, LLC (“Shamrock” or “Patent Owner”) for this *Inter Partes* Review IPR2023-00842 of U.S. Patent No. 9,535,454 (the “’454 Patent” or the “Patent”) filed by Western Digital Technologies, Inc. (“Western Digital” or “Petitioner”).

2. In the Petition, I understand that Petitioner is challenging the validity of Claims 1-4, 6, 7, and 15-17 (the “Challenged Claims”) of the ’454 Patent (Ex. 1001), on the grounds of obviousness. I have been asked to prepare the present Declaration to be submitted with Patent Owner’s Response.

3. I have been asked to consider whether the Challenged Claims of the ’454 Patent would have been obvious to a person of ordinary skill in the art (“POSITA”) as of the date of the invention. I was also asked to review and comment on several technical statements made by Petitioner in the Petition and by its expert, Dr. Michael Caloyannides.

4. In this Declaration, all emphasis in boldface has been added unless otherwise noted.

II. QUALIFICATIONS

5. This section summarizes my educational background, career history, publications, and other relevant qualifications. My curriculum vitae is attached as

Patent Owner's Exhibit 2001, which includes my detailed employment background, professional experience, and list of publications.

6. I have over 40 years of industry experience in integrated circuits, PCB design, and memory subsystem design. This experience includes both product development and product definition with a long-term focus on next generation DRAM technologies.

7. I received a Bachelor of Science Degree in Electrical Engineering from the University of Missouri - Columbia in 1983, a Master of Science Degree in Electrical Engineering from the University of Southern California - Los Angeles in 1985, and a Master's Degree in Business Administration from the University of Washington - Seattle in 1010.

8. Since 2016, I have been the President and Subject Matter Expert at IdaTech Consulting, which provides electrical engineering expert services to intellectual property firms. These services include research, reverse engineering, claim charting, patent assessment, and patent litigation support. My previous industry experience includes 26 years of employment at Micron Technology, Inc., a Fortune 500 memory semiconductor company.

9. From 1980 to 1983, I was employed by Texas Instruments, Inc, as a co-op student for CMOS integrated circuit design.

10. From 1983 to 1989, I was employed by Hughes Aircraft Co. as an electrical engineer and engineering manager. My responsibilities included RF digital integrated circuit design, and subsystem design for aerospace and electronic warfare systems. This work included the development of an advanced Digital RF Memory (DRFM) subsystem and custom bipolar integrated circuits for serializing and deserializing data words.

11. From 1989 to 2015, I was employed at Micron Technology, Inc. in various engineering roles of progressive responsibility. These roles involved DRAM Integrated Circuit Design, Applications Engineering, DRAM Product Definition, Strategic Marketing, Signal Integrity, Logic Design, PCB Design for memory modules (DIMMs) and Solid State Drives (SSDs), advanced semiconductor packaging, engineering tester development, and Managed NAND development.

12. While working for Micron, I was involved in many standards-setting groups and industry consortiums focused on defining the next-generation DRAM products. My first involvement with the JEDEC Solid State Technology Association (formerly known as Joint Electron Device Engineering Council, or just JEDEC) was in 1992. In 1996, I became involved in the SLD RAM Consortium (formerly known as SyncLink) for developing the specification of a next generation high-speed DRAM technology. SLD RAM pioneered many

technologies, which were later adopted into DDR SDRAM, DDR2 SDRAM, and DDR3 SDRAM. Around 1999, I was involved in the Advanced DRAM Technology Consortium (ADT) for defining a next-generation DRAM technology. The consortium included several DRAM manufacturers and Intel Corporation. ADT pioneered several technologies that became adopted in DDR2 DRAM and DDR3 DRAM. At different times, I was the Technical Leader for both SDRAM and ADT. My area of focus and expertise in these organizations involved high-speed signaling and bus communication, physical design, and memory subsystem design.

13. At various times while working for Micron, I was a member of IEEE and part of the Memory Paper Selection Subcommittee for IEEE International Solid State Circuits Conference (ISSCC). I gave various presentations and keynote addresses in support of memory industry conferences while at Micron. I also served time on Micron committees for patent disclosure approval and for the Technical Career Ladder.

14. In June 2015, I retired from Micron and started my own company in 2016 to provide expert services in support of intellectual property firms.

15. I am the named inventor of 97 patent families, all of which are directed in general to memory integrated circuits, memory module design, high-speed interfaces, and packaging.

16. I have served as an expert witness for the Federal Trade Commission, European Commission and for patent litigation.

III. COMPENSATION

17. I am being compensated for my work in this case at the rate of \$400 per hour. I am also reimbursed for all reasonable expenses that I incur during the course of this case. My compensation does not depend upon the results of my analysis or the substance of my testimony, nor does my compensation depend on the outcome of this or any related proceeding. I have no personal interest in the outcome of this matter. I have no financial interest in Patent Owner or affiliation with any of the real parties in interest, the Patent Owner, or the named inventor of the '454 Patent. It is conceivable that I may own mutual funds whose portfolios include stock in Petitioner. If this is the case, the value of such holding would not constitute a material part of my net worth.

18. The statements made and opinions provided in this Declaration are based on my own personal knowledge and, if called as a witness, I could and would testify in a manner consistent with this Declaration.

IV. MATERIALS CONSIDERED

19. I have reviewed the following documents in forming the opinions expressed in this Declaration and have considered all the materials cited herein and all materials filed in this IPR up to and including January 7, 2024.

- a. U.S. Patent No. 9,535,454 to Ahern et al. (“‘454 patent”)
- b. Prosecution History of the '454 patent
- c. Declaration of Michael Caloyannides, Ph.D.
- d. U.S. Patent No. 6,289,386 to Vangemert ("Vangemert")
- e. U.S. Patent No. 5,179,661 to Copeland, III et al. ("Copeland")
- f. U.S. Patent No. 5,717,870 to Dobson (“Dobson”)
- g. U.S. Patent No. 6,005,414 to Reay (“Reay”)
- h. U.S. Patent No. 6,088,752 ("‘752 patent")
- i. U.S. Patent No. 6,070,214 ("‘214 patent")
- j. U.S. Patent No. 8,060,675 ("‘675 patent")
- k. Excerpt from Prosecution History of the '675 patent

20. I have also relied on my education, skill, training, and experience in the relevant fields of technology in forming my opinions. I have further considered the viewpoint of a person of ordinary skill in the art (“POSITA”) as of the time of the inventions of the ‘454 Patent.

21. I reserve the right to supplement my opinions as expressed in this Declaration to address any new information obtained in the course of this proceeding, or based on any new positions taken by Petitioner.

V. SUMMARY OF MY OPINIONS

22. It is my understanding that the following table summarizes the grounds of challenge to the Challenged Claims raised in the Petition:

Ground	Invalidity Challenge	References	Challenged Claims
1	§ 103	Vangemert, Copeland	1, 2, 3, 4, 7, 15, 16, 17
2	§ 103	Vangemert, Copeland, Dobson	4
3	§ 103	Vangemert, Copeland, Reay	6
4	§ 103	Dobson, Copeland	1, 2, 3, 4, 7, 16
5	§ 103	Dobson, Copeland, Reay	6
6	§ 103	Dobson, Copeland, Vangemert	15, 17

23. After a review of the alleged prior art asserted by Petitioner and the Dr. Michael Caloyannides (Ex. 1003), it is my opinion that Petitioner has not demonstrated that the claims are unpatentable.

24. My opinions, and the bases, therefore, are detailed throughout this Declaration.

VI. LEGAL PRINCIPLES

25. Counsel for Patent Owner has informed me of the legal principles that apply in this proceeding.

26. I have been informed that Petitioner has the burden to prove that the claims challenged in the Petition are not patentable by a preponderance of the evidence, which I understand to be just enough evidence to make it more likely than not that Petitioner's argument is correct.

27. A claim is invalid as obvious under 35 U.S.C. § 103 if the differences between the subject matter sought to be patented and the prior art are such that the subject matter of the claim as a whole would have been obvious at the time of the invention to a person having ordinary skill in the art to which the invention pertains. I have been informed that the following factors are used to determine whether or not the claimed subject matter would have been obvious: (i) the scope and content of the prior art; (ii) the differences, if any, between the prior art and the claimed invention; (iii) the level of ordinary skill in the field of the invention; and (iv) any relevant objective considerations of non-obviousness.

28. A party asserting obviousness based on a combination of prior art references must demonstrate that one of ordinary skill in the art would have been motivated to combine the teachings of those references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so. It is my understanding that it is not enough to show that one skilled in the art could combine elements of multiple references, but instead there must be some reason that would have prompted a person of ordinary skill in the art to combine the elements in the way the claimed invention does. I understand that there must be some reasoned explanation as to why one of ordinary skill in the art would combine the references.

29. I understand that obviousness may be shown by considering more than one item of prior art and by considering the knowledge of a person of ordinary skill in the art, and that obviousness may be based on various rationales including: (i) combining prior art elements according to known methods to yield predictable results; (ii) simple substitution of one known element for another to obtain predictable results; (iii) use of known techniques to improve similar devices in the same way; (iv) applying a known technique to a known device ready for improvement to yield predictable results; (v) “obvious to try” – choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; (vi) known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; and (vii) some teaching, suggestion, or motivation in the prior that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

30. A combination of references would not have been obvious if the alleged modification(s) to be made to the reference(s) would have been inconsistent with the reference’s stated goals or method of operation or would have rendered the combination inoperable for its intended purpose. I further understand that for something to have been obvious, the party asserting

obviousness must explain why a POSITA would have selected components for combination in the manner claimed.

31. I further understand that a claim is not obvious over a combination of prior art references if such references “teach away” from the claimed combination, if there is no motivation to combine such references, or if the combination would lead to waste and inefficiencies not present in one or more of the references in isolation.

32. It is my further understanding that obviousness cannot be based on a hindsight combination of components selected from prior art references. For example, the challenged Patent itself cannot be used as a basis for combining prior art references absent such a teaching in the Patent. I also understand that an invention would not have been obvious simply because all of the elements of the invention may have been known separately in the prior art.

33. The concept of incorporation by reference was also explained to me. Incorporation by reference is a key mechanism in U.S. patent law that allows a patent applicant to incorporate material from an external document into a patent application by referring to it. This method is beneficial for including material that is too extensive to be included directly in the application or has already been described in a prior patent or application. The aim is to simplify the application process and eliminate the need for repeating information. The material

incorporated by reference is considered an integral part of the patent application, which helps in fully describing the invention and supporting the patent's claims.

VII. LEVEL OF ORDINARY SKILL IN THE ART

34. I understand that the hypothetical Person of Ordinary Skill in the Art (POSITA) is considered to have the normal skills and knowledge of a person in a certain technical field, as of the time of the invention at issue. I understand that factors that may be considered in determining the level of ordinary skill in the art include: (1) the types of problems encountered in the art; (2) the prior art solutions to those problems; (3) rapidity with which innovations are made; (4) the sophistication of the technology; and (5) the education level of active workers in the field. I also understand that “the person of ordinary skill” is a hypothetical person who is presumed to be aware of the universe of available prior art. Therefore, no such person actually exists.

35. Petitioner has proposed that a POSITA for the ‘454 Patent “would have had at least a bachelor’s degree in electrical engineering, or equivalent, with about three years of professional work experience in serial and parallel data communication within or between computing devices or systems.” (Petition, 11) Petitioner further alleged that “More education can supplement practical experience and vice versa.” (Id).

36. I agree with Petitioner that these would be the minimum qualifications for a POSITA.

37. The opinions I express herein are from the viewpoint of such a POSITA as of August 6, 1998.

VIII. CLAIM CONSTRUCTION

38. After having read the ‘454 Patent in its entirety, I find no reason that any of the claim terms should be construed to have a meaning that is different than the plain and ordinary meaning as understood by a POSITA at the time of invention. Accordingly, when evaluating the terms recited in the claims of the ‘454 Patent, I have applied their plain and ordinary meaning.

IX. SUMMARY OF THE ‘454 PATENT

A. ‘454 Patent

39. U.S. Patent 9,535,454, issued on January 3, 2017, presents a computing module designed for asynchronous, serial exchange of parallel system bus data with other modules in a computer system. The invention allows for the transfer of parallel bus data between modules of a computer system in a serial manner, both sending and receiving data asynchronously. This approach enables more flexible and efficient data communication within computer systems, particularly in scenarios where multiple modules need to interact and exchange data.

40. The background of the invention addresses the limitations of contemporary computer systems, particularly the restrictions imposed by the prevalent PCI bus technology. The PCI bus, widely used in the industry, is limited in its ability to be extended over distances or connected to remote devices due to physical and loading constraints. This limitation hampers the division of computer systems into modular components and their adaptability to evolving computing environments. The invention seeks to overcome these limitations by enabling more flexible connections and communication between various computer components and peripherals, such as storage devices, displays, and input devices like keyboards and mice.


41. The technical advantages of '454 Patent are achieved through a modular computer system that includes a universal connectivity station (UCS). This station is designed to connect and route data via serial data links to various devices using proprietary Split-Bridge™ technology. The system comprises a computer core and a UCS, which are spatially separate but functionally integrated. The UCS enables the connection of the processor module to a range of devices and other computer modules, all while maintaining transparency in terms of data bus architecture. This modular approach allows for easy upgrades and modifications to the system's performance module, such as enhancements in processor speed and other key components, without necessitating changes to the

rest of the computer system. This innovation offers a novel approach to computer architecture and upgradeability, addressing the needs of diverse computing environments and security-sensitive situations.

42. The '454 Patent incorporates U.S. Patent Application Ser. No. 09/130,057, filed on August 6, 1998, which granted as the '214 Patent. As explained below, the disclosures from the '214 Patent include an example of the flow control technique claimed in the '454 Patent. Claim 1 captures aspects of the disclosed flow control example. In that example, in response to the fill amount of a first buffer being at or above a threshold, the module that includes the first buffer generates a flow control signal which is forwarded through the serial link without passing through a second buffer.

43. The '454 Patent incorporates U.S. Patent Application Ser. No. 09/130,058, filed on August 6, 1998, which granted as the '752 Patent.

44. In the following Family tree of the '454 Patent (Ex. 2004), the '214 Patent, '752 Patent, '675, and '454 Patent, are highlighted in Blue, Red, Orange, and Green respectively.




B. '214 Patent

45. U.S. Patent 6,070,214, issued on May 30, 2000, describes a technology for a serially linked bus bridge that expands access from a first bus to a second bus. This bridge is designed to be accessible by a host processor and can connect various bus-compatible devices, including memory and input/output devices. The key components of the bridge include a link and two interfaces: the first interface couples the first bus to the link, and the second interface connects the second bus to the link. These interfaces function as a single bridge and are capable of sending information serially through the link in a format distinct from that of the first and second buses.

46. The bridge operates without waiting for incoming acknowledgment over the link before transferring information, enhancing efficiency and speed. The first interface and the second interface work together to approve initial exchanges between the first and second buses based on pending transactions. Additionally, the bridge allows the host processor, communicating through the first bus, to individually address different selectable devices on the second bus. This is achieved using the same type of addressing as on the first bus, without the need for an intervening device on the second bus.

47. Specifically, the '214 Patent discloses an exemplary implementation of a serial data link in which frames produced by a first encoder are forwarded

through a first transmitter via a first simplex link to a first receiver, which supplies the serial information to a first decoder. Likewise, a second encoder forwards serial information through a second transmitter along a second simplex link to a second receiver, which supplies the serial information to a second decoder. Ex. 1010, 6:39-46. The ‘214 Patent depicts this exemplary implementation as shown in the following annotated version of Figure 1.



48. The ‘214 Patent further explains that “[f]rames produced by encoder 28 are forwarded through transmitter 44 along simplex link 46 to receiver 48, which supplies the serial information to decoder 30. Likewise, encoder 26 forwards serial information through transmitter 38 along simplex link 40 to

receiver 42, which supplies the serial information to decoder 32.” Ex. 1010, 6:39-46.

49. The ‘214 Patent identifies an overflow danger relating to registers (or buffers) 22 and 24, and discloses a flow control technique for addressing that danger. Id., 6:47-48. For example, the ‘214 Patent describes that each register generates a “threshold detect” signal when it is “almost full.” Id., 6:49-50. The “threshold detect” signal is then forwarded via a simplex link to the decoder on the side of the link. That decoder then issues a “threshold stop” signal to an encoder, which then stops forwarding information to the register that originally generated the “threshold detect” signal. Id., 6:49-58. The ‘214 Patent depicts the signal flow of the “threshold detect” and “threshold stop” signals as shown below in the following annotated version of Figure 1. The diagram contains an obvious error, in that the arrow for signal 36 should point in the opposite direction.

51. Accordingly, the ‘214 Patent outlines various embodiments related to the bridge’s functionality, focusing on the hierarchical exchange of information between buses, processing of transactions, and compatibility with different types of bus systems. The disclosure provides an advanced method for expanding access and communication between different bus systems, which is essential in computer and electronic systems where multiple devices need to be connected efficiently.

C. ‘752 Patent

52. U.S. Patent 6,088,752, issued on July 11, 2000, describes a docking system that enables a portable computer to access a second bus via a first bus. This system comprises a link, along with a first and a second interface. The first interface connects the first bus to the link, and the second interface connects the second bus to the link. Both interfaces operate to send bus-related information through the link in a format different from that of the first and second buses, allowing the portable computer to address bus-compatible devices on the second bus using the same addressing as on the first bus.

53. The patent details a method for docking a portable computer with a first bus to communicate with a second bus over a link, each bus being capable of connecting to various bus-compatible devices. The method involves sending bus-related information through the link in a different format from the buses and

enabling the portable computer to address devices on the second bus using the first bus's addressing system. This disclosure provides an improved system for transferring information between buses, utilizing duplex links formed from simplex links, and managing data with FIFO (first-in, first-out) registers for efficient transmission.

54. This technology is particularly relevant for scenarios involving a portable computer and a docking station. In a preferred embodiment, the system can be implemented using two application-specific integrated circuits (ASICs) connected by a cable. These ASICs can operate in different modes depending on the associated bus, with one ASIC including an arbiter for controlling the secondary bus. This setup is ideal for docking stations, providing a bridge between different bus types like CardBus and PCI, and supporting various peripheral devices including video cards and processing circuits for monitors.

X. SUMMARY OF CITED REFERENCES

A. Vangemert

55. U.S. Patent 6,289,386, issued on September 11, 2001, details an innovative approach to buffer credit calculation in high-speed serial communication networks, particularly for systems using the Fibre Channel protocol. The Patent introduces a network interface unit equipped with a low latency algorithm to manage buffer credits more efficiently. This unit comprises a

receive buffer, a receive controller, a transmit controller, and a buffer credit manager. The buffer credit manager is responsible for calculating committed buffer space based on data stored in the receive buffer and signals from both the receive and transmit controllers. It uses a comparator to assess whether the available buffer space exceeds the committed, unoccupied space plus the maximum frame size, ensuring efficient and timely flow control in the network.

56. Specifically, Vangemert describes that “[t]he buffer credit manager iteratively determines the additional committed buffer space (i.e., the space committed but not yet occupied by receive data) by counting the number of receive and transmit flags, and generates the credit signal using a comparator which compares the sum of additional committed buffer space and the maximum frame size to the available receive buffer storage capacity and asserts the credit signal if the available buffer space is greater than or equal to the sum of the maximum frame size and the additional committed buffer space.” (Ex. 1005, 2:19-29) Vangemert’s buffer credit manager is graphically depicted in Figure 5 (reproduced below).




FIG. 5

57. With regard to Figure 5, Vangemert expressly describes that:

Since the data storage required for completely received frames is accounted for in the available buffer space, the buffer credit manager doesn't need to continue reserving space for it. Since the actual amount of stored data is being used in this iterative buffer-credit management strategy, more efficient use of the buffer space is possible. When the size of the received data frames is less than the maximum frame size, the extra space can add up to allow additional buffer credits to be sent. Using this strategy, the number of credits made available by the buffer credit manager is at least equal to the maximum buffer capacity divided by the frame size, but it may be greater when received data frames are smaller than the maximum frame size. Importantly, the credits are generated quickly since only additions, subtractions, and comparisons are used. This incremental “one-more-credit?”

evaluation is well suited to the serial communication channel since only one credit can be sent at a time anyway. The hardware implementation of this buffer credit manager is advantageously fast, simple and inexpensive. (Ex. 1005, 4:52-5:4)

58. Accordingly, Vangemert describes a system where the buffer credit manager issues credits in proportion to the available space in the receiver's buffer. If there is no space in the buffer of the receiver, the transmitter cannot transmit. As a result, the buffer credit manager makes buffer overflows impossible.

B. Copeland

59. U.S. Patent 5,179,661, issued on January 12, 1993, describes a method and apparatus for automatically controlling the flow of serial data from a remote transmitter to a host device. The system includes a serial communications card with a UART (Universal Asynchronous Receiver/Transmitter) that has a buffer. A counter tracks the number of bytes received by the UART since the host last read all data in the buffer. If this count exceeds a predefined portion of the buffer capacity, it triggers a control signal to the remote transmitter to stop sending data, thus controlling data flow without the host's intervention.

60. The Patent outlines a mechanism within an asynchronous serial communications adapter card for host computers. It features a 16550A UART and a microprocessor, which can enable or disable the automatic flow control feature.

The system includes circuitry to perform this automatic flow control, along with gates allowing the microprocessor to manage the flow control circuitry.

61. The UART uses a BAUDOUT signal, a clock at 16 times the baud rate, and an inverted -RXREADY signal, which activates when there's at least one byte of serial data in the UART's 16-byte FIFO buffer. The flow control is executed through these signals and a counter. When the -RXREADY signal goes low, indicating the FIFO buffer isn't empty, the counter is enabled and starts counting the BAUDOUT cycles. Once a predetermined number of cycles is counted, it signals the remote transmitter to halt data transmission.

62. Specifically, Copeland describes that:

After a predetermined number of BAUDOUT cycles are counted without the FIFO buffer being empty an output of the counter is enabled which disables the counter. This output deasserts the DTR and/or RTS signals to the remote transmitter. This signals the remote transmitter to stop sending. The enabling of the counter output thus results in the remote transmitter halting its transmission of data. (Ex. 1006, 2:56-63).

63. Copeland's process that automatically manages data flow, preventing overflow and data loss, is graphically depicted in Figure 3 (reproduced below).




Fig - 3

64. Specifically, Copeland teaches a system where “data flow control is activated when the FIFO buffer is forty percent full.” (Ex. 1006, 5:39-40)

Copeland further explains that the “ data flow control drives either the DTR (Data Terminal Ready) or the RTS (Request To Send) line of the serial card to a “not ready” status to halt the flow of serial data from the remote transmitter.” (Ex. 1006, 2:12-16) By way of example, Copeland explains that “The count of 1024 BAUDOUT cycles represents 64 bit times or, at 10 bits per character, 6.4 character times” (Ex. 1006, 5:33-35) and “[t]his leaves at least 9.6 character times remaining for the remote transmitter connected to serial data port 20 to respond before the FIFO buffer overflow.” (Ex. 1006, 5:35-38) Accordingly, Copeland teaches a system where transmission of data is halted when the receiver’s buffer

reaches a limit so that the amount of inflight data does not cause a buffer overflow.

C. Dobson

65. U.S. Patent 5,717,870, issued on February 10, 1998, presents an advanced controller for serial ports in computers, particularly enhancing functionality with universal asynchronous receiver/transmitters (UARTs). The invention includes an expanded buffer, which connects to either the transmitter or receiver buffer of the UART. This setup significantly increases the data storage capacity for transmitted or received data through the serial port.


66. The patent describes a scalable trigger circuit associated with the expanded buffer. This circuit detects when the data in the expanded buffer reaches a predetermined level and produces a service request signal. Furthermore, the patent introduces a timer that responds to the service request signal to provide an interrupt signal to the computer. After an interrupt is serviced, the timer resets the interrupt signal and starts a predefined timeout period. During this period, the system is non-responsive to subsequent service requests, preventing repetitive interrupts. This mechanism reduces the load on the computer's processing resources by managing the interrupt signals more effectively.

67. Specifically, Dobson describes that:

Basically stated, the present invention, which is embodied as a part of this device, is a set of expanded of transmit and

receive FIFOs, together with scalable trigger level circuits to control the generation of service requests for data in the buffers. In particular, the present invention provides expanded buffers with scalable trigger levels for generation of service requests that are likewise expanded to approximately mimic the conventional selectable trigger levels of the type 16550 UART. This allows the present invention to be used with applications and devices that select one of the conventional trigger levels (i.e., 1, 4, 8, or 14 bytes in a 16 byte buffer) to physically select a larger number of bytes as the trigger level, in a way that will be essentially transparent to the application running on the computer. Therefore, the present invention can be used with much higher data rate modems in concurrent processing environments with respect to a conventional type 16550 UART, yet reliably handle high speed incoming and outbound data and servicing interrupts generated by the received data buffer without the loss of data. (Ex. 1007, 3:33-52)

68. Dobson's system of managing interrupt signals more efficiently is graphically depicted in Figure 6 (reproduced below).



69. Accordingly, Dobson teaches a system where interrupt signals are calculated and transmitted to a computer.

D. Reay

70. U.S. Patent 6005414A, titled “Mixed-mode multi-protocol serial interface driver,” describes a driver that operates in both current-mode and voltage-mode. It is designed to conform output signals to various electrical interface standards, such as CCITT/EIA standards V.35, V.11/RS-422, V.28/RS-232, and V.10/RS-423. The Patent features a mode-select input signal that enables the selection of a particular standard, activating the appropriate circuitry for that standard. When certain parts of the circuitry are not in use, they are placed in a high impedance state to prevent interference with the active sections.

XI. THE PETITION FAILS TO DEMONSTRATE THAT THE CHALLENGED CLAIMS ARE UNPATENTABLE

A. The Petition Fails to Prove That Independent Claim 1 is Obvious Over Vangemert and Copeland

71. The Petition fails to prove that Independent Claim 1 is Obvious over Vangemert and Copeland for at least the following reasons.

72. First, a combination of Vangemert and Copeland to improve data flow control is neither obvious nor logical given that Vangemert teaches a flow control scheme based on buffer credits, which makes overflow impossible, while Copeland teaches using a flow control signal to stop transmission to avoid a possible overflow, (i.e., in a system where overflow is possible). Accordingly, there would be no purpose in using a flow control signal to stop potential overflow when overflow is not possible, and therefore, there is no motivation to combine Vangemert's and Copeland's flow control solutions.

73. Second, a combination of Vangemert and Copeland would render Vangemert inoperable for its intended purpose of providing a low “a low latency algorithm for buffer credit calculation” for use with the Fibre Channel protocol. (Ex. 1005, 1:65-2:1),

74. Third, the Petition fails to demonstrate that the module is “configured to transmit the flow control signal through the serial link and connector without passing the flow control signal through the second buffer,” as

recited in the Challenged Claims. Rather, Vangemert teaches that “A transmit controller 412 retrieves control words and data frames from transmit buffer 410 and forwards them to encoder 402” (Ex. 1005, 3:49-52), and “the transmit controller 412 begins transmitting buffer credit control words in response to a credit signal from buffer credit manager 418.” (Ex. 1005, 4:14-16)

75. Fourth, the Petition also fails to demonstrate that Copeland transmits a flow control signal through the serial link; instead, Copeland’s flow control signal is an out-of-band parallel signal transmitted through a separate signal line and not through the serial link.

76. Fifth, the motivation to combine Vangemert and Copeland and minimize latencies is incorrect. Copeland explains, “The present invention, therefore, automatically signals the remote transmitter to stop sending when the FIFO buffer fills to a certain point and to resume sending once the FIFO buffer has been completely emptied,” (Ex. 1006, 3:3-6) which has a detrimental effect on serial link latency and performance to Vangemert’s solution, so the motivation to combine Vangemert with Copeland and improve latencies is illogical. In contrast, the flow control solution of ‘454 mitigates the latency impact of flow control and achieves higher performance.

1. There is No Motivation to Combine Vangemert and Copeland

77. A combination of Vangemert and Copeland to improve data flow control is neither obvious nor logical given that Vangemert teaches a flow control scheme based on buffer credits, which makes overflow impossible, while Copeland teaches using a flow control signal to stop transmission to avoid a possible overflow.

78. Buffer credits are a vital flow control mechanism in high-speed networks like Fibre Channel, designed to manage data transfer efficiently and prevent loss. Each device in the network has limited buffer space for incoming data, and buffer credits track this space's availability. This tracking is crucial for maintaining an orderly data transfer process and ensuring that the receiver's buffer isn't overrun by incoming data.

79. In practical terms, buffer credits operate as a communication tool between the sender and receiver. The receiver sends a specific number of buffer credits to the sender, indicating the amount of data it can handle at a time. Each credit typically represents a unit of data, such as a frame. As the receiver processes the data and frees up space, it sends additional credits back to the sender. This dynamic exchange allows the sender to adjust its data transmission rate in real-time, aligning with the receiver's capacity to process the incoming data.

80. The importance of buffer credits becomes evident in maintaining network efficiency and data integrity. By enabling the sender to match its transmission rate with the receiver's processing ability, buffer credits prevent potential data overflow and loss. This system is especially crucial in high-performance networks, where balancing fast data transfer rates with the need for reliable and orderly data processing is essential. Buffer credits thus play a key role in ensuring the smooth operation of these advanced communication systems.

81. For example, the Asynchronous Transfer Mode (ATM) Forum proposed a "specification for a credit-based scheme called Quantum Flow Control... which lets a node send data only when its downstream partner has enough credits to receive the data." (See, Ex. 2008). As noted by Network World, one of the limitations of credit-based system is that "you could be told to hold back completely" when there are no credits available to transmit the data. (Ex. 2008)

82. In another example system that uses buffer credits entitled "Scalable Protocol Engine For High-Bandwidth Communications," the buffer control credits are explained as:

The flow control mechanism is also maintained in this module. Fibre Channel uses a back pressure mechanism based on buffer credits to regulate the traffic flow. The buffer credit of the receiver is maintained at its upstream node. A buffer credit

counter of an upstream node is incremented each time a packet is sent to the downstream node. This process is suspended when the maximum buffer credit is reached. At that point, the upstream node has to wait for the arrival of the receiver ready signal before the credit counter can be decremented and a new packet be sent. (Ex. 2007)

83. Accordingly, a POSITA would understand that when two devices are communicating, the sender must know how much buffer space is available on the receiver's end to avoid overrunning its capacity. In a credit-controlled system, the receiver communicates this information by sending back a certain number of buffer credits to the sender. As a result, when the sender receives buffer credits, it knows how many units of data it can send without overwhelming the receiver's buffer.

84. In the present case, Vangemert teaches a revised buffer credit-based system for flow control to avoid buffer overflow. The general concept of a credit-based system discussed above is consistent with Vangemert's description of the Description of Related Art. Specifically, the inventor provides an example of such a system from the Fibre Channel protocol: "Under this methodology, before a first network node (initiating node) can transmit a data frame to a second network node (responding node), the initiating node must first receive a buffer credit from the responding node, indicating that the responding node has available space in its

receive buffer for a data frame.” (Ex. 1005, 1:49-54). Accordingly, since the initiating node is required to receive the buffer credit in advance of sending the data, the possibility of buffer overflow is eliminated in such a credit-controlled system.

85. Vangemert attempts to improve the buffer credit-based system by making more efficient use of the received buffer: “Further, predetermination of the number of buffer credits is not possible due to the programmability of the maximum frame size, and assuming the worst case scenario (as is necessary for predetermination) results in extremely inefficient use of the receive buffer. A low-latency implementation of buffer-credit determination is needed.” (Ex. 1005, 1:58-63)

86. Vangemert avoids receiver buffer overflow with a buffer credit-based system by getting advance permission from the receiving node that sufficient buffer space is available before a transmitter can transmit frames: “The remote node is configured to transmit a data frame to the local node only if the remote node receives a buffer credit which indicates that the local node has available receive buffer space for a data frame” (Ex. 1005, 2:3-6) and “[t]he buffer credit manager iteratively determines the additional committed buffer space (i.e. the space committed but not yet occupied by receive data) by counting the number of receive and transmit flags, and generates the credit signal using a comparator

which compares the sum of additional committed buffer space and the maximum frame size to the available receive buffer storage capacity and asserts the credit signal if the available buffer space is greater than or equal to the sum of the maximum frame size and the additional committed buffer space.” (Ex. 1005, 2:19-29) Vangemert teaches the relationship of control signals to control words used for flow control at “the buffer credit manager 418 is enabled, and the transmit controller 412 begins transmitting buffer credit control words in response to a credit signal from buffer credit manager 418” (Ex. 1005, 4:13-16) and “the fabric switch 114 transports data frames and control words from the transmit channel of a source node to the receive channel of a destination node.” (Ex. 1005, 2:66-3:1) From these statements, a POSITA would understand that buffer overflow will not occur because a transmitter will not be enabled to transmit unless there is adequate receiver buffer space to receive a data frame.

87. Further, a POSITA would also understand that the buffer credit manager calculates available space, sends a buffer credit signal if sufficient space is available, and buffer control words are then sent to the transmit node via buffer control words, and that these buffer control words or “credits” enable the transmitter to send data frames.” Without buffer credits, the transmitter is not allowed to send data frames and overflow the receiving buffer.

88. On the other hand, Copeland teaches a flow control method that attempts to stop a transmitter from sending data as the receive buffer starts to fill. Copeland relies on a round-trip latency from receiver-to-transmitter to stop the transmitter from sending after the receive buffer is partially full, or forty percent full in the preferred embodiment. Specifically, Copeland states “UART 14 has another signal, BAUDOUT, which is a clock at 16 times the baud rate. When - RXREADY goes low counter 151 counts the BAUDOUT cycles. When 1024 BAUDOUT cycles pass without the FIFO buffer being empty the Q11 output of the counter 151 goes high. The count of 1024 BAUDOUT cycles represents 64 bit times or, at 10 bits per character, 6.4 character times. This leaves at least 9.6 character times remaining for the remote transmitter connected to serial data port 20 to respond before the FIFO buffer overflow. Thus, in the preferred embodiment of the present invention, data flow control is activated when the FIFO buffer is forty percent full” (Ex. 1006, 5:28-40), and “which sends a low for the RTS and DTR signals on serial port connector 20. This signals the remote transmitter to stop sending data, thus shutting down the flow of data from the serial data port 20 to the UART 14 FIFO buffer.” (Ex. 1006, 5:51-55) Therefore, a POSITA would understand that as the buffer becomes relatively full, a signal is sent to the transmitter to stop transmitting data frames so receive buffer overflow could be avoided, and that a certain amount of latency, 9.6 character times in the preferred

embodiment, is provided for the transmitter to respond to avoid the receive buffer overflow. As a result, a POSITA would that in Copeland's system, overflow is indeed possible because the transmitter is not required to obtain advance permission from the receiver prior to transmitting data.

89. It is illogical, if not nonsensical, to combine Vangermert with Copeland because Vangemert doesn't allow buffer overflow to occur, so adding the flow control mechanism of Copeland to respond to pending overflow would have no purpose. Simply put, no purpose would be served by supplementing Vangemert's system where an overflow is impossible, with Copeland's flow control techniques designed to minimize overflow risk in systems where overflow is possible. Therefore, there would be no motivation to combine Vangemert's and Copeland's flow control solutions.

90. An analogous combination argument to Vangemert and Copeland was made in the '454 prosecution history by the Patent Examiner. The Patent Examiner cited U.S. Patent 5,613,141, "Szatkowski" in view of Daines et al. (U.S. Patent 6,192,422, "Daines." (Ex. 1002, 182-186). This argument was overcome by the Applicant during final prosecution, "Szatkowski teaches a "flow control" mechanism that prevents buffer overflow. Szatkowski's use of an acknowledgement-based protocol may eliminate most or all occurrences of buffer overflow or "near" overflow", and "Daines teaches: If (the amount of data in an

input buffer crosses a first threshold level), the flow control mechanism inserts congestion data into an output buffer for the port to alert the associated node to cease transmitting data packets to the port” (Ex. 1002, at 226-227), and “Further, the acknowledgement-based protocol of Szatkowski may eliminate most or all occurrences of buffer overflow, or “near” overflow, rendering unmodified Szatkowski superior to Daines in terms of flow control. Thus, the proposed modification in view of Daines may actually be detrimental to Szatkowski, e.g., by eliminating the advantageous acknowledgement-based protocol of unmodified Szatkowski.” (Ex. 1002, at 227-228). Subsequent to this final argument by the Applicant, the ‘454 Patent was allowed and issued.

2. The proposed combination of Vangemert and Copeland, would render Vangemert inoperable for its intended purpose of providing a low latency algorithm for buffer credit calculation according to the Fibre Channel protocol.

91. Even if a POSITA were motivated to make the illogical combination of Vangemert and Copeland, the out-of-band RS-232 signals taught in Copeland are fundamentally incompatible with the addressable data frames taught in Vangemert. Copeland describes the use of the DTR (Data Terminal Ready) or the RTS (Request To Send) signals to signal the transmitter. The DTR signal is a key component in RS-232 serial communication, used to indicate the readiness of a data terminal, like a computer or modem, for communication. In RS-232, a

standard for serial data transmission, the DTR signal is sent from the data terminal equipment (DTE) to the data communication equipment (DCE). When the DTE is prepared to establish a connection, it sets the DTR signal to a high voltage level, signaling its operational status. Conversely, a low voltage level indicates that the DTE is not ready or is turned off. This signal plays a pivotal role in the handshaking process, where it interacts with the Data Set Ready (DSR) signal from another device to control data flow. Communication can only proceed when both DTR and DSR signals are active. (See Ex. 2009)

92. Similarly, the Request To Send (RTS) signal is a critical control signal in RS-232 serial communication used for managing data flow between devices. It operates primarily between data terminal equipment (DTE), like a computer, and data communication equipment (DCE), such as a modem. The essence of RTS is to signal the intent of the DTE to send data. When the DTE has data ready for transmission, it activates the RTS signal by setting it to a high voltage level. This action signals the DCE to prepare for receiving data. The DCE, in turn, acknowledges this by sending back a Clear To Send (CTS) signal, confirming its readiness to accept the data. This interaction is part of hardware flow control, ensuring that data transmission occurs only when both devices are prepared, thereby preventing data loss and ensuring synchronization. (See Ex. 2009)

93. Accordingly, both RTS and CTS are concepts in RS232. RS-232, short for Recommended Standard 232, is a standard for serial communication that is widely used for connecting computers to various peripheral devices. Developed in the 1960s, it was one of the earliest protocols for serial data exchange between a terminal and a modem or other peripheral devices. RS-232 operates by transmitting data one bit at a time over a single wire, although the standard, as a whole, uses a variety of signals and control pins in a connector, most commonly with DB-9 or DB-25 interfaces. The primary use of RS-232 includes connecting devices like printers, modems, mice, and data storage devices to computers. Due to its simplicity and widespread adoption in the early days of computing, RS-232 became a foundational standard in computer communications. (See, Ex. 2009)

94. RS-232 operates in a full duplex mode. One serial link can be used for transmitting data (simplex) while simultaneously receiving data on the other serial link. Data is serially transmitted one symbol at a time on the TXD (Transmit Data) line and received serially on the RXD (Receive Data) line. The serial data is typically encoded for transmission and decoded upon reception. (See, Ex. 2009)

95. The technical specifications of RS-232 are designed for short-distance communication, typically up to a maximum of 15 meters, due to signal degradation over longer distances. The standard defines both the electrical characteristics and the connector types, ensuring interoperability between devices.

In RS-232 communication, the voltage levels define logical states, where a voltage in the range of +3 to +15 volts represents a logical 0 (“off” or “space” state), and -3 to -15 volts represents a logical 1 (“on” or “mark” state). This inverse logic might seem counterintuitive, but it was part of the design to accommodate certain technical aspects of early electrical engineering. (See Ex. 2009

96. DTE, DCE, CTS, and RTS are asynchronous level sensitive signals which indicate the status of the equipment based on the voltage level. Level sensitive signals differ from a serial bitstream where the symbol on the line can change sequentially with each symbol time. (See Ex. 2009)

97. On the other hand, Vangemert teaches an I/O bus that uses a higher level protocol, such as SCSI, that is transmitted using the Fibre Channel protocol. SCSI, which stands for Small Computer System Interface, is a set of standards for connecting and transferring data between computers and peripheral devices. Developed in the early 1980s, SCSI is notable for its flexibility and high performance, making it a popular choice for connecting a wide range of peripherals, including hard drives, printers, and scanners. Unlike simpler communication protocols like RS-232, SCSI can support multiple devices in a daisy-chain or bus configuration, allowing several devices to be connected to a single SCSI port. The standard has evolved through various versions, each

offering improvements in speed, device support, and cable length. SCSI interfaces use parallel communication, meaning they transmit multiple bits of data simultaneously across several channels, leading to significantly higher data transfer rates compared to serial communication methods. (See, Ex. 2010)

98. SCSI is recognized for its robust performance, particularly in server environments and high-performance computing applications. It offers several advantages, such as fast data transfer rates (ranging from 5 MBps in early SCSI-1 to several hundred MBps in later versions like Ultra-640 SCSI), the ability to connect multiple devices, and reliable data integrity checks. Setting up SCSI systems, however, can be more complex than simpler interfaces due to the need for proper termination at the end of the SCSI chain and the configuration of unique device IDs for each connected device. The cables and connectors used in SCSI systems vary depending on the version, with earlier versions using 50-pin connectors and later ones employing 68-pin or 80-pin connectors to accommodate more data channels. (See, Ex. 2010)

99. Further, a POSITA would have understood that the Fibre Channel protocol is a high-speed network technology primarily used for storage networking, known for its ability to transport large volumes of data reliably and at high speeds. Developed as a standardized protocol (ANSI X3.230-1994), Fibre Channel can operate at speeds ranging from 1 to 128 gigabits per second, making

it particularly suited for connecting computer data storage to servers in storage area networks (SANs). It supports various topologies, including point-to-point, switched fabric, and arbitrated loop, offering flexibility in network design. This protocol is widely used in data center environments for its robust performance, particularly in scenarios requiring fast access to large amounts of data, such as in database management, backup, and recovery operations. (See, Ex. 2012)

100. An analogous combination argument to Dobson and Copeland was made in the '454 prosecution history by the Patent Examiner. The Patent Examiner cited U.S. Patent 5,613,141, "Szatkowski" in view of Daines et al. (U.S. Patent 6,192,422, "Daines." Ex. 1002, 182-186. This argument was overcome by the Applicant by highlighting the non-analogous art of dedicated, point-to-point serial links versus higher level protocols used in local area networks, "The proposed combination of Szatkowski in view of Daines is improper, as Szaatkowski and Daines are non-analogous art. Szatkowski is directed to "dedicated, point to point... serial links" (Szatkowski, Abstract, inter alia). In contrast, Daines is directed to "a local area network" of Daines. For example, the "local area network" of Daines requires "arbitration" and "an arbitration mechanism" (Daines, Abstract, inter alia), wherein the "dedicated, point to point... serial links" of Szatkowski have no such need." Ex. 1002, at 226-228.

Subsequent to this argument by the Applicant, the ‘454 patent was allowed and issued.

101. Accordingly, a POSITA would have appreciated that it would not have been possible to transmit the signaling (e.g., +3 to +15 volts) taught with regard to Copeland’s CTS and DTR signal over the higher level I/O bus, such as Fibre Channel protocol or SCSI, that transmits addressable data frames as taught in Vangemert. For example, a POSITA would appreciate that systems utilize CTS and DTR of RS232 are incompatible with the lower layers of the Fibre Channel protocol, FC-0 and FC-2 and are not operable in Fibre Channel communication links.

3. Petition Fails to Demonstrate That Vangemert or Copeland Transmits the Flow Control Signal Through the Serial Link and Connector Without Passing the Flow Control Signal Through the Second Buffer

A. The Petitioner Fails to Demonstrate That Vangemert Transmits the Flow Control Signal Without Passing the Flow Control Signal Through the Second Buffer.

102. Vangemert makes a clear distinction between buffer credit signals and buffer credit control words. Specifically, Vangemert describes that “Using control words, an initiating node first arbitrates to get control of the loop, then notifies a responding node of its desire to transmit data frames.” (Ex. 1005, 3:10-13) Vangemert further explains that the “Transmit controller 416 creates or

updates exchange information in exchange tables 416 in response to commands and data from the node's i/o bus, generates **control words** and data frames in transmit buffer 410, and forwards them to encoder 402.” (Ex. 1005, 3:57-61) In addition, Vangemert describes that “when receive controller 406 receives **a control word** indicating that a remote node wishes to transmit data, the buffer credit manager 418 is enabled, and the transmit controller 412 begins transmitting buffer credit **control words** in response to **a credit signal** from buffer credit manager 418.” (Ex. 1005, 4:11-16). A POSITA would understand that the “credit control words” are specific type of the control words described in Vangemert and transmitted by the Transmit controller 416 via the transmit buffer 410.

103. Accordingly, Vangemert clearly explains that the transmit controller transmits **credit control words** to devices connected to the I/O bus via the transmit buffer (410) in response to the **internal credit control signal** generated by the comparator (502). The reference is wholly silent with regard to the internal credit **control signal** being sent over the I/O bus. In fact, such a **control signal** could not be physically communicated over the I/O bus taught in Vangemert because Vangemert teaches that I/O bus uses a higher-level protocol such as SCSI.

104. The Petitioner states “Vangemert does not disclose or suggest that the buffer credit signal or associated buffer credit control words are ever stored in

or passed through the transmit buffer 410 (second buffer).” (Petion, 31) This statement is false.

105. Vangemert clearly states that **“A transmit controller 412 retrieves control words and data frames from transmit buffer 410 and forwards them to encoder 402.”** (Ex. 1005, 3:49-52) Vangemert states this again as follows: **“Transmit controller 416 creates or updates exchange information in exchange tables 416 in response to commands and data from the node’s i/o bus, generates control words and data frames in transmit buffer 410, and forwards them to encoder 402.”** (Ex. 1005, 3:57-61). A POSITA would understand that Vangemert's transmit buffer 410 is simply a memory used for temporary storage which could not on its own originate control words. In interpreting Vangemert's disclosure that transmit controller 416 generates control words "in transmit buffer 410," a POSITA would therefore understand that transmit controller 416 is generating control words that are input to transmit buffer 410 before being forwarded to encoder 402.

106. The path of the transmit buffer, which contains the credit control words to the output transmitter, is shown in Vangemert’s Figures 3 and 4. Vangemert makes a clear distinction between internal buffer credit signals and buffer credit control words: **“the buffer credit manager 418 is enabled, and the transmit controller 412 begins transmitting buffer credit control words in response to a credit signal from buffer credit manager 418.”** (Ex. 1005, 4:13-16) Vangemert clearly

states that data frames and control words are transported on the transmit channel from source (transmitter) to destination (receiver): “It is sufficient to stipulate that the fabric switch 114 transports data frames and control words from the transmit channel of a source node to the receive channel of a destination node” (Ex. 1005, 2:65-3-1), and “Receive controller stores received control words and data frames in receive buffer 408” (Ex. 1005, 3:46-48). The coupling of the control words from the encoder to the transmit channel is described as: “Transceiver 300 also includes a parallel-to-serial converter 308 which receives a stream of codewords from an encoder and uses a clock signal from clock multiplier 310 to convert the stream of codewords into a serial bitstream. Transmitter 312 then converts the serial bitstream into signals suitable for transport by the transmit channel.” (Ex. 1005, 3:31-36) Vangemert discusses no other payloads on the transmit channel besides data frames and control words.

107. A POSITA would understand without doubt that the buffer credit control words provide the flow control signals for transport on the link. It is clearly stated that the transmit controller generates control words in the transmit buffer and redundantly stated that the transmit controller retrieves control words from the buffer and forwards them to the encoder for eventual transport on the transmit channel. A POSITA would understand that the Vangemert’s flow control signal is not transmitted over the I/O bus. Additionally, it is reiterated that the transmit

controller retrieves these control words from the buffer and forwards them to the encoder for transmission on the transmit channel. From this, a POSITA would clearly comprehend that Vangemert's credit control signal is not transmitted over the I/O bus (serial link).

B. Petitioner Fails to Demonstrate That Copeland Transmits the Flow Control Signal Through the Serial Link as recited in the Challenged Claims


108. The '454 Patent and its related patents, '214, '675, and '752 all have the same lead inventor, Frank Ahern ("Ahern"). As discussed above, the '214 and the '752 Patents are incorporated by reference. The serial link contemplated by '454 Patent family is the duplex serial connection that passes data, also commonly referred to as a serial data link. This has been consistently shown by context in the related patent specifications which are incorporated by reference in the '454 Patent. For example, the '214 Patent states that: "In one preferred embodiment, two buses communicate over a duplex link formed with a pair of simplex links, each employing twisted pair or twin axial lines (depending on the desired speed and the anticipated transmission distance). Information from the buses are first loaded onto FIFO (first-in first-out) registers before being serialized into frames for transmission over the link" (Ex. 1010, 4:18-24), and "Frames produced by encoder 28 are forwarded through transmitter 44 along simplex link 46 to receiver 48, which

supplies the serial information to decoder 30. Likewise, encoder 26 forwards serial information through transmitter 38 along simplex link 40 to receiver 42, which supplies the serial information to decoder 32.” (Ex. 1010, 6:40-46).

109. The ‘214 Patent further describes that “Flow control may be necessary should FIFO registers 22 or 24 be in danger of overflowing. For example, if FIFO register 22 is almost full, it supplies a threshold detect signal 36 to encoder 26, which forwards this information through link 40 to decoder 32” (Ex. 1010, 6:47-51). A POSITA would understand that the flow control signal is encoded and transmitted as part of a data frame, and then decoded. The flow control signal is transmitted with the data on a serial data link.

110. For further support of the meaning of a serial link, ‘the 214 provides Figure 4 and its description, “Referring to FIG. 4, the previously mentioned simplex links 40 and 46 are shown as twin axial lines 40A and 46A, wrapped with individual shields 40B and 46B. A single shield 122 encircles the lines 40 and 46. Four parallel wires 124 are shown (although a greater number may be used in other embodiments) mounted around the periphery of shields 122 for various purposes. These wires 124 may carry power management signals, dock control signals or other signals that may be useful in an interface between a docking station and a portable computer. While twin axial lines offer high performance, twisted pairs or other transmission media may be used in other embodiments where the transmission distance is not as

great and where the bit transfer speed need not be as high. While a hard wire connection is illustrated, in other embodiments, a wireless or other type of connection can be employed instead.” (Ex. 1010, 9:3-18).



111. Figure 4 (reproduced above) of the ‘214 Patent provides a clear visual representation of the twin axial lines forming each simplex link. The duplex link is formed by a receive simplex link and a transmit simplex link. A POSITA would understand that twin axial cables and twisted pair cables are used for transmitting a differential signal or at least a single-ended signal plus a ground reference.

112. In ‘the 214 Patent, like in the ‘454 Patent, any signal transmitted through the serial link must go through one of these two simplex links, depending on the direction of communication (transmit or receive). There are no other links in the description or figure, but ‘214 Patent does mention the possibility of including

other parallel wires in the cable, which can be useful for interfacing a docking station and a portable computer. While these parallel wires are part of the cable and part of the connector, they are not included in the serial link as contemplated by Ahern. To be clear, any signal transmitted through the serial link described in the ‘454 Patent must be serialized and transmitted along the simplex link.

113. Finally, the ‘214 Patent, like the ‘454 Patent, also contemplates serial links implemented in optical fiber or wireless technology: “FIG. 4 is a block diagram of the modular computer system of the present invention utilizing a universal connectivity station (UCS) communicably coupled to a plurality of devices via serial links, such as the Split-Bridge™ technology serial links employed using fixed wire, optical, or wireless communication links.” A POSITA would understand that signals transmitted through these media optical or wireless communication links would all be in-band and communicated through the serial link, not through a parallel wire.

114. The Petitioner states that “In Copeland, data flow control is activated when the FIFO buffer is forty percent full.” The data flow control is executed by sending a low Request to Send (“RTS”) signal over the connector 20 to instruct a remote device to stop sending data upon the FIFO buffer reaching the fill amount. (Petition, 28-29). The Petition contends that “the RTS signal is a flow control signal

because it is generated to instruct a remote device to stop sending data.” (Petition, 28-29).

115. Copeland describes that “Serial port connectors 20 and 21 preferably define an EIA RS-232-D serial interface. UARTs 14 and 15 convert data between a serial data format on connectors 20 and 21 and a parallel data format on bus 40” (Ex. 1006, 4:33-37), and “The preferred embodiment of the present invention has a 16550A UART.” (Ex. 1006, 2:41-42). A POSITA would know that RS-232 was a common serial interface specification, and the 16550A was commonly used to provide serialization and deserialization for an RS-232 application.

116. A POSITA would also know that RTS is a parallel out-of-band handshake signal used in RS-232. The RS-232 RTS signal is on a separate out-of-band connector pin from the serial data signals, RXD or TXD. Therefore, the flow control signal for Copeland is not transmitted through the serial link.

117. Accordingly, the Petition seems to confuse the serial link with the connector 20. As discussed, the RTS signal is not part of the serial link, even if it was included in the same connector.

4. Petitioner’s Stated Motivation to Combine Vangemert and Copeland and Minimize Latencies is Incorrect

118. Petitioner states “If Copeland’s teachings were used to improve Vangemert, a PHOSITA would similarly strive to minimize latencies by routing the

interrupt signal directly to the transmitting device, as disclosed by Copeland.” (Petition, 31).

119. This statement is incorrect because Copeland’s teachings do not minimize latencies over Vangemert; rather, Copeland’s teachings increase latencies. Copeland’s solution, “data flow control is activated when the FIFO buffer is forty percent full” (Ex. 1006, 5:39-40), and “if the buffer is not emptied then, after a predetermined count, the DTR and/or RTS lines are de-asserted, thereby stopping the flow of data from the remote transmitter until the buffer has been emptied” (Ex. 1006, 7:23-26) has a negative effect on serial link latency and performance.

120. Copeland adds “The present invention therefore automatically signals the remote transmitter to stop sending when the FIFO buffer fills to a certain point and to resume sending once the FIFO buffer has been completely emptied.” (Ex. 1006, 3:3-6) Therefore, once a buffer reaches its certain fill point, the data transmission must be suspended until the buffer is completely emptied. The transmission latency is dominated by the buffer empty time. Large fill points will result in longer latencies, and small fill points will result in inefficient use of the buffer and higher costs. As previously stated, combining Vangemert and Copeland would add a flow control signal to prevent pending receiver buffer overflow, whereas Vangemert’s solution eliminates receiver buffer overflow, so the

combination makes no sense. Even if the combination were pursued, the incorporation of Copeland's flow control would increase latencies in contrast to what the Petitioner has proposed. There would be no motivation to combine Vangemert and Copeland with regard to latencies.

121. An analogous combination argument to Vangemert and Copeland was made in the '454 prosecution history by the Patent Examiner. The Patent Examiner cited U.S. Patent 5,613,141, "Szatkowski" in view of Daines et al. (U.S. Patent 6,192,422, "Daines." (Ex. 1002, 182-186) The Applicant overcame this Argument during the final prosecution. " Szatkowski teaches a "flow control" mechanism that prevents buffer overflow. Szatkowski's use of an acknowledgement-based protocol may eliminate most or all occurrences of buffer overflow or "near" overflow," and "Daines teaches: If (the amount of data in an input buffer crosses a first threshold level), the flow control mechanism inserts congestion data into an output buffer for the port to alert the associated node to cease transmitting data packets to the port," (Ex. 1002, at 226-227) and "Further, the acknowledgement-based protocol of Szatkowski may eliminate most or all occurrences of buffer overflow, or "near" overflow, rendering unmodified Szatkowski superior to Daines in terms of flow control. Thus, the proposed modification in view of Daines may actually be detrimental to Szatkowski, e.g., by eliminating the advantageous acknowledgement-based protocol of unmodified

Szatkowski.” (Ex. 1002, at 227-228) Subsequent to this final argument by the Applicant, the ‘454 patent was allowed and issued.

122. Finally, the proposed combination also fails to achieve the desired goal of ‘454 Patent, which explains that “the present applicant, actually had to go back to the drawing board and design a far simpler serial protocol to allow a marriage to the PCI bus, because none of the existing implementations could coexist without substantial loss of performance” Ex. (1001, 2:54-58). A POSITA would have known that the “performance” of a serial communication network includes attributes of data rate, latency, and transmission efficiency. In contrast, the ‘454 Patent mitigates the latency impact of flow control and achieves higher performance.

B. Petition Fails to Prove That Independent Claim 1 is Obvious Over Dobson and Copeland

123. The Petition fails to prove that Independent Claim 1 is Obvious over Dobson and Copeland for at least the following reasons.

124. First, the proposed combination of Dobson and Copeland does not teach or suggest “wherein the module is further configured to transmit the flow control signal through the serial link and the connector without passing the flow control signal through the second buffer,” as recited in the challenged claims.

125. Second, the Petitioner’s proposed mapping of the recited “second buffer” is inconsistent and, therefore, does not show how the proposed combination of Dobson and Copeland teaches every element recited in the challenged claims.

1. Petition Fails to Demonstrate That Dobson or Copeland teaches a module configured to transmit the flow control signal through the serial link and the connector without passing the flow control signal through the second buffer,” as recited in the challenged claims.

126. Dobson describes a system where a “level detect circuit 130 is provided with the write address bus 55’ and read address bus 65’.” (Ex. 1007, 8:32-33) Dobson further describes that:

The level detect circuit 130 determines when data is to be written from the receive FIFO buffer 100, as is discussed below. The level detect circuit 130 is provided with the write address bus 55’ and read address bus 65’. The level detect circuit 130 outputs receive data ready signal on line 132 to bit 0 in the UART line status register 134. Bit 0 in the UART line status register 134 is passed to the host computer 8 on line 136 as an interrupt request. The host computer 8 responds to the interrupt request on line 136 by reading data from the receive HFO buffer 100, as set forth below. Ex. (1007, 8:30-39)

127. Dobson further describes that “Received data is passed from the receive FIFO buffer 100 to the host personal computer 8 via the receive data bus 26d, bus interface circuit 20, and data bus 12.” (Ex. 1007, 7:61-64)

allows a special program, an interrupt handler or interrupt service routine, to run instead. This mechanism is a critical part of a computer's architecture, allowing it to respond to events from hardware devices like keyboards, mice, network adapters, and disk drives. (See, Ex. 2011)

131. When a device needs the CPU's attention, it sends an IRQ signal. The CPU responds by suspending its current activities, saving its state, and executing the interrupt handler for that specific IRQ. This allows the device to be serviced immediately, which is crucial for time-sensitive tasks. After the interrupt handler has completed its task, the CPU resumes its previous activity. (See, Ex. 2011)

132. There are a limited number of IRQ lines in a computer, and historically, managing these lines (especially in systems like the IBM that were commercially available in 1998) could lead to conflicts if multiple devices were set to use the same IRQ line. (See, Ex. 2011)

133. Accordingly, it makes perfect sense that Dobson teaches to use a separate line (136) that is discrete from the serial bus (14) or the parallel received data bus (26d) to alert the computer. Therefore, Dobson does not expressly or inherently teach a module configured to transmit the flow control **signal through the serial link** and the connector without passing the flow control signal through the second buffer," as recited in the challenged claims.


134. The additional teachings of Copeland do not cure the shortcomings of Dobson. Instead, as discussed above, Copeland's flow control signal is an out-of-band parallel signal transmitted through a separate signal line and not through the serial link.

135. Therefore, if Dobson were to be combined with Copeland as proposed by the Petitioner, the out of band signaling taught in Copeland (e.g., CTS and DTR signal) would be sent over the separate IRQ line and not through the serial link. Alternatively, if Copeland were combined with Dobson, the interrupt taught in Dobson would be transmitted over the CTS and/or DTR lines taught in Copeland and not through the serial link as recited in the Challenged Claims.

136. As a result, Dobson, alone or in combination with Copeland, fails to teach or suggest "a module configured to transmit the flow control signal through the serial link and the connector without passing the flow control signal through the second buffer," as recited in the challenged claims.


2. Petitioner's proposed mapping of the recited "second buffer" is inconsistent and, therefore, does not show how the proposed combination of Dobson and Copeland teaches every element recited in the challenged claims.

137. On Pg. 58 of the Petition, the Petitioner maps the "second buffer" to TX FiFO Buffer (40) of Dobson. Petitioner's proposed mapping is shown in their annotation of Dobson's Figure 2 (reproduced below).



Dobson FIG. 2 (Annotated)

138. Then, on page 66 of the Petition, Petitioner changes the proposed mapping of the “second buffer” to Copeland’s Buffer 13. Petitioner’s proposed mapping is illustrated in Petitioner’s annotated version of Copeland’s Figure 1 (reproduced below)



Copeland FIG. 1 (Annotated)

139. The Petition states that “Copeland’s buffer 13 equates to the claimed “second buffer” because buffer 13 controls data flow to and from the host, including to UART 14, for converting parallel data to serial data on line 40.” (Petition, 67). However, the Petitioner and the Petitioner’s expert are both completely silent as to how Copeland’s buffer 13 is the equivalent of Dobson’s TX FiFO Buffer (40).

140. In fact, Copeland expressly teaches that “Host data buffer 13 is a **bidirectional** buffer which controls the flow of data to and from host 10.” (Ex. 1006, 4:7-9) Accordingly, Petitioner has offered no explanation for why a POSITA would equate Dobson’s unidirectional buffer (40) with Copeland’s bidirectional buffer. In fact, the only reason for a POSITA to equate the two buffers would be to use the ‘454 Patent as a blueprint.

XII. CONCLUSIONS

141. Relying on my education, experience, skill, and training in the relevant field of art, and after reviewing the materials identified above, it is my opinion that:

- Petitioner has not demonstrated that the Challenged Claims are obvious based on the grounds identified in the Petition.

142. I hereby 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.

143. I reserve the right to supplement my opinions in the future to respond to any arguments that Petitioner raises and to take into account new information as it becomes available to me.

144. In signing this Declaration, I understand that the Declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I acknowledge that I may be subject to cross examination in the case and that cross examination will take place within the United States. If cross examination is required of

me, I will appear for cross examination remotely within the United States during the time allotted for cross examination.

Executed January 10, 2024 at McCall, Idaho.


Terry R. Lee