

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner

IPR2023-01047
U.S. Patent No. 7,921,323

**DECLARATION OF DANIEL BLUMENTHAL
UNDER 37 C.F.R. § 1.68 IN SUPPORT OF PETITION
FOR *INTER PARTES* REVIEW**

TABLE OF CONTENTS

I.	INTRODUCTION	4
II.	QUALIFICATIONS AND PROFESSIONAL EXPERIENCE	7
III.	LEVEL OF ORDINARY SKILL IN THE ART	14
IV.	RELEVANT LEGAL STANDARDS	15
V.	BACKGROUND	16
VI.	OVERVIEW OF THE '323 PATENT	36
A.	Summary of the Patent	36
B.	Prosecution History of the '323 Patent	37
VII.	CLAIM CONSTRUCTION	38
VIII.	CLAIMS 27-28, 31, AND 33 ARE UNPATENTABLE	43
IX.	GROUND 1: HARTSELL RENDERS OBVIOUS CLAIMS 27-28, 31, AND 33	45
A.	Detailed Analysis of Claims	45
1.	Claim 27	45
2.	Claim 28	78
3.	Claim 31	79
4.	Claim 33	83
X.	GROUND 2: HARTSELL AND HAUCK RENDERS OBVIOUS CLAIMS 27-28, 31, AND 33	84
A.	Reasons to Combine Hartsell and Hauck	84
B.	Detailed Analysis of Claims	85
1.	Claim 27	85
2.	Claim 28	92
3.	Claim 31	92
4.	Claim 33	93

XI.	GROUNDS 3 AND 4: THE COMBINATION OF HARTSELL AND RAJ AND THE COMBINATION OF HARTSELL, HAUCK, AND RAJ RENDER OBVIOUS CLAIMS 27-28, 31, AND 33	93
A.	Reasons to Combine Hartsell and Raj.....	93
B.	Detailed Analysis of Claims.....	94
1.	Claim 27.....	94
2.	Claim 28.....	105
3.	Claim 31	107
4.	Claim 33.....	110
XII.	CONCLUSION.....	111

I, Daniel Blumenthal, do hereby declare as follows:

I. INTRODUCTION

1. I am making this declaration at the request of Cisco Systems, Inc. in the matter of the *Inter Partes* Review of U.S. Patent No. 7,921,323 (“the ’323 Patent”) to Yancey et al.

2. I am being compensated for my work in this matter at my standard hourly rate. I am also being reimbursed for reasonable and customary expenses associated with my work and testimony in this proceeding. My compensation is not contingent on the outcome of this matter or the specifics of my testimony.

3. I have been asked to provide my opinions regarding whether the subject matter of claims 27-28, 31, and 33 (“the Challenged Claims”) of the ’323 patent would have been obvious to a person having ordinary skill in the art (“POSITA”) at the time of the alleged invention, in light of the prior art. It is my opinion that the Challenged Claims would have been obvious to a POSITA.

4. In the preparation of this declaration, I have relied on:

- the ’323 patent, Ex. 1001;
- the prosecution history of the ’323 patent (“’323 File History”), Ex. 1002;
- U.S. Patent Publication No. 2002/0059274 to Hartsell et al. (“Hartsell”), Ex. 1005;

- U.S. Patent No. 6,496,291 to Raj et al. (“Raj”), Ex. 1007; and
- Scott Hauck, “The Roles of FPGA’s in Reprogrammable Systems,” Proceedings of the IEEE, vol. 86, no. 4 (April 1998).

5. In forming the opinions expressed below, I have considered the documents listed above; the relevant legal standards, including the standard for obviousness; and my own knowledge and experience based upon my work in the field of field of communication networks as described below, as well as portions of the following additional materials:

- U.S. Patent No. 6,907,595 to Curd et al., Ex. 1006;
- Microsoft Press, *Computer Dictionary*, Third Edition, 1997, Ex. 1009;
- “Virtex-II Pro™ Platform FPGAs: Introduction and Overview,” Xilinx, March 24, 2003, Ex. 1010;
- “Virtex-II Pro™ Platform FPGA User Guide,” Xilinx, June 30, 2003, Ex. 1011;
- U.S. Patent No. 7,680,054 to Acharya, Ex. 1013;
- U.S. Patent No. 7,039,057 to Acharya et al., Ex. 1014;
- U.S. Patent No. 7,382,787 to Barnes et al., Ex. 1015;

- L. Kazovsky et al., “Optical Fiber Communication Systems,” 1996, Ex. 1021;
- J. Goodman, “Optical Interconnection in Microelectronics,” SPIE Vol. 456 Optical Computing, 1984, Ex. 1022;
- R. Michalzik, “Optical Backplanes, Board and Chip Interconnects,” Fiber Optic Data Communication: Technological Trends and Advances, 2002, Ex. 1023;
- K. Hahn et al., “Gigabyte/s Data Communications with the POLO Parallel Optical Link,” Electronics Components and Technology Conference, 1996, Ex. 1024;
- D. Blumenthal, “Routing Packets with Light,” Scientific American, Jan. 2001, Ex. 1025;
- M. O’Mahony et al., “The Application of Optical Packet Switching in Future Communication Networks,” IEEE Communications Magazine, Mar. 2001, Ex. 1026;
- N. Sherwani et al., “Introduction to Multichip Modules,” 1995, Ex. 1027;
- C. Maxfield, “The Design Warrior’s Guide to FPGAs: Devices, Tools, and Flows,” 2004, Ex. 1028;

- U.S. Patent No. 5,898,801 to Braun , Ex. 1029;
- J. Lockwood et. al., “Reprogrammable Network Packet Processing on the Field Programmable Port Extender (FPX),” FPGA 2001 (Feb. 11-12, 2001), Ex. 1031
- U.S. Patent No. 6,891,397 to Brebner, Ex. 1032; and
- U.S. Patent Publication No. 2004/0249964 to Mougel, Ex. 1033.

6. Unless otherwise noted, all **emphasis** in any quoted material has been added. Claim terms are *italicized*.

II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE

7. My complete qualifications and professional experience are described in my *Curriculum Vitae*, a copy of which can be found in Exhibit 1004. The following is a brief summary of my relevant qualifications and professional experience.

8. My education and my experience in this and related areas span 41 years of professional and academic experience.

9. I am considered a pioneer in optical communications, packet switching, and integrated photonic and optical devices and several other optical application areas, each of which have some relationship to the subjects involved.

10. I received a B.S. degree in Electrical Engineering from the University of Rochester, Rochester, New York, in 1981, an M.S. degree in Electrical

Engineering from Columbia University, New York, in 1988, and a Ph.D. degree in Electrical Engineering from the University of Colorado, Boulder, in 1993. From 1981 until 1984, I worked as an engineer at StorageTek, Louisville, Colorado, on one of the world's first optical data storage products. From 1984 to 1985 I worked as a Senior Engineer at Sievers Research in the area of optical spectroscopy. From 1985 - 1988, I worked at Columbia University as a research engineer in the Center for Telecommunications, in the area of ultra-fast optical fiber communications and optical packet switching while pursuing at the same time my Masters degree in Electrical Engineering. While at Columbia I built one of the world's first optical packet switches. The Telecommunications Center was one of the National Science Foundation centers of excellence that launched the modern form of telecommunications networks layered communications and spawned a wide range of research from multi-media communications, to packet communications, to optical communications and network layered architectures. From 1988 to 1990 I worked in the Distributed Computing Systems Laboratory at the University of Pennsylvania. From 1990 I worked as a graduate student researcher while pursuing my PhD in the Optoelectronic Computing Center at the University of Colorado Boulder, where I built one of the world's first self-routing optical packet switches. From 1993 to 1997, I was an assistant professor in the School of Electrical and Computer Engineering at

the Georgia Institute of Technology in Atlanta where I built a program in optical communications and packet switching and my focus on research in optical packet switching, optical network, optical wavelength conversion using semiconductor optical amplifiers (SOAs), optical label switching and other related technologies led to groundbreaking innovations in these areas.

11. I am currently a Distinguished Professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara (UCSB). I also serve as the Director of the Terabit Optical Ethernet Center (TOEC). My research areas are in optical communications, photonic packet switching and all-optical networks, all-optical wavelength conversion and regeneration, ultra-fast communications, indium phosphide Photonic Integrated Circuits (PICS) and nano-photonic device technologies.

12. I have authored or co-authored over 525 journal and conference publications, hold 23 issued US patents, and am a co-author of one book, Tunable Laser Diodes and Related Optical Sources (New York: IEEE–Wiley, 2005) as well as seven other book chapters. I have published two invited papers in the prestigious, broadly read, Proceedings of the IEEE and one invited paper in the prestigious broadly read Scientific American magazine. The Proceedings of the IEEE is the leading IEEE journal that provides in-depth review, survey, and tutorial coverage of

the technical developments across all engineering areas including electronics, electrical and computer engineering, and computer science and is ranked as one of the top journals by Impact Factor, Article Influence Score and serves as a resource for engineers around the world. The first invited paper published in the Proceedings was entitled “Photonic Packet Switches: Architectures and Experimental Implementations” in the 1994 special issue on Optical Computing Systems and is the highest cited paper in that special issue. The second invited paper in the Proceedings is in the 2018 Special Issue on Silicon Photonics entitled “Silicon Nitride in Silicon Photonics.” In 2000, I was invited to write a seminal paper entitled “Routing Packets with Light,” that appeared in the January 2001 issue of Scientific American, a popular science magazine that has had many famous scientists contribute, and with 170 years is the oldest continuously published monthly magazine in the United States.

13. I have served as the principal investigator for multiple large-scale research programs funded by the Defense Advanced Research Projects Agency (DARPA).

14. At the national scale, I contributed to multiple US based fiber optic Internet backbone projects. I served on the Board of Directors for National LambdaRail (NLR), a nation-wide fiber backbone, from 2003-2010. I was an elected

Board Member for the Internet-2 Architecture & Operations Advisory Council from 2008-2012.

15. In addition to my industry experience at StorageTek and Sievers, I have co-founded two companies, Packet Photonics, Inc. and Calient Technologies, Inc. Packet Photonics, Inc. provided network and data center operators with photonic integrated wavelength tunable solutions to implement, install, operate, and maintain networks. Calient Technologies, Inc. provides software enabled adaptive MEMS mirror based photonic switching systems that enable dynamic optical layer switching and optimization.

16. I was elected as a Fellow of the National Academy of Inventors (NAI) in 2017. From the NAI website: “The program has 1,060 Fellows worldwide representing more than 250 prestigious universities and governmental and non-profit research institutes. Collectively, the Fellows hold more than 38,000 issued U.S. patents, which have generated over 11,000 licensed technologies and companies, and created more than 36 million jobs. In addition, over \$1.6 trillion in revenue has been generated based on NAI Fellow discoveries. With the induction of the 2018 class, there are now more than 125 presidents and senior leaders of research universities and non-profit research institutes, 502 members of the National Academies of Sciences, Engineering, and Medicine; 40 inductees of the National Inventors Hall of

Fame, 57 recipients of the U.S. National Medal of Technology and Innovation and U.S. National Medal of Science, 34 Nobel Laureates, 3 Queen Elizabeth Prize for Engineering recipients, 304 AAAS Fellows, 200 IEEE Fellows, and 164 Fellows of the American Academy of Arts & Sciences, among other awards and distinctions.”

17. In 2020 I was the recipient of the prestigious Optical Society of America (now OPTIFCA) C.E.K. Mees medal. The medal was established in 1961 in memory of charter member C. E. K. Mees, to recognize achievements that exemplifies the thought that "optics transcends all boundaries". Prior to 2017 the medal was presented every two years since 1962, and it has been presented annually since 2017. Former recipients of the medal include Charles H. Townes, inventor of the laser and Nobel Laureate.

18. I am a Fellow of the IEEE (only 0.1% of the world-wide membership can be elected as Fellows in a given year: “According to IEEE Bylaw I-305.5, "The total number of Fellow recommendations in any one year must not exceed one-tenth of one percent of the IEEE voting membership on record as of 31 December of the year preceding.") and a Fellow of the Optical Society of America (OSA). I am the recipient of a 1999 Presidential Early Career Award for Scientists and Engineers (PECASE) from the White House. I also received a National Science Foundation Young Investigator (NYI) Award in 1994, and an Office of Naval Research Young

Investigator Program (YIP) Award in 1997. The Presidential White House PECASE Award citation reads “For outstanding research and creative impact in multi-wavelength optical techniques.”

19. I have had extensive contributions to the international scientific community in the form of internationally renowned special issue journals and international conferences. I served as the primary Guest Editor for the 2018 JSTQE Special Issue on “Ultra Low Loss Planar Waveguides and Their Applications,” and have served as the primary Guest Editor for other special issues including the 1998 IEEE Journal of Lightwave Technology special issue on “Photonic Packet Switching Systems” and the 2003 IEEE Journal Of Selected Areas In Communications special issue on “High-Performance Optical/Electronic Switches/Routers for High-Speed Internet.” I also served as one of the Guest Editors for the 2006 JSTQE Special Issue on “Ultra-Fast Integrated Photonics.” In terms of regular journal editing, I was an Associate Editor for the IEEE Photonics Technology Letters and an Associate Editor for the IEEE Transactions on Communications for many years. I have organized and served on committees for internationally renowned conferences and topical meetings including as General Program Chair for the 2001 OSA Topical Meeting on Photonics in Switching and as Program Chair for the 1999 Meeting on Photonics in Switching. I have also served on numerous other technical program committees including the

Conference on Optical Fiber Communications OFC (2018 – Present and 1997 – 2000), the Conference on Lasers and Electrooptics CLEO (1999-2000), the European Conference on Optical Communications ECOC (2004-2005) and SIGCOMM 2006.

20. In forming the opinions expressed in this Declaration, I have relied upon my education and my approximately 42 years of professional and academic experience.

III. LEVEL OF ORDINARY SKILL IN THE ART

21. I understand there are multiple factors relevant to determining the level of ordinary skill in the pertinent art, including (1) the levels of education and experience of persons working in the field at the time of the invention; (2) the sophistication of the technology; (3) the types of problems encountered in the field; and (4) the prior art solutions to those problems.

22. A POSITA in the field of the '323 patent, as of its earliest possible priority date of May 11, 2004, would have been someone knowledgeable and familiar with communications infrastructures and the integration of ASIC devices on a fabric interface. Such a POSITA would have a master's degree in computer science, computer engineering, electrical engineering, or equivalent training, and approximately four to six years of experience working in the field of communication networks. Additional work experience can substitute for specific educational background, and vice versa.

23. For purposes of this Declaration, in general, and unless otherwise noted, my statements and opinions, such as those regarding my own experience and what a POSITA would have understood or known generally (and specifically related to the references I consulted herein), reflect the knowledge that existed in the relevant field as of the priority date of the '323 patent.

IV. RELEVANT LEGAL STANDARDS

24. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the '323 patent, I am relying on certain basic legal principles that Cisco's counsel has explained to me.

25. I understand that prior art to the '323 patent includes patents and printed publications in the relevant art that predate the priority date of the '323 patent. For purposes of this Declaration, I am applying May 11, 2004, as the priority date of the '323 patent.

26. I have been informed by Cisco's counsel that a claimed invention is unpatentable as obvious under 35 U.S.C. § 103 if the differences between the claimed invention and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a POSITA. I have also been informed by Cisco's counsel that the obviousness analysis considers factual inquiries, including the level of ordinary skill in the art, the scope and content of the prior art, and the differences between the prior art and the claimed subject matter.

27. I have been further informed by Cisco's counsel that there are several recognized rationales for combining references or modifying a reference to show obviousness. These rationales include: (a) combining prior art elements according to known methods to yield predictable results; (b) simple substitution of one known element for another to obtain predictable results; (c) use of a known technique to improve a similar device (method, or product) in the same way; (d) applying a known technique to a known device (method, or product) ready for improvement to yield predictable results; (e) choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and (f) some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify the prior art or to combine prior art teachings to arrive at the claimed invention.

28. Also, I have been informed and understand that obviousness does not require physical combination/bodily incorporation, but rather consideration of what the combined teachings would have suggested to a POSITA at the time of the alleged invention.

V. BACKGROUND

29. Optical communications has proved a cost-efficient way to modulate and receive high-speed optical data (signals) over short to long distances since at least the 1970s. The basic elements of a point-to-point optic transmission link consist of an optical transmitter and optical receiver connected over an optical fiber or an optical

waveguide (See Fig. 3.1 Kazovsky1996 (Ex.1021)). Depending on the link length (distance) and data capacity (number of bits per second) optical amplification may be used during transmission. The optical fiber and its compact chip-based counterpart the optical waveguide, have proved to be a cost and space efficient way to move high-speed data over many distances, from on a chip, to between chips on a circuit board, to between circuit boards, to between data equipment racks, and beyond to national and global distances.




Figure 3.1 Block diagram of the basic IM-DD optical binary communication system.

Fig. 3.1 Kazovsky1996 (Ex.1021)

30. The most basic fiber link is an intensity modulated direct detection (IM-DD) binary communication optic link that consists of a transmitter, an optical fiber and a receiver, as shown in Fig. 3.1 Kazovsky1996 (Ex.1021). Electronic data to be transmitted either drives a semiconductor laser directly or drives an optical modulator that has a continuous wave (cw) or pulsed light source (e.g., a laser or light emitting diode LED) connected to its optical input. The current into the laser or the signal into the optical modulator impresses the data from the information source onto the optical

carrier in the form of binary on-off light modulation, and the modulated optical signal is input to an optical waveguide or an optical fiber. At the receiving end of the fiber a photodetector converts the on-off modulated light to an electrically modulated signal (current) that is converted to an electrically modulated voltage representative of the data. This analog received voltage signal is amplified and equalized to optimize the signal to noise ratio (SNR) and minimize signal distortion. The amplified-filtered signal is then passed to a digital binary decision circuit that converts the received analog voltage to a digital binary data stream that is input to an information sink.

31. The capacity of a fiber optic link can be increased by transmitting multiple of these IM-DD data streams in parallel over the same fiber by transmitting each channel on a different wavelength (optical frequency). Such a system is called wavelength division multiplexed (WDM) or frequency division multiplexed fiber optic transmission system. The single wavelength IM-DD link in Fig. 3.1 Kazovsky1996 can be replicated over multiple different wavelengths to build a WDM link as shown in Fig. 7.15 Kazovaky1996 (Ex.1021). Each of N transmitters, each operating at a different wavelength λ_i for $i = 1-N$, is connected to a wavelength division multiplexer that combines the modulated wavelengths onto the same fiber. A wavelength demultiplexer is used at the fiber output, and sends each wavelength to a dedicated receiver for that wavelength. The transmitter and receiver design

considerations for this architecture are somewhat the same as a single wavelength IM-DD link, except for details of added loss from the multiplexer/demultiplexer and other wavelength crosstalk and nonlinearity considerations. But essentially, each receiver detector is receiving only one data modulated wavelength and can be optimized for the power in that channel. When the number of optical wavelengths (frequencies) are increased to beyond approximately 32 in a single fiber, these systems are referred to as Dense WDM or DWDM. The parallelism of fiber optics can also be used to increased capacity by transmitting over parallel optical fibers, called space-division multiplexing (SDM), and combinations of WDM and SDM are also employed.




Figure 7.15 Diagram of a simple WDM system.

Figure 7.15 Kazovsky1996 (Ex.1021)

32. The capacity can be further increased using coherent optical communications techniques and more sophisticated modulation techniques for each

wavelength stream. A basic single wavelength coherent optical link, is illustrated in Fig. 4.1 Kazovsky1996 (Ex.1021). There are well-known tradeoffs between incoherent (described above) and coherent links, including cost, complexity, and power consumption. For example, a coherent link can tolerate higher fiber loss and can transmit at higher bit-rates over longer distances than an incoherent link.




Figure 4.1 Block diagram of a coherent optical communication system.

Figure. 4.1 Kazovsky1996 (Ex.1021)

33. It has been well-known since the early 1980s that optical transmission can provide high-speed connections on and between electronic chips. For example, optical waveguide interconnections were proposed to support high-speed on-chip communications delivery by Goodman in 1984 (see Figure 4, Goodman1984) (Ex.1022).




Figure 4. Index-guided interconnections.

Figure 4, Goodman1984 (Ex.1022)

34. Since before 2000 it has also been well-known that optical communications can be used for high-speed data delivery over a wide range of applications and distances from on-chip, as in Goodman1984, to multi-chip modules (MSMs) where one module contains multiple chips, chip-to-chip, board-to-board, rack-to-rack and beyond. For example Michalzik2000 summarizes such applications in quotes below, and is illustrated in Figure 6.1 Michalzik2000 (Ex.1023).

It has become customary to classify optical interconnects into distinct, albeit overlapping, categories. Some of those are illustrated in Fig. 6.1, ranging from the longest to the shortest transmission distance: ~ Rack-to-rack, also called frame-to-frame; ~ Board-to-board; ~ Multi-chip module (MCM)-to-MCM or intraboard; ~ Chip-to-chip on a single MCM; ~ Intra- or on-chip.




Fig. 6.1 Major optical interconnect categories.

Figure 6.1 Michalzik2000 (Ex.1023)

35. Examples of high-speed waveguide and fiber board-level optical interconnects have been well-known since at least 1996, as described for parallel fiber SDM links as shown in optic Figure 1 Hahn1996 (Ex.1024). In this type of

implementation, the optical transmitters and receivers are implanted at the circuit-board level.




Figure 1. Schematic of POLO-1 module

Figure 1. Hahn1996 (Ex.1024)

36. The implementation of parallel fiber SDM transmitters and receivers in small form factor optical pluggable modules has been driven by the need to provide reduced size, power, and cost in a field-replaceable modular form factor as described by Michalzik2000 in Figure 6.3 (Ex.1023).




Fig. 6.3 Infineon Technologies' parallel optical link (PAROLI) module with and without an attached 12-fiber ribbon cable [19] (left, © 2000 IEEE) and bottom view of an open module [19a] (right).

Figure 6.3 Michalzik2000 (Ex.1023)

37. While optical transmitters and receivers provide the necessary interface between electronic data switching and routing equipment and optical waveguides/fiber, these interfaces consume power and space and can be costly as the bandwidth and capacity scale. As the data capacity increases between electronic switching and routing equipment, it may be desirable to use methods to switch or route data, packets and connections using all-optical techniques instead of the power and space consuming conversion that occurs between electrical and optical signals.

38. Areas in optical communications that focus on achieving this goal include optical routing and optical packet switching. An example of an optical packet switching architecture and technology is described in Blumenthal2001 and interconnect processing devices that were used, as illustrated in Figure 1 Blumenthal2001 (Ex.1025). Packetized data that originates from electronic packet

Ex. 1003

switches and routers, that are transmitted on a optical fiber or waveguide, contains headers in front of the data that specify information about where the packet is to go, where it came from, as well as other information. When entering the optical router (Figure 1 Blumenthal2001 (Ex.1025)) the header is optically removed (erased) without converting the signal back to electronics and a portion of the signal is photodetected and received by control electronics to compute a new header, which is then added back to the optical packet while it is moving through the router. The “updated” packet with new header is then converted all-optically to a new wavelength of light and routed to an output fiber or waveguide.





Figure 1. Blumethal2001 (Ex.1025)

39. Such an optical packet router can be used to build larger networks of routers that are connected by, for example fiber optic links, as shown in Figure 2

O'Mahony2001 (Ex.1026). Such optical packet routers, also called optical packet switches (OPS) can be combined with other well-known all-optical circuit routing technologies like the optical crossconnect (OXC) to realize flexible circuit and packet routing optical transport networks (OTN) that interconnect electronic routing, switching, and transport equipment such as SONET/SDH, edge OPS, and label switched routers among others.




■ Figure 2. Applications of OPS as core and edge routers.

Figure 2 O'Mahony2001 (Ex.1026)

40. Optical nodes in the OTN and at the edge of the OTN can be built with a hierarchical combination of all-optical routing, switching, crossconnect and transport technologies as shown in Figure 3 O'Mahony2001 (Ex.1026). These solutions have offered ways and design choices with which to provide reconfigurability and packet and circuit connection granularity, to high-speed fiber

optic links connecting multiple electronic packet routing equipment and systems, without requiring intermediate conversion to electronics for the data.



■ Figure 3. Interfacing of the OPS with the OXC.

Figure 3 O'Mahony2001 (Ex.1026)

41. Field Programmable Gate Arrays (FPGAs) are integrated circuits that can realize large and complex digital functions and enable rapid, low cost testing of hardware and software concepts, with the power of many millions of logic gates and today with billions of transistors. FPGAs provide important flexibility of programmability that can be done in prototyping, engineering development, and out in the field. FPGAs are programmable on-chip network connections between various input and output logic and internal logic gates. Different vendors such as Xilinx and Altera offer FPGAs utilizing different types of on-chip interconnection network.

Ex. 1003

Since FPGAs can be reprogrammed to handled different types of processing and memory functions, they play an important role in communications equipment such as data switches and routers as well as fiber optic communications equipment such as transport interfaces and transmitters and receivers, and many other digital systems.

42. FPGAs are composed of two main components, gates and interconnects. FPGAs also have other special purpose gates, especially for time domain conversion and regeneration including elastic memory, 3R amplifiers and clock distribution functions. FPGAs can be used to realize almost any function, from simple Boolean operators (XOR, NOT, etc) to complex lookup tables. Look up tables (LUTs) enable arbitrary mapping between any specified inputs and outputs. Generic electrical FPGAs can have many 1000s of logic blocks which can be flexibly configured any logical function of a number of input e.g. NOT, NOR, XOR, NOR, muxes, comparators, registers and latches. An example logic block and interconnect are shown in Figure 1.

43. With FPGAs, the size of logic block vs. performance comes about from the size of logic block defines density area utilization in the FPGA and the performance of the FPGA can be affected by size of logic block. The optimal size of logic block is one that uses lesser number of logic blocks for a functional implementation while consuming as little space as possible. Typically the used

(active) logic area is typically less than total built-in logic area. The industry defines the effective logic area contains active logic and area of programmable connections, where the routing area is typically 70 to 90 percent of total area. In case of lookup table based FPGA, many designs have converged on a 4-input lookup table as it gives optimal results in terms of logic synthesized and area consumed.




Figure 1. Example FPGA (left) logic block and (right) interconnect matrix.

44. A top down view of the FPGA is shown in Figure 5.1 Maxfield 2004 (Ex.1028)), with the programmable logic blocks (see example left side of Figure 1) and the programmable interconnect.




Figure 5-1. Top-down view of simple FPGA architecture.

Figure 5.1 Maxfield2004 (Ex.1028)

45. The FPGA has been commercially available in its first simple programmable logic device form from Altera in 1984 and in a first generation of network interconnected logic gates FPGA, the XC2064 by Xilinx, in 1985. Xilinx, a well-known FPGA company, has continued to innovate and commercialize FPGA technology since the early days in 1985 through the 1990s and 2000s and to today. The Xilinx family of FPGAs, started to focus around Virtex-II Pro FPGA family architecture, which led to successive generations including the Virtex 5 and 6 to today Virtex 7 and Ultrascale+ designs. The FPGA design evolved to contain more, faster logic gates, on-chip system clock speeds, and faster input/output data rates. For a given transistor node technology (e.g. 45 nm nodes), the transistors can be fabricated at a certain density and run at a certain capacity and power dissipation, and with decreasing node size all of these metrics improve including density and power

efficiency. However, The FPGA on-chip electronic circuitry runs, over parallel busses, at a relatively low system clock rate compared to the input/output bandwidth (e.g. 100 MHz to order 1 GHz), and uses massive parallelism to increase the total processing capacity. In order for the FPGA to communicate with external chips and circuit board transmission lines/traces, electrical transceivers (transmitters and receivers) are used to communicate between the on-chip circuitry and the external circuitry as illustrated in Figure 4-24 Maxfield2004 (Ex.1028). In order match the FPGA internal bus speed with increasingly external high speed devices and communications systems such as optical communications, electrical speed multiplexers and demultiplexers, called SERDES, are used. The SERDES doubles and halves the bit rate, using multiple stages to decrease and increase by factors of powers of 2.




Figure 4-24: Using high-speed transceivers to communicate between devices.

46. Increased speed SERDES require the integration of higher speed transistor technology which typically lags behind the current generation FPGA

transistors technology. As such, monolithic (single chip designs) that can support capacity increase of the FPGA tends to lag the ability to make higher speed SERDES on separate chips. Hence, a higher speed generation of FPGA can more quickly be realized by utilizing multiple current generation FPGAs with their exiting SERDES and input/output (I/O) data rate interfacing to higher speed busses using new generation SERDES.

47. The use of multiple chips to solve this problem and emulate a single chip solution using a mix of current and next generation technologies, is handled using an approach and technology called multi-chip modules (MCMs). MCMs have been around since before 1995 and involve packaging multiple electronic die onto a common substrate (e.g semiconductor, ceramic) with internal wiring to make the multiple chips operate as if a single chip. The MCM can be packaged in a large electronics package with its own pins, operating as what looks externally like a single chip. In fact, today, many high performance chips, including Xilinx FPGAs, that communicate with high speeds and high capacities, are multichip modules that contain multiple interconnected building blocks including microprocessors, memory, interconnection chips, and high speed I/O or SERDES.




Figure 1.3 Sherwani1995 (Ex.1027)




Figure 1.19 Sherwani1995 (Ex.1027)

It was well-known that FPGAs, as described above, would incorporate elements of packet routers to process and route packets between devices. By the early 2000s

FPGAs reached the level of maturity and performance to allow the design and incorporation of a wide range of functions including packet routing related functions.

48. A well-known and widely used FPGA-based platform was developed at Washington University and Stanford and published in well-known conferences and journals, enabling researchers, students, educators, and industry to implement a wide range of programmable packet routing functions as described in Lockwood2001. An FPGA was used to implement “Reprogrammable Network Packet Processing” on a field programmable port extender (FPX) for “processing of packets at the edge of a *multi-gigabit-per-second* network switch,” as described in Lockwood2001 (Ex.1031). This system, the Field Programmable Port Extender (FPX), “enables packet processing functions to be implemented as modular components in reprogrammable hardware.” As described in Lockwood2001 “All logic on the on the FPX is implemented in two Field Programmable Gate Arrays (FPGAs). Packet processing functions in the system are implemented as dynamically loadable modules.”

49. In Lockwood2001, the FPX is implemented on an FPGA called the Networking Interface Device (NID). The NID is able to transmit and receive packets over a network and as individual traffic flows. The NID is able to implement a wide range of switching functions as shown in Figure 6 Lockwood2001 (Ex.1031). These

functions include passing packets between a line card interface and a switch, processing of packets as they exit a router, implementing ingress cell processing using virtual circuits from a line card, loopback testing, and selective packet forwarding.




Figure 6: NID Switching Functionality

Ex. 1031, Fig. 6.

VI. OVERVIEW OF THE '323 PATENT

A. Summary of the Patent

50. The '323 patent describes “[r]econfigurable communications infrastructures...implemented to interconnect ASIC devices,” e.g., Field

Programmable Gate Arrays (“FPGAs”), “using high bandwidth interconnection mediums.” Ex. 1001, Abstract. “[H]igh speed serial I/O connections...may be employed to interconnect a pair of ASICs to create a high bandwidth, low signal count connection.” Ex. 1001, 2:53-57. Such systems “may be implemented using a predictable and uniform or standardized interface across the boundaries between each pair of board-level components (e.g., FPGAs, ASICs, general-purpose processors, etc.) to help promote consistent communications[.]” Ex. 1001, 4:15-19.

51. As I will explain below, these concepts were already known as of the priority date of the ’323 patent.

B. Prosecution History of the ’323 Patent

52. I understand that the ’323 patent issued on April 5, 2011, from U.S. Patent Application No. 11/600,934, filed November 16, 2006, as a continuation-in-part of U.S. Patent Application No. 10/843,226, filed May 11, 2004. Ex. 1002, 256, 259. The application faced a single office action, rejecting all claims under 35 U.S.C. § 112 because the use of “ASIC device” included “any device that has at least one ASIC chip and/or at least one form of programmable logic chip.” Ex. 1002, 61-63. All claims were also rejected as obvious by U.S. Patent No. 7,506,297 to Mukherjee, with the examiner taking official notice of missing elements. Ex. 1002, 64-67.

53. In response, the applicant made significant amendments to the independent claims. Ex. 1002, 40-42. The changes to independent claim 27 are provided below:

27. (Currently Amended) A communications infrastructure, comprising two or more ~~ASIC devices~~ separate signal processing circuits, each one of said two or more ~~ASIC devices~~ signal processing circuits including multiple ASIC devices that each itself includes ~~comprising~~ a packet router, said packet router of each one of said ~~two or more~~ ASIC devices of each given one of said ~~respective two or more signal processing circuits~~ being coupled through respective first and second common interfaces and an intervening high speed serial optical link to a each ~~respective~~ packet router of each of the ASIC devices of each ~~other of said two or more ASIC devices~~ signal processing circuits by an interconnection comprising a high speed serial optical link with no other processing device intervening between the high speed optical link and said ASIC devices of each of said two or more signal processing circuits.

Ex. 1002, 40-41.

54. In addition, the applicant argued that the use of the term “ASIC device” in the claims was clear and consistent with the definition of the term. Ex. 1002, 44-45, 51-52 (supporting arguments with “ASIC Definition from PC Magazine Encyclopedia”). The applicant also argued against the examiner’s use of official notice in the obviousness rejection. Ex. 1002, 46-49. A Notice of Allowance then issued with no additional statements or reasons for allowance provided by the examiner. Ex. 1002, 15.

VII. CLAIM CONSTRUCTION

55. It is my understanding that in order to properly evaluate the ’323 patent, the terms of the claims must first be interpreted. It is my understanding that for the

purposes of this *inter partes* review, the claims are to be construed under the so-called *Phillips* standard, under which claim terms are given their ordinary and customary meaning as would have been understood by a POSITA in light of the specification and prosecution history, unless the inventor has set forth a special meaning for a term. I have also been informed that claim terms only need to be construed to the extent necessary to resolve the obviousness inquiry.

56. I have reviewed the entirety of the '323 patent, as well as its prosecution history. It is my opinion that for purposes of applying the prior art presented herein to evaluate the patentability of the claims no term requires express construction. Accordingly, I have analyzed the claims consistent with their ordinary and customary meaning as would have been understood by a POSITA.

57. However, for sake of clarity, I note the following example embodiments of the claim term "*ASIC device*," as further supported by statements made during the prosecution of the '323 patent.

58. Claims 27 and 31 recite "*signal processing circuits including multiple ASIC devices*." ASIC stands for "Application Specific Integrated Circuit," and the '323 Patent explains that ASIC devices include "Field Programmable Gate Arrays ('FPGAs') or other field programmable devices ('FPDs') or programmable logic

devices ('PLDs')." Ex. 1001, 2:45-50; *see also* Ex. 1001, 4:50-53, 10:3-6, 14:49-54, 19:28-35.

59. During prosecution, the filed claims were rejected because the use of "ASIC device" included "any device that has at least one ASIC chip and/or at least one form of programmable logic chip," such as FPGAs. Ex. 1002, 61-63. In response, the applicant argued that the use of the term "ASIC device" in the claims was clear and consistent with the definition of the term, which encompasses FPGAs. Ex. 1002, 44-46 (applicant arguing that ASIC devices include FPGAs). In addition, the applicant provided a definition for ASIC from "PC Magazine Encyclopedia" to support ASIC devices as encompassing FPGAs. 51-52. That definition is provided below:

(Application Specific Integrated Circuit) Pronounced "a-sick." A chip that is custom designed for a specific application rather than a general purpose chip such as a microprocessor. The use of ASICs improve performance over general-purpose CPUs, because ASICs are "hardwired" to do a specific job and do not incur the overhead of fetching and interpreting stored instructions. However, a standard cell ASIC may include one or more microprocessor cores and embedded software, in which case, it may be referred to as a "system on a chip" (SoC).

A full custom ASIC chip is the most costly, and like standard cell ASICs, use a custom-designed mask for every layer in the chip. Unlike standard cells, designers of a full custom device have total control over the size of every transistor forming every logic gate, so they can “fine tune” each gate for optimum performance. Thus, a full custom ASIC performs electronic operations as fast as it is possible to do so, providing that the circuit design is efficiently architected.


Several Alternatives

Today, full custom ASICs represent a small percentage of the ASIC market because gate arrays, structured ASICs and standard cells turn circuit designs into working chips much faster and at much less cost. Such chips have greatly improved in speed over the years and provide the necessary performance for many applications. The speed advantage of a full custom ASIC is not as relevant as it was in the past. It is used primarily for devices such as microprocessors that must run as fast as possible and will be produced in huge quantities.

Also promoting the decline of full custom ASICs are chip manufacturers that make generic chips containing all the necessary functions for specific mass market products such as DVDs, CDs, digital cameras, etc. See structured ASIC,

gate array, standard cell, ASSP, PLO, adaptive computing
and CSIC.

Ex. 1002, 51. The applicant also provided the following figure with the above
definition of ASIC:



Ex. 1002, 52.

60. In addition, the '323 Patent provides an example where ASIC devices, provided in the form of FPGA devices, are implemented as “Xilinx Virtex-II Pro® XC2VP50 FPGA device[s].” Ex. 1001, 10:3-6, 10:22-24; *see also* Ex. 1001, 13:12-19. To the extent Patent Owner would seek to read this example into the claim

language, the references cited in this Declaration teach this limitation, as outlined in the analysis below.

61. The definition offered by the applicant in prosecution, discussed above, also defined an ASIC (and specifically, an FPGA) as a “chip.” To the extent Patent Owner would seek to limit an ASIC device to a single semiconductor chip, the references cited in this Declaration teach this interpretation.

VIII. CLAIMS 27-28, 31, AND 33 ARE UNPATENTABLE

62. This Declaration provides a detailed analysis of how the Challenged Claims are rendered obvious by the asserted prior art references, either alone or in combination, as follows:

- Ground 1: Hartsell renders obvious claims 27-28, 31, and 33.
- Ground 2: The combination of Hartsell and Hauck renders obvious claims 27-28, 31, and 33.
- Ground 3: The combination of Hartsell and Raj renders obvious claims 27-28, 31, and 33.
- Ground 4: The combination of Hartsell, Hauck, and Raj renders obvious claims 27-28, 31, and 33.

63. As part of my analysis, I have considered, and discuss in detail, the scope and content of the prior art and any differences between the alleged invention and the

prior art. It is my opinion that Hartsell, Hauck, and Raj are analogous art to the '323 patent because each is at least (1) in the same field of endeavor (reconfigurable communication infrastructures); or (2) reasonably pertinent to a problem with which the inventor of the '323 patent was involved (e.g., high speed serial I/O connections between ASIC devices, such as FPGAs, using a fabric interface). *Compare* Ex. 1001, Abstract (“Reconfigurable communications infrastructures may be implemented to interconnect ASIC devices”), 1:14-17 (relating to “interconnection of multiple ASIC devices” such as “Field Programmable Gate Array (FPGA) devices”), 2:41-50 *with* Ex. 1005, [0047], [0053] (providing an infrastructure using “multiple system engines...communicatively coupled” to perform “one or more dedicated tasks or sub-tasks associated with the delivery or transmittal of content from one or more content sources to one or more networks”); Ex. 1007, 1:55-67 (connecting circuits through high speed optical serial links); Ex.1030, Abstract (“Reprogrammable systems based on field programmable gate arrays are revolutionizing some forms of computation and digital logic.”), 618 (“With this rich sea of routing resources, the Xilinx 4000 series is able to handle fairly arbitrary routing demands, though mappings emphasizing local communication will still be handled more efficiently.”).

64. It is my opinion that the alleged invention recited in the Challenged Claims would have been obvious in view of the teachings of the asserted prior art and the knowledge of a POSITA.


IX. GROUND 1: HARTSELL RENDERS OBVIOUS CLAIMS 27-28, 31, AND 33

A. Detailed Analysis of Claims

1. Claim 27

- a.** [27.1] *A communications infrastructure, comprising two or more separate signal processing circuits,*

65. First, Hartsell teaches a content delivery system 1010 that “receives requests for content from remote browser clients via the network, accesses a file system to retrieve the requested content, and delivers the content to the client.” Ex. 1005, [0047]. Content delivery system 1010 “includes multiple system engines...communicatively coupled,” which perform “one or more dedicated tasks or sub-tasks associated with the delivery or transmittal of content from one or more content sources to one or more networks.” Ex. 1005, [0053]. Fig. 1A, below, provides an illustrated example of Hartsell’s content delivery system:



Ex. 1005, Fig. 1A (annotated).

66. As described in further detail below, Hartsell’s content delivery system teaches the claimed *communications infrastructure*.

67. **Second**, content delivery system 1010 contains multiple “processing modules that may be configured to operate as content delivery engines 1030, 1040, 1050, 1060, and 1070 communicatively coupled via distributive interconnection 1080.” Ex. 1005, [0143]. “Content delivery engines 1030, 1040, 1050, 1060 and 1070 are present to independently perform selected sub-tasks” where “any one or more of such subtasks may be combined and performed by a single engine, or subdivided to

be performed by more than one engine.” Ex. 1005, [0055]. Various numbers and configurations of content delivery engines are employed in content delivery system 1010 “to fit particular needs of a given application or installation.” Ex. 1005, [0054]. For example, Hartsell’s content delivery system includes one or more transport processing engines 1050, each of which is a *signal processing circuit*. Ex. 1005, [0071]-[0072], [0123].

68. It would have been obvious to implement Hartsell’s transport processing engine 1050 as a *circuit* because Hartsell explains that each of its content delivery engines, including transport processing engine 1050, “generally includes any hardware...capable of **performing one or more dedicated tasks** or subtasks[.]” Ex. 1005, [0053]; *see also* Ex. 1009, 90 (Microsoft Computer Dictionary defining *circuit* as “[a] combination of electrical components interconnected to perform a particular task”). It was also well-known and common to implement such computer-centric technologies using a *circuit*. Ex. 1015, 15:56-59 (listing a “forwarding engine” as a type of “integrated circuit...which determines the destination of a packet”); *see also* Ex. 1009, 90 (Microsoft Computer Dictionary noting that “[a]t one level, a computer consists of a single circuit; at another, it consists of hundreds of interconnected circuits”). Hartsell elaborates that transport processing engine 1050 “may employ one or more independent processor modules (e.g., CPU modules)

having independent processor and memory subsystems and **suitable for performance of a given function[.]**” Ex. 1005, [0055].

69. Hartsell’s transport processing engine 1050 “performs transport and protocol processing” and “may be any hardware...subsystem suitable for TCP/UDP processing, other protocol processing, transport processing, etc.” Ex. 1005, [0071]-[0072]. Hartsell further explains that transport processing engine 1050 facilitates and engages in data communications activities, since “content may also be routed through transport processing engine 1050,” so as “to fulfill requests from one or more clients[.]” Ex. 1005, [0123]. Because data is commonly communicated within and between computer systems (e.g., delivering requested content) in the form of a signal, it would have been obvious to a POSITA that Hartsell’s transport processing engine 1050 performs *signal processing*.


70. Accordingly, Hartsell’s transport processing engine 1050, which comprises hardware configured to perform transport and protocol processing or to route TCP/UDP content data signals, renders obvious a *signal processing circuit*.

71. Hartsell also describes a “network interface processing engine 1030” that performs a variety of network protocol stack operations, such as “handling queue management, buffer management, TCP connect sequence, checksum, IP address lookup, internal load balancing, packet switching, etc.” Ex. 1005, [0063]; *see also*

Ex. 1005, [0075]-[0084]. Hartsell suggests combining the functions of network interface processing engine 1030 with those of the transport processing engine 1050. Ex.1005, [0084] (“[T]he transport engine 1050 and the network interface engine 1030 may be combined into a single engine.”). Thus, as an alternative mapping of the claim language to the prior art, the combination of network interface processing engine 1030 and transport processing engine 1050 render obvious a *signal processing circuit*[].

72. Third, Hartsell’s content delivery system comprises two transport processing engines that operate individually (*two or more separate signal processing circuits*). Hartsell’s content delivery engines utilize “processor modules 1001 and 1003” that are configured “to implement the content delivery processing engines” 1030, 1040, 1050, 1060, and 1070. Ex. 1005, [0150]-[0151]. In other words, the processor modules 1001 and 1003 are flexibly reconfigurable to perform a variety of functions, including operating as transport processing engines 1050. Figure 1C, below, shows “one possible network content delivery system hardware configuration” where “[f]our processor modules 1001 may be configured to operate as...the transport processing engine 1050,” and “[t]wo processor modules 1003 may operate as...the transport processing engine 1050.” Ex. 1005, [0020], [0143]. As with the transport processing engine 1050, Hartsell contemplates having “two or more

individual network interface processing modules 1030a and 1030b.” Ex.1005, [0105]. Hartsell further explains that the “processor modules 1001 and 1003 of FIG. 1C may be configured in alternative manners to implement the content delivery processing engines such as the network interface processing engine 1030....” Ex. 1005, [0151]. Thus, Hartsell contemplates multiple processor modules acting together as one transport processing engine, and Hartsell further contemplates that the content delivery system may include multiple transport processing engines and network processing engines.



Ex. 1005, Fig. 1C (combining Figs. 1C' and 1C'', annotated).

73. Hartsell explains that the processor modules of each transport processing engine 1050 and network interface processing engine 1030 have “independent processor and memory subsystems...suitable for performance of a given function[.]” Ex. 1005, [0055]. Thus, each processor module 1001 or 1003 has its own independent components “allowing independent operation without interference from other engines or modules.” Ex. 1005, [0055]. Further, Hartsell notes that “[a]ny two modules of a given engine may communicate with **separate**

Ex. 1003

modules of another engine or may communicate with the same module of another engine.” Ex. 1005, [0120]. Accordingly, “content delivery system 1010 is capable of providing multiple dedicated and independent processing engines...each of which is substantially self-contained.” Ex. 1005, [0053].

74. Since each processor module 1001 or 1003 operates independently, it would have been obvious to a POSITA that one transport processing engine 1050 (composed of processor modules 1001) would operate independently from another transport processor engine 1050 (composed of processor modules 1003). Hartsell’s two transport processing engines 1050, which operate independently, therefore render obvious *two or more separate signal processing circuits*. As previously noted, Hartsell contemplates that transport processing engine 1050 may be integrated with the functionality of the network interface processing engine 1030. Ex.1005, [0084].

75. In summary, Hartsell’s teaching of content delivery system 1010 coupled to a network (*communications infrastructure*) including two transport processing engines 1050 (alone or together with network interface processing engines 1030) that each employ respective processor modules to allow independent operation from other system engines (*comprising two or more separate signal processing circuits*) renders obvious this limitation.

- b. [27.2] *each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router,*

76. As discussed in [27.1], Hartsell's two transport processing engines 1050 (*signal processing circuits*) each include a plurality of processor modules 1001 and 1003. These plurality of processor modules render obvious *multiple ASIC [(Application-Specific Integrated Circuit)] devices* because, as detailed further below, each processor module is a device that includes an FPGA (*ASIC device*).


77. **First**, Hartsell's processor modules¹ 1001 and 1003 are *ASIC devices* because they are application-specific devices that include an FPGA. The '323 Patent explains that it relates to "interconnection of multiple ASIC devices, for example, multiple Field Programmable Gate Array (FPGA) devices." Ex. 1001, 1:14-17. While the '323 Patent does not provide an explicit definition of *ASIC device*, it

¹ Hartsell uses "processor modules" to refer to modules 1001 and 1003, and "processing modules" to refer to implementations of the processor modules as, e.g., modules 1030a and 1030b, 1040a and 1040b, and 1050a and 1050b. For consistency, this Petition solely uses the term "processor modules" outside of directly quoted language.


consistently equates ASIC devices and FPGA devices. *E.g.*, Ex. 1001, 5:32-36, 10:3-6, 16:26-29, 19:44-50, 21:6-12, 23:15-24, 25:26-32.

78. Hartsell's processor modules 1001 and 1003, as shown in Fig. 1, comprise (1) a field-programmable gate array ("FPGA"), connected to (2) a Pentium Chipset/Pentium III (or a PowerPC processor), and (3) memory.² Figure 1C, below, provides examples of processor module 1001 and processor module 1003:

² Hartsell describes processor module 1003 as also including various fibre channel components that are used "when the processor module 1003 is utilized as a storage processing engine 1040." Ex. 1005, [0149]. Since these fibre channel components are not present in processor module 1001, it would have been obvious to a POSITA that the fibre channel components are not needed as part of processor module 1003 when processor module 1003 is used as a transport processing engine 1050.



Ex. 1005, Fig. 1C (partial view, annotated).



Ex. 1005, Fig. 1C (partial view, annotated).

79. Hartsell discloses that each of the processor modules may use “any processor suitable for adapting to endpoint processing...such as...the PowerPC processor.” Ex. 1005, [0055], [0072] (“In one embodiment transport engine 1050


Ex. 1003

may be a dedicated TCP/UDP processing module based on an INTEL PENTIUM III or MOTOROLA POWERPC 7450 based processor....”). The ’323 Patent similarly contemplates embodiments where PowerPC processors are embedded on an FPGA device (*ASIC device*), along with memory and additional programmable gates that can be configured to perform different tasks such as a data interface. Ex. 1001, 1:33-50 (describing “system-level FPGAs with embedded Power PC” in the context of “current commercial products”), 10:22-35 (“FPGA devices 102, 104, 106, and 108 may be a Xilinx Virtex-II Pro® XC2VP50 FPGA device”).

80. To the extent that Patent Owner argues that the claimed *ASIC devices* need to be implemented on a single integrated chip, well-known FPGA devices at the time of the ’323 Patent also include the components found in Hartsell’s processor modules 1001 and 1003, such as memory and a processor, incorporated into an integrated solution. For instance, the ’323 Patent points to Xilinx Virtex-II Pro FPGAs as an example of an ASIC device. Ex. 1001, 10:3-6 (“[M]ultiple ASIC devices may be provided on a single circuit card, in this exemplary embodiment in the form of four FPGA devices”) 10:22-28 (“[E]ach of [the four] FPGA devices...may be a Xilinx Virtex-II Pro® XC2VP50 FPGA device”). Xilinx documentation confirms that the Virtex-II Pro FPGA family contains the same components contemplated by Hartsell for processor modules 1001 and 1003 in a


single FPGA device. *See* Ex. 1010, 1-2 (describing the technical specifications of the Xilinx Virtex-II Pro FPGA family). The Virtex-II Pro FPGA family members include “user-configurable supporting circuitry,” an “on-chip processor,” and “on-chip memory.” Ex. 1011, 17-18 (Virtex-II Pro FPGA “incorporate[s] both serial transceiver technology and a hard processor core,” e.g., a PowerPC Processor, “within a general-purpose FPGA device.”); Ex. 1010, 1 (providing technical specifications for the ten members of the Virtex-II Pro family); *see also* Ex. 1006, 1:14-31 (describing conventional FPGAs), 1:31-39 (describing Xilinx documentation relating to the Virtex-II Pro family in relation to conventional FPGAs).

81. For instance, Curd (Ex. 1006) illustrates the major components of a Virtex-II Pro FPGA device, including input/output interfaces, multiple multi-gigabit transceivers, configurable logic blocks and configuration logic (together providing the device’s field-programmable functionality), and a PowerPC processor. These components are reflected in Figure 1, below:



Ex.1006, Fig. 1 (annotated)

82. The components substantially correspond to the same components described by Hartsell as being part of its processor modules 1001 and 1003, such as, a PowerPC processor, memory, and input/output (I/O) interfaces, multi-gigabit transceivers, configurable logic blocks, and configuration logic (collectively forming an FPGA).



Ex. 1005, Fig. 1C (partial view, annotated).

83. In short, Curd confirms that the Xilinx Virtex-II Pro FPGA device is a commercially available product that includes all of the components and functionality required by Hartsell's processor modules 1001 and 1003. Each component of Hartsell's processor modules 1001 and 1003 corresponds to functionality included in a Xilinx Virtex-II Pro FPGA device. Accordingly, it would have been obvious to a POSITA that to the extent the claimed *ASIC devices* need to be implemented on a single integrated chip, Hartsell's processor modules 1001 and 1003 would have been implementable as single integrated chip devices such as the known Xilinx Virtex-II Pro FPGA devices. Further, by implementing Hartsell's processor modules as FPGA devices, the processor modules necessarily are *ASIC devices* because FPGAs are *ASIC devices*, as confirmed by the '323 Patent and '323 File History. See Ex.1001,

2:45-50 (noting that FPGAs are examples of *ASIC devices*); Ex.1002, 44-46 (applicant arguing same during prosecution).

84. A POSITA would have been motivated to implement Hartsell's processor modules as single integrated chip FPGA devices, such as a Xilinx Virtex-II Pro implementation of Hartsell's processor modules, because such an implementation would lower the manufacturing cost (e.g., using an already-designed and -tested embedded processor on an FPGA device instead of creating a custom integration of a separate CPU such as Pentium III or PowerPC and a separate FPGA device), save circuit area, and thus improve production and operation efficiency. Ex.1011, 18 (“[U]sing a separate processor externally requires hundreds of additional interface pins, which degrades system performance and significantly increases FPGA I/O requirements and overall board costs.”), 17 (“cost savings com[e] from faster time-to-market, reduced printed circuit board (PCB) complexity, and lower component count”); Ex.1008, 175 (discussing cost advantages of Xilinx Virtex-II Pro FPGAs). Additionally, use of an off-the-shelf single chip integrated circuit FPGA, such as a Xilinx Virtex-II Pro FPGA device, would result in reduced complexity and faster time-to-market by employing an off-the-shelf component instead of designing, prototyping, testing, and manufacturing custom modules as originally suggested by Hartsell. Ex.1011, 17 (discussing Virtex-II Pro FPGA technology providing costs

savings from “faster time-to-market, reduced printed circuit board (PCB) complexity, and lower component count”).

85. The use of single integrated chip FPGA devices, such as Xilinx Virtex-II Pro FPGAs, to implement Hartsell’s content delivery system 1010 is also consistent with Hartsell’s discussion of using “modular” components. Hartsell discusses “modular” multi-task configuration of its content delivery system 1010 for easy customization of the system. Ex.1005, [0045], [0055]. Hartsell teaches the use of “modular processing engines,” each “provided with one or more discrete subsystem modules.” Ex.1005, [0045]. This modular configuration “allows the number and/or type of content delivery engines and processors to be selected or varied to fit the needs of a particular application.” Ex.1005, [0055].

86. Thus, a POSITA would have been motivated to implement each “modular” processor module 1001 or 1003 of Hartsell using a single integrated chip FPGA device, such as a Xilinx Virtex-II Pro FPGA, such that each processing module is a single integrated circuit (e.g., a single semiconductor chip). *See* Ex.1006, Fig. 1. Doing so would support and further enhance the modular structure Hartsell intends to achieve. Each “modular” processing module integrated on a single FPGA device can be easily installed or replaced in the content delivery system when the “number

and/or type of content delivery engines and processors” varies to “fit the needs of a particular application.” Ex.1005, [0055].

87. In addition, a POSITA would have been motivated to implement processor module 1001 or 1003 as a single integrated chip FPGA device, such as a Xilinx Virtex-II Pro FPGA, because such FPGAs supported superior circuit performance at the time of the '323 Patent. Ex.1010, 17-18. By implementing Hartsell's processor modules 1001 and 1003 as single integrated chip FPGA devices, such as Xilinx Virtex-II Pro FPGAs, Hartsell's content delivery system 1010 would achieve high transmission rates and compatibility of a variety of I/O technologies. Ex.1010, 2 (Virtex-II Pro devices support transmission rates of “622 Mbps to 3.125 Gbps per channel”), 4 (noting the advantages of Xilinx FPGAs “for customers who still need to support their differential or single-ended legacy bus architectures as they transition to serial architectures” with compatibility and support for “all the popular standards”). Such benefits would have motivated a POSITA to implement Hartsell's processor modules 1001 and 1003 as single integrated chip devices such as the known Xilinx Virtex-II Pro FPGA devices.

88. In addition, Hartsell's processor modules 1001 and 1003 (*ASIC devices*) are each dedicated to perform specific tasks. Content delivery system 1010 is configurable such that “any given processing engine of content delivery system

1010,” such as transport processing engine 1050 or network interface processing engine 1030, is “provided with multiple processing modules so as to enable parallel or redundant processing of data and/or communications.” Ex. 1005, [0105]. These processor modules are application specific because they are dedicated and optimized for specific functions. Ex. 1005, [0107] (noting that in some embodiments, processor modules are “specialized to specific applications”). For instance, “two or more **individual dedicated TCP/UDP processing modules** 1050a and 1050b,” shown in Fig. 1C as processor modules 1001 or 1003, “may be provided for transport processing engine 1050[.]” Ex. 1005, [0105], [0053] (“content delivery system 1010 is capable of providing **multiple dedicated** and independent processing engines that are optimized for networking”).

89. Hartsell’s content delivery engines employ processor modules (each with their own processor and memory subsystems) to perform particular sub-tasks and operate independently. Ex. 1005, [0055]. For instance, Hartsell’s transport processing engine off-loads tasks from other processing engines to handle application-specific tasks and to free up resources of other processing engines:

The transport processing engine 1050 may perform time out checks for each network connection, session management, data reordering and retransmission, data queuing and flow control, packet header generation, etc.


off-loading these tasks from the application processing engine or the network interface processing engine. The transport processing engine 1050 may also handle error checking, likewise freeing up the resources of other processing engines.

Ex. 1005, [0073]; *see also* Ex. 1005, [0135] (“Each engine may also be tailored in hardware and/or software to **perform specific content delivery task[s]**, thereby providing increasing content delivery speeds while requiring less system resources.”).

90. Accordingly, Hartsell’s processor modules 1001 and 1003, configured to operate as transport processing engines 1050 (and, alternatively, further configured to operate as network interface processing engines 1030), are application-specific devices that contain FPGAs (*ASIC devices*).

91. Second, Hartsell’s processor modules (*ASIC devices*) *each itself includes a packet router* because each processor module controls packet routing.

92. Hartsell explains that “[a]ny two modules of a given engine,” such as processor modules 1003 on the left-hand side in FIG. 1C’, “may communicate with separate modules of another engine or may communicate with the same module of another engine,” such as processor modules 1001 on the right-hand side in FIG. 1C”. Ex. 1005, [0120].



Ex. 1005, Fig. 1C (combining Figs. 1C' and 1C'', annotated).

93. A POSITA would have recognized that Hartsell's processor modules 1001 and 1003 would have operated as a *packet router* in the same way as the FPGAs described in the '323 Patent because the processor modules 1001 and 1003 are used to process packets and communicate packets to other processor modules 1001 and 1003. Indeed, Hartsell explains that its transport processing engines 1050 (which are implemented via multiple processor modules 1001 or 1003) "may perform time out checks for each network connection, session management, data reordering and

Ex. 1003

retransmission, data queuing and flow control, packet header generation, etc.”

Ex. 1005, [0073]. Each processor module also includes an FPGA, and a POSITA would have configured the FPGAs of processor modules 1001 and 1003 to perform desired functionality, including performing routing functions, in light of Hartsell’s disclosure of the transfer of packets between engines in content delivery system 1010.


Ex. 1005, [0078] (discussing the transfer of packets between engines).

94. Hartsell further contemplates processing engine configurations that are “distributively interconnected across a network,” including configurations where processing engines are “physically remote” from one another. Ex. 1005, [0275]. Data is communicated between Hartsell’s transport processing engines in the form of data packets. Ex. 1005, [0072]-[0073], [0078] (discussing TCP/UDP processing and the transfer of packets “between the two engines”). “Routing of requests...may be controlled by transport and/or interface processing modules,” such as Hartsell’s processor modules 1001 and 1003. Ex. 1005, [0107]. Each processor module performs the routing of requests in parallel with one or more additional processor modules. Ex. 1005, [0105]. For example, Hartsell describes filtering requests, indicating that some processor modules decide whether to forward a request packet or to drop the request packet, for example, because the request cannot be fulfilled. Ex. 1005, [0103] (“The network interface engine may also be used to filter requests

that are not supported....”), [0100], [0169] (“to filter out requests that the receiving system is not capable of processing”), [0190] (“...the request for content may be filtered at step 110 by, for example, one or more processing engines or modules...”).

95. Hartsell’s disclosure—with processor modules that comprise network interface and transport processing engines that control the routing of packetized data requests—renders obvious a *packet router* included within each processor module (*ASIC device*). As each processor module 1001 or 1003 is capable of routing and processing content requests (e.g., packets) in parallel, it would have been obvious for each transport processing module 1001 or 1003 to include a *packet router*. See Ex. 1005, [0123] (discussing routing content through transport processing engine 1050). For example, it would have been obvious for the processor (e.g., PowerPC) of each processor module 1001 or 1003 to act as a *packet router* because the processor module 1001 or 1003 operates to send packets to other processor modules 1001 or 1003, or even to a different system altogether. Hartsell illustrates in Figs. 1G and 1H, for example, that multiple systems 1010 can be coupled together and to a network 1020 via network interface processing engines 1030, which are implemented via processor modules 1001 or 1003. Ex. 1005, [0151]. Hartsell describes this as a “cluster arrangement” of systems 1010. Ex. 1005, [0112]. Hartsell also contemplates

that the network interface processing engine 1030 may be combined with the transport processing engine 1050. Ex. 1005, [0084].



Hartsell, Fig. 1H.

96. Alternatively, the incorporation of the network interface processing engine 1030 into the transport processing engine 1050 further renders obvious the claimed *packet router*. Ex. 1005, [0084]. Hartsell's network interface processing engine 1030 interfaces with a network "by receiving and processing requests for content and delivering requested content" to the network. Ex. 1005, [0063]. Network interface processing engine 1030 "may be suitable for handling queue management, buffer management, TCP connect sequence, checksum, IP address lookup, internal load balancing, packet switching, etc." Ex. 1005, [0063]. Network interface

processing engine 1030 utilizes a processor which “receives a packet from a port, verifies fields in the packet header, and decides on an outgoing port to which it forwards the packet.” Ex. 1005, [0066]. The “network interface engine 1030 may process[] all or some of the TCP/IP protocol stack.” Ex.1005, [0075]. The network interface engine 1030 receives TCP/IP packets and strips their header information, replacing it with an identifier or tag that is used to route the packet within Hartsell’s system. Ex.1005, [0066], [0080], [0126].

97. The ’323 Patent’s disclosure confirms that Hartsell’s processor module 1001 or 1003 would operate as a *packet router*. The ’323 Patent explains that any routing scheme, including TCP/IP and Infiniband, would involve a *packet router*. Ex. 1001, 22:49-23:8 (“[I]t will be understood that any other data packet configuration and routing scheme suitable for allowing a sender or source to determine the packet's destination within a reconfigurable communications infrastructure 1300 may be employed (e.g., TCP/IP, Fibrechannel, XAUI, Ethernet, Infiniband, Rapid I/O, etc.)”). Similarly, Hartsell teaches communications among the processor modules 1001 and 1003 may use of any interconnect protocol, such as Infiniband. Ex. 1005, [0275]. A POSITA would have understood that the Infiniband protocol transmits data in the form of packets. *See* Ex. 1013, Abstract (discussing

switching Infiniband data packets); Ex. 1014, Abstract (discussing converting ATM cells to Infiniband packets).

98. Thus, the combination would result in processor modules 1001 or 1003 performing routing of Infiniband packets, which also renders obvious that each processor module 1001 or 1003 *includes a packet router*.


99. Therefore, Hartsell renders obvious this limitation.

- c.** *[27.3] said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits being coupled through respective first and second common interfaces and an intervening high speed serial optical link to a respective packet router of each of the ASIC devices of each other of said two or more signal processing circuits*

100. First, as discussed in [27.1]-[27.2], Hartsell teaches that one transport processing module (such as processor module 1001 or 1003) in a transport processing engine is **communicatively coupled** to another transport processing module in another transport processing engine. Ex. 1005, [0143]. A POSITA would have understood that these processor modules are *ASIC devices*. And each processor module includes a *packet router* for routing requests and data between the transport processing engines. *See supra*, [27.2].

101. Second, Hartsell renders obvious that the processor modules (each of which comprises a *packet router*) are coupled via an *intervening high speed serial optical link*, as further detailed below.

102. Hartsell's processor modules are coupled to other processor modules by a node-to-node link in interconnect 1080. Ex. 1005, [0058] ("distributive interconnection 1080 may be any...high speed interchange system[] such as a switch fabric"), [0144] ("the distributive interconnect 1080 may be a switch fabric based interconnect"). As shown in Fig. 1C, below, interconnect 1080 connects transport processing engines 1050 and their processor modules 1001 and 1003. *See* Ex. 1005, [0058] ("a switch fabric establishes independent paths from node to node"), [0144] ("the switch fabric 1081 may be coupled to the individual processor modules through interface converter circuits 1082").



Ex. 1005, Fig. 1C (combining Figs. 1C' and 1C'', annotated).

Interconnect 1080 includes node-to-node path links that connect one processing module to another processing module and transmit data bit-by-bit serially over “data aligned serial link interfaces.” Ex. 1005, [0143]-[0144]. This node-to-node path link is a *high speed serial* link because interconnect 1080 is an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode. Ex. 1005, [0144].

103. Further, it would have been obvious to a POSITA to use a *serial optical link* as the node-to-node path link connecting one processor module to another. Hartsell teaches that its processing engines can be interconnected via “optical interconnect technology,” such as Infiniband. Ex. 1005, [0275]; *see also supra*, [27.2]. Connecting Hartsell’s processor modules via “optical interconnect technology” renders obvious the claimed *intervening high speed serial optical link*.

104. Alternatively, to the extent the claimed *intervening high speed serial optical link* requires a single bidirectional link as described in embodiments of the ’323 Patent, it would have been obvious to implement interconnect 1080 as a bus. A POSITA would have been motivated to connect the processing engines in Hartsell’s content delivery system 1010—and thus each of Hartsell’s processor modules 1001 and 1003—using a single bidirectional bus. Hartsell encourages such an implementation because it teaches that “distributive interconnection 1080 may be any multi-node I/O interconnection hardware or hardware/software system” including a “bus architecture.” Ex.1005, [0058]. Hartsell explains that “a ‘bus’ is any system bus that carries data in a manner that is visible to all nodes on the bus.” Ex.1005, [0058]. Thus, in a bus architecture, a single link would be used for communications between any two processor modules. Skilled artisans were using bus architectures to facilitate communications between and among FPGAs prior to the priority date of the ’323

Patent to couple FPGAs together. Ex.1030, 628 (describing a “linear array of FPGA’s [] augmented with a global bus”).

105. A POSITA would have coupled each of Hartsell’s processor modules 1001 and 1003 (*ASIC devices* containing FPGAs) using a bus because Hartsell provides a teaching, suggestion, or motivation to implement interconnect 1080 as a bus architecture. Ex.1005, [0058] (teaching that processing engines are interconnected by a “bus architecture”). Buses provide a single bidirectional link, such as is described in embodiments of the ’323 Patent. Ex.1029, 1:54-2:4 (providing “an advanced bus structure” for bidirectional optical routing of information); *see also* Ex.1001, Fig. 13, 19:24-28 (describing a “high bandwidth interconnection medium 1350 to enable high bandwidth bi-directional communications”). A POSITA would have recognized that a bus architecture would have been capable of interconnecting Hartsell’s processor modules using “optical interconnect technology.” *See* Ex.1005, [0275]; Ex.1029, 2:33-41 (describing an “optical bus interface module” to facilitate device communication over fiber optic transmission lines). Bus architectures present advantages because they are easier to deploy and result in cost savings over alternative architectures. And Hartsell confirms that a bus architecture is a “high speed interchange system[],” thus maintaining the high speed advantages of Hartsell’s switch fabric implementation discussed directly above. Ex.1005, [0058].

106. A POSITA would have had a reasonable expectation of success in implementing Hartsell's distributive interconnect 1080 as a bus architecture because, as explained above, such architectures were known in the art, and Hartsell intended for its content delivery systems to be implementable using bus architectures. Ex.1005, [0058]. Accordingly, it would have been obvious to a POSITA to implement content delivery system 1010 such that each processing engine and corresponding processor modules 1001 and 1003 were interconnected using a single bidirectional bus.

107. Third, Hartsell renders obvious that each of its processor modules (each comprising a *packet router*) would be coupled via an *intervening high speed serial optical link* to another processor module through fabric interfaces (*first and second common interfaces*).

108. Hartsell describes an interface 34 located between interconnect 1080 and each processor module 1001 and 1003. Ex. 1005, [0144] ("interface 34 provides an interface between the processor module bus and the UDASL switch interface converter circuit 1082 [within interconnect 1080]"), Fig.4. Hartsell's interface 34 for each processor module 1001 or 1003 renders obvious the *first and second common interfaces*. Alternatively, Hartsell's UDASL switch interface converter circuit 1082 may be considered part of each processor module's *common interface* rather than part of the *high speed serial... link*. In view of Hartsell's suggestion to modify

interconnect 1080 (in either a switch fabric or bus implementation) to use optical communication technology, it would have been obvious to incorporate suitable optical transceiver functionality into interface 34, UDASL switch interface converter circuit 1082, or both.

109. Therefore, Hartsell renders obvious this limitation.

- d. [27.4] *with no other processing device intervening between the high speed optical link and said ASIC devices of each of said two or more signal processing circuits.***

110. As discussed in [27.3], Hartsell renders obvious each processor module (*ASIC device*) connected to interconnection 1080 (*intervening high speed serial optical link*) via interface 34 (*common interface*), *with no other processing device intervening between.*

111. Hartsell teaches the processing engines “communicate with each other directly as peers.” Ex. 1005, [0125]. Further, Hartsell depicts no other processing devices between the processor modules 1001 and 1003 and the interconnect 1080 in its switch fabric implementation. *See* Ex. 1005, Figs. 1C', 1C"; *see also* Ex. 1005, [0058] (“a switch fabric establishes **independent paths from node to node** and data

is specifically addressed to a particular node on the switch fabric”).³ Similarly, no other processing devices would intervene in a bus implementation of interconnect 1080 because in such implementations, each processing engine or processor module 1001 and 1003 of Hartsell’s content delivery system 1010 would be coupled to the bus via interfaces with no other processing devices between the respective common interfaces and the bus.⁴ Accordingly, in the architecture suggested by Hartsell, the processor modules are interconnected by optical communication links without any intermediary processing device.

112. Therefore, Hartsell renders obvious this limitation.

³ I have provided this analysis consistent with the position that the claim’s “*no other processing device intervening...*” limitation does not foreclose the presence of processing devices within the *high speed serial optical link*.

⁴ I have provided this analysis consistent with how Figures 13, 14, and 15 of the ’323 Patent depict FPGA devices and signal processing circuits as being interconnected. Ex.1001, Figs. 13-15.

2. Claim 28

- a. [28.1] *The communications infrastructure of claim 27, wherein said interconnection further comprises a high bandwidth interconnection medium.***

113. Hartsell renders obvious that its *intervening high speed serial optical link* would comprise a *high bandwidth interconnection medium* (to the extent there are differences between the two terms). Claim 28 recites “said interconnection,” but claim 27 does not recite an “interconnection.” *See also* Ex. 1002, 41 (amending claim 27 to remove “an interconnection comprising a high speed serial optical link” and to add “an intervening high speed serial optical link”). The following analysis interprets the recited “interconnection” to refer to the “intervening high speed serial optical link” from claim 27.

114. As noted in limitation [27.3], Hartsell’s interconnect 1080 includes node-to-node path links that connect one processing module to another processing module and transmit data bit-by-bit serially over “data aligned serial link interfaces.” Ex. 1005, [0143]-[0144]. Hartsell explains that interconnect 1080 is an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode. Ex. 1005, [0144]. A POSITA would have understood that Hartsell’s interconnect 1080 is a *high bandwidth interconnection medium* due to data transmission capabilities of at least 1.77 Gbps. Ex.1005, [0144]. This is confirmed by the ’323 Patent, which explains that “[e]xamples of suitable high bandwidth interconnection

Ex. 1003

mediums include...any interconnection mediums...having a data transmission capability of greater than or equal to about 1 gigabit per second (Gbps).” Ex.1001, 6:2-7, 13:25-28 (describing an embodiment supporting “data transmission rates between 622 Mbps and 3.125 Gbps”).

115. Because Hartsell describes its interconnect 1080 as an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode—each supporting a data transmission capability greater than 1 Gbps—Hartsell renders obvious this limitation. Ex. 1005, [0144]; Ex. 1001, 6:2-7.

3. Claim 31

a. [31.1] *A method, comprising providing two or more separate signal processing circuits,*

116. As discussed in [27.1], Hartsell provides a communications infrastructure comprising *two or more separate signal processing circuits*. See *supra*, [27.1]. And as discussed in further detail below in [31.5], Hartsell renders obvious a method involving transferring data packets between the signal processing circuits. See *infra*, [31.5]. Accordingly, Hartsell renders obvious this limitation.

b. [31.2] *each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router,*

117. See analysis at [27.2].

- c. ***[31.3] said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits being coupled through respective first and second common interfaces and an intervening high speed serial optical link to a respective packet router of each of the ASIC devices of each other of said two or more signal processing circuits***

118. See analysis at [27.3].

- d. ***[31.4] with no other processing device intervening between the high speed optical link and said ASIC devices of each of said two or more signal processing circuits;***

119. See analysis at [27.4].

- e. ***[31.5] selectively transferring at least one data packet from each said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits to each respective packet router of said at ASIC devices of each of the other of said two or more signal processing circuits through said first and second common interfaces and said intervening high speed serial optical link without routing said data through any other intervening processing device between each respective ASIC device and the high speed serial optical link.***

120. As discussed in [27.3], the prior art renders obvious that data may be transmitted between one *packet router* in a processor module 1001 or 1003 (*ASIC device*) of a transport processing engine 1050 (*signal processing circuit*) to another processor module 1001 or 1003. And as discussed in [27.2], a POSITA would have found it obvious for each of Hartsell's processor modules 1001 or 1003 to operate as

(or to include) a *packet router* because in the combination, each transport processing module would route request packets to their intended destinations over peer-to-peer optical links available through interconnect 1080. Ex. 1005, [0045] (“high speed distributive interconnect... that allows peer-to-peer communication between individual subsystem modules”), [0107] (“Routing of requests... may be controlled by transport and/or interface processing modules”), [0239]. Further, as discussed in [27.3], Hartsell’s processor modules 1001 and 1003, configured as transport processing engines 1050, are coupled through a first interface 34 (*first common interface*), interconnection 1080 (*high bandwidth interconnection medium*), and a second interface 34 (*second common interface*). See Ex. 1005, [0144], Figs. 1C, 4.

121. Further, Hartsell teaches *selectively transferring at least one data packet* by disclosing the data communicated is in the form of data packets. See *supra*, [27.2]. Hartsell discloses that transport processing engines 1050 (*signal processing circuits*) have routing functions to route data between processor modules. “Routing of requests...may be accomplished using a distributive interconnect and may be controlled by transport and/or interface processing modules.” Ex. 1005, [0107]. It would have been obvious to a POSITA that routing requests among processor modules includes *selectively transferring at least one data packet*. Hartsell further explains that “information management flow may advantageously occur between any

two or more processing engines that may be present on a network and interconnected via a virtual distributed inter connect backplane.” Ex. 1005, [0276].

122. Hartsell’s transport processing engine 1050 performs “network transport protocol sub-tasks, such as processing content requests received from network interface engine 1030.” Ex. 1005, [0072]. A POSITA would have understood that the received requests would include packets, because network interface engine 1030 “process[es] high speed streams of network packets,” Ex. 1005, [0066], requests are passed to network transport processing engine 1050, and transport processing engine 1050 performs “packet header generation” as part of its tasks. Ex. 1005, [0072]-[0073]; *see also* Ex. 1005, [0127] (“each processing engine” manipulates data packet headers by, e.g., setting priority levels), [0138] (processing engines reduce “the packet size transmitted across the distributed interconnect” through protocol layer manipulation).

123. Further, Hartsell discusses using Infiniband, as noted above. *See supra*, [27.2]-[27.3]; Ex. 1005, [0231], [0275]. A POSITA would have understood that the Infiniband protocol transmits data in the form of packets. *See* Ex. 1013, Abstract (discussing switching Infiniband data packets); Ex. 1014, Abstract (discussing converting ATM cells to Infiniband packets). Accordingly, a POSITA would have

found it obvious for Hartsell's transport processing engine 1050 to *selectively transfer[] at least one data packet.*

124. Therefore, Hartsell renders obvious this limitation.

4. Claim 33

- a. **[33.1] *The method of claim 31, wherein said interconnection further comprises a high band width interconnection medium.***

125. See analysis at [28.1]. Like claim 28, claim 33 recites “said interconnection,” though claim 31 does not recite an “interconnection.” The analysis presented above at [28.1], which interprets the recited “interconnection” to refer to the “intervening high speed serial optical link,” similarly applies to claim 33.

**X. GROUND 2: HARTSELL AND HAUCK RENDERS OBVIOUS
CLAIMS 27-28, 31, AND 33**

A. Reasons to Combine Hartsell and Hauck

126. To the extent that it is argued that the claimed “*packet router*” within each ASIC device implies that the ASIC devices have point-to-point interconnects to other ASIC devices as illustrated in the ’323 patent’s Fig. 2 (see Ex.1001, 11:56-66), it would have been obvious to a POSITA to modify Hartsell’s architecture to provide such direct, point-to-point interconnects between the FPGAs on each of Hartsell’s processing modules 1001 and 1003. Such an architectural change would have eliminated the need for Hartsell’s switch fabric 1080 by providing more direct communications paths between pairs of processing modules 1001 and 1003. Many such suitable multi-FPGA interconnection schemes were known in the art. For example, it would have been obvious to a POSITA to try a “mesh” or “crossbar” interconnection architecture among the FPGAs of the processing modules because those were the “most common topologies” for interconnecting multiple FPGAs in a multi-FPGA system. Ex.1030, 626. As Hauck notes, many two-dimensional and three-dimensional mesh-based systems had been built as early as 1998. Ex.1030, 626. Techniques for routing in such mesh-based multi-FPGA systems had also been long known, as evidenced by Hauck’s bibliography. E.g., Ex.1030, 636 at [57] (“S.

Hauck, G. Borriello, and C. Ebeling, “Mesh routing topologies for multi-FPGA systems,” in *Proc. Int. Conf. Computer Design*, Oct. 1994, pp. 170-177”).

127. Additional reasons to combine teachings from each of the references are provided in the analysis below.

B. Detailed Analysis of Claims

128. The claim analysis for this ground relies on the same analysis and evidence as provided for Ground 1, except for where additional analysis is provided, as noted for respective limitations of the claims below.

1. Claim 27


- a.** [27.1] *A communications infrastructure, comprising two or more separate signal processing circuits,*

129. Ground 2 relies on the same analysis and evidence as provided above in Ground 1 for limitation [27.1]. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

- b.** [27.2] *each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router,*

130. Ground 2 relies on the same analysis and evidence as provided above in Ground 1 for limitation [27.2] regarding Hartsell’s disclosure of *ASIC devices*. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

131. Additionally, it would have been obvious to a POSITA to implement Hartsell's content delivery system 1010 such that the processor modules 1001 and 1003 (*ASIC devices*) would have each *include[d] a packet router* allowing for direct connection between each processor module, where such implementation would have interconnected each processor module 1001 and 1003 in the same way as depicted in the figures of the '323 Patent. Such interconnections between the FPGAs of processor modules 1001 and 1003 would have been obvious and well-known to a POSITA by the priority date of the '323 Patent, as evidenced by Hauck. Hauck confirms the practice of directly connecting FPGA devices to route packets between the FPGAs was used "for both routing and rerouting of signals[.]" Ex.1030, 632, 633 (describing a system where FPGAs "are used for both routing and rerouting of signals"), Fig. 25(a).



Ex. 1030, Fig. 25(a).

132. Like Hartsell, Hauck describes the connection of devices, such as a processor, memory, or other necessary elements to implement a system, connected to the FPGAs depicted in Hauck's Fig. 25(a), above, just as Hartsell's processor modules 1001 and 1003 include both an FPGA as well as elements including a processor and memory. Ex.1030, 632-633. Accordingly, it would have been obvious to modify Hartsell's content delivery system 1010 to directly interconnect Hartsell's processor modules 1001 and 1003, rather than having them connect through the switch fabric, because doing so would have been no more than applying a known technique (directly connecting FPGAs) to a known device (Hartsell's content delivery system 1010) ready for improvement to yield predictable results (directly

connecting the FPGAs of each of Hartsell's processor modules 1001 and 1003 to route data between the processor modules). A POSITA would have had a reasonable expectation of success in making such modifications to Hartsell because, as noted above, the practice of directly connecting FPGAs such that FPGAs serve as packet routers to directly route data to other FPGAs was well-known in the art prior to the priority date of the '323 Patent.

133. And to the extent that each of Hartsell's processor modules 1001 and 1003 are implemented as single integrated chip FPGA devices, a POSITA would have recognized that each of the FPGA devices would be directly interconnected. As an example, Xilinx Virtex-II Pro FPGAs, described above, were well-known to have port modules capable of routing communications to other FPGAs. *See* Ex.1032, 1:40-43 (Xilinx patent FPGA with "port modules having configurable logic configured to process communications for routing the communications"); Ex.1033, Abstract, [0004] (explaining that FPGAs are designed "with multi-gigabit serial transceivers that support a communications protocol...to enable communications of data between FPGAs at high speed").

134. Thus, the modification would result in processor modules 1001 or 1003 performing routing of packets directly to other processor modules through the

respective FPGAs, which renders obvious that each processor module 1001 or 1003 *includes a packet router*.

- c. [27.3] *said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits being coupled through respective first and second common interfaces and an intervening high speed serial optical link to a respective packet router of each of the ASIC devices of each other of said two or more signal processing circuits*

135. **First**, as discussed in [27.1]-[27.2], Hartsell teaches that one transport processing module (such as processor module 1001 or 1003) in a transport processing engine is **communicatively coupled** to another transport processing module in another transport processing engine. Ex. 1005, [0143]. A POSITA would have understood that these processor modules are *ASIC devices*. And each processor module includes a *packet router* for routing requests and data between the transport processing engines. *See supra*, [27.2].

136. **Second**, Hartsell renders obvious that the processor modules (each of which comprises a *packet router*) are coupled via an *intervening high speed serial optical link*, as noted above in [27.2] and further detailed below.

137. In the modification described above in [27.2], Hartsell's processor modules are directly coupled to other processor modules via the FPGA in each processor module. *See supra*, [27.2]. The practice of directly connecting FPGA devices to route packets between the FPGAs was used "for both routing and rerouting

of signals[.]” Ex.1030, 632, 633 (describing a system where FPGAs “are used for both routing and rerouting of signals”), Fig. 25(a). This was also well-known to be implemented with single integrated chip FPGA devices, such as the Xilinx Virtex-II Pro FPGA family. *See* Ex.1020, 1:40-43 (Xilinx patent FPGA with “port modules having configurable logic configured to process communications for routing the communications”); Ex.1021, Abstract, [0004] (explaining that FPGAs are designed “with multi-gigabit serial transceivers that support a communications protocol...to enable communications of data between FPGAs at high speed”).

138. This node-to-node path link is a *high speed serial* link because interconnect 1080 is an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode. Ex. 1005, [0144]. In modifying Hartsell to use direct, FPGA-to-FPGA interconnects as suggested by Hauck, it would have been obvious to a POSITA to maintain the speed and bandwidth of Hartsell’s original communications capabilities, and thus to use interconnection technologies capable of communicating signals at least as fast as in Hartsell’s original disclosure in order to maintain Hartsell’s original capabilities.

139. Further, it would have been obvious to a POSITA to use a *serial optical link* as the node-to-node path link connecting one processor module to another. Hartsell teaches that its processing engines can be interconnected via “optical

interconnect technology,” such as Infiniband. Ex. 1005, [0275]; *see also supra*, [27.2]. Specifically, because Hartsell discloses that its processor modules use optical interconnect technology and optical protocols for packet transfer, in the modification to Hartsell, each processor module 1001 and 1003 would be connected to other processor modules 1001 and 1003 using optical interconnect technology. Connecting Hartsell’s processor modules via “optical interconnect technology” renders obvious the claimed *intervening high speed serial optical link*.

140. Third, in the modification, Hartsell renders obvious that each of its processor modules (each comprising a *packet router*) would be coupled via an *intervening high speed serial optical link* to another processor module through *first and second common interfaces* connected to the respective FPGAs of each processor module.

141. Hartsell describes an interface 34 located between interconnect 1080 and each processor module 1001 and 1003. Ex. 1005, [0144] (“interface 34 provides an interface between the processor module bus and the UDASL switch interface converter circuit 1082 [within interconnect 1080]”), Fig.4. A POSITA would have recognized that by modifying Hartsell to directly connect processor modules 1001 and 1003 (as suggested by Hauck), Hartsell’s interface 34 for each processor module 1001 or 1003 would instead be used as *first and second common interfaces* to

incorporate suitable optical transceiver functionality to transfer packets via an optical link between the respective processor modules.

142. Therefore, Hartsell renders obvious this limitation.

- d. [27.4] *with no other processing device intervening between the high speed optical link and said ASIC devices of each of said two or more signal processing circuits.***

143. Ground 2 relies on the same analysis and evidence as provided above in Ground 1 for limitation [27.4]. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

2. Claim 28

- a. [28.1] *The communications infrastructure of claim 27, wherein said interconnection further comprises a high bandwidth interconnection medium.***

144. Ground 2 relies on the same analysis and evidence as provided above in Ground 1 for limitation [28.1]. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

3. Claim 31

- a. [31.1]-[31.4] (elements recited in claim 31)**

145. As noted above in Ground 1, the analysis for elements [31.1]-[31.4] of claim 31 are substantially similar to elements [27.1]-[27.4] of claim 27. Thus, these elements are obvious for the same reasons discussed above for claim 27 in Ground 2.

- b. ***[31.5] selectively transferring at least one data packet from each said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits to each respective packet router of said at ASIC devices of each of the other of said two or more signal processing circuits through said first and second common interfaces and said intervening high speed serial optical link without routing said data through any other intervening processing device between each respective ASIC device and the high speed serial optical link.***

146. Ground 2 relies on the same analysis and evidence as provided above in Ground 1 for limitation [27.1]. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

4. Claim 33

- a. ***[33.1] The method of claim 31, wherein said interconnection further comprises a high band width interconnection medium.***

147. Ground 2 relies on the same analysis and evidence as provided above in Ground 1 for limitation [33.1]. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

XI. GROUNDS 3 AND 4: THE COMBINATION OF HARTSELL AND RAJ AND THE COMBINATION OF HARTSELL, HAUCK, AND RAJ RENDER OBVIOUS CLAIMS 27-28, 31, AND 33

A. Reasons to Combine Hartsell and Raj

148. Reasons to combine teachings from each of the references are provided in the analysis below.

B. Detailed Analysis of Claims

149. The claim analysis for Grounds 3 and 4 rely on the same analysis and evidence as provided for Ground 1 (Hartsell without Hauck) and Ground 2 (Hartsell with Hauck), respectively, expect for where additional analysis is provided, as noted for respective limitations of the claims below. The additional analysis provided below applies for both Grounds 3 and 4.

1. Claim 27

- a.** [27.1] *A communications infrastructure, comprising two or more separate signal processing circuits,*

150. Ground 3 relies on the same analysis and evidence as provided above in Ground 1 for limitation [27.1]. The arguments and evidence discussed above with respect to Ground 1 are incorporated herein.

- b.** [27.2] *each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router,*

151. Ground 3 relies on the same analysis and evidence as provided above in Grounds 1 or 2 for limitation [27.2]. The arguments and evidence discussed above with respect to Grounds 1 and 2 are incorporated herein.

- c. [27.3] *said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits being coupled through respective first and second common interfaces and an intervening high speed serial optical link to a respective packet router of each of the ASIC devices of each other of said two or more signal processing circuits*


152. **First**, as discussed in [27.1]-[27.2], Hartsell teaches that one transport processing module (such as processor module 1001 or 1003) in a transport processing engine is **communicatively coupled** to another transport processing module in another transport processing engine (which may incorporate the functionality of network interface processing engines). Ex. 1005, [0143]. A POSITA would have understood that these processor modules are *ASIC devices*. And each processor module includes a *packet router* for routing requests and data between the transport processing engines. *See supra*, [27.2].

153. **Second**, Hartsell in view of Raj renders obvious that the processor modules (each of which comprises a *packet router*) are coupled via an *intervening high speed serial optical link*, as further detailed below.

154. Hartsell's processor modules are coupled to other processor modules by a node-to-node link in interconnect 1080. Ex. 1005, [0058] ("distributive interconnection 1080 may be any...high speed interchange system[] such as a switch fabric"), [0144] ("the distributive interconnect 1080 may be a switch fabric based interconnect"). As shown in Fig. 1C, below, interconnect 1080 connects transport

Ex. 1003

processing engines 1050 and their processor modules 1001 and 1003. *See* Ex. 1005, [0058] (“a switch fabric establishes independent paths from node to node”), [0144] (“the switch fabric 1081 may be coupled to the individual processor modules through interface converter circuits 1082”).




Ex. 1005, Fig. 1C (combining Figs. 1C' and 1C'', annotated).

Interconnect 1080 includes node-to-node path links that connect one processing module to another processing module and transmit data bit-by-bit serially over “data aligned serial link interfaces.” Ex. 1005, [0143]-[0144]. This node-to-node path link is a *high speed serial* link because interconnect 1080 is an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode. Ex. 1005, [0144].

155. Further, it would have been obvious to a POSITA to use a *serial optical link* as the node-to-node path link connecting one processor module to another. Hartsell teaches that its processing engines can be interconnected via “optical interconnect technology,” such as Infiniband. Ex. 1005, [0275]; *see also supra*, [27.2]. To the extent that Hartsell’s suggestion to use Infiniband and optical interconnect technology is insufficient by itself, a POSITA would have found it obvious to consult other references for more information about using optical technologies such as Infiniband to provide optical interconnections between the processing modules. Raj provides such information, teaching that multiple circuit units can be interconnected via a switch built with *high speed serial optical links*. Ex. 1007, 1:55-67.

156. Raj teaches a “server or other processor-based device” that includes input/output controllers coupled to one another via fabric services. Ex. 1007, 1:50-55. Raj is described in the context of Infiniband, but its teachings are “applicable to

optical links in general.” Ex. 1007, 1:63-67. Raj’s Fig. 1, below, provides a depiction of the device’s fabric services, while Fig. 2 provides further detail of the configuration of the optical links:



Ex. 1007, Fig. 1 (partial view, annotated), Fig. 2 (annotated).

Raj’s device-to-device optical serial link, depicted in Figs. 1 and 2 above as links 122 and switch 124, has an optical interface to communicate optical signals with electrical

unit 132. Ex. 1007, Abstract, 2:24-31. Alternatively, to the extent that Patent Owner argues that the claimed “*optical link*” is limited to a fiber optic cable, it would have been obvious to incorporate Raj’s electrical unit 132 and optical interface 134 into Hartsell’s processing modules as part of each module’s interface 34 (“*common interface*”).

157. It would have been obvious to a POSITA to combine concepts from Hartsell’s interconnection 1080 with Raj’s optical switch 124 to interconnect multiple processor modules via optical links. A POSITA would have been motivated to combine features from Hartsell and Raj because Hartsell provides an explicit teaching to implement an optical serial link connection. Ex. 1005, [0275] (separate processing engines may be interconnected using technologies such as “**optical interconnect technology** (e.g., in conjunction with optic/optic interface-based systems)”). Raj teaches the implementation of optical interconnections using optical serial links. Ex. 1007, 1:64-65. Therefore, using Raj’s optical serial link connection between processor modules would facilitate Hartsell’s content delivery system to be compatible with an optical network environment, as suggested by Hartsell. Ex. 1005, [0275]. The combination of Hartsell and Raj would achieve an optical interconnection and configuration for Hartsell’s content delivery system, which would improve system data rates.

158. Further, a POSITA would have had a reasonable expectation of success in combining Raj's optical interconnection technology with Hartsell's interconnect 1080 because Hartsell teaches that its interconnect 1080 can be "any multi-node I/O interconnection hardware or hardware/software system" such as "a switch fabric or bus architecture." Ex. 1005, [0058]. Raj's optical switch 124 is a multi-node I/O interconnection hardware as the switch 124 connects multiple I/O controllers 130. Ex. 1007, 1:56-59. In addition, both Hartsell and Raj teach their interconnect technology can operate under the Infiniband specification, which was a well-known serial optical interconnection standard whose use was admitted to be within the POSITA skillset. Ex. 1005, [0231], [0275]; Ex. 1007, 1:65-67; Ex. 1001, 11:2-9, 14:41-48, 21:12-16, 22:63-23:1. Therefore, it would have been obvious to implement Hartsell's interconnect 1080 using concepts from Raj's optical switch 124 and links 122, and a POSITA would have had a reasonable expectation of success in doing so.

159. In the resulting combination, each pair of Hartsell's processor modules would be connected via an optical serial link as taught by Raj, which renders obvious the claimed *intervening high speed serial optical link*.

160. Third, Hartsell in view of Raj renders obvious that each of Hartsell's processor modules (each comprising a *packet router*) would be coupled via an

intervening high speed serial optical link to another processor module through fabric interfaces (*first and second common interfaces*).

161. Hartsell describes an interface 34 located between interconnect 1080 and each processor module 1001 and 1003. Ex. 1005, [0144] (“interface 34 provides an interface between the processor module bus and the UDASL switch interface converter circuit 1082 [within interconnect 1080]”), Fig.4. However, Hartsell provides limited details for how interface 34 would operate when the interconnect 1080 uses optical communication technology such as Infiniband. Raj teaches an electro-optical interface that converts optical signals from the optical switch fabric to electrical signals. Ex. 1007, 2:24-31. It would have been obvious to a POSITA to implement Hartsell’s interface 34 using optical communication components, such an electro-optical interface like that discussed in Raj.

162. A POSITA would have been motivated to combine the teachings because Hartsell describes using optical technology, and when implementing optical technology, an optical-to-electrical (and vice versa) interface would have been needed between the optical serial links and an electrical circuit unit such as the processor modules. Raj’s optical switch 124 has a similar need, as discussed immediately above, and Raj describes its implementation. By combining Raj’s optical-to-electrical interfaces (e.g., electrical interface 140 and laser driver 138) with

Hartsell's fabric interface 34, the combination would facilitate data transmission between Hartsell's processor modules over an optical fabric (or optical bus interconnecting Hartsell's processor modules, as discussed in further detail above in Ground 1). A POSITA would have had a reasonable expectation of success in combining Hartsell's interface 34 with Raj's optical-to-electrical interfaces, because Hartsell itself suggests such a combination (Ex. 1005, [0113]) and because such technology was well-known and standardized. *See* Ex. 1007, 1:5-26 (describing known electro-optical interfaces, such as an Infiniband host channel adapter).

163. In the resulting combination, Hartsell's interface 34 for each processor module 1001 or 1003, implementing Raj's electrical interface 140 and laser driver 138 that converts electrical signals to optical signals, renders obvious the *first and second common interfaces*.

164. Therefore, the prior art renders obvious this limitation.

d. [27.4] *with no other processing device intervening between the high speed optical link and said ASIC devices of each of said two or more signal processing circuits.*


165. As discussed in [27.3], the Hartsell-Raj combination renders obvious each processor module (*ASIC device*) connected to interconnection 1080, implementing Raj's optical serial link, (*intervening high speed serial optical link*) via

interface 34, implementing Raj's electrical interface 140 and laser driver 138, (*common interface*), with no other processing device intervening between.

166. Hartsell teaches the processing engines “communicate with each other directly as peers.” Ex. 1005, [0125]. Further, Hartsell depicts no other processing devices between the processor modules 1001 and 1003 and the interconnect 1080. *See* Ex. 1005, Figs. 1C', 1C"; *see also* Ex. 1005, [0058] (“a switch fabric establishes **independent paths from node to node** and data is specifically addressed to a particular node on the switch fabric”). Accordingly, in the architecture suggested by the Hartsell-Raj combination, the processor modules are interconnected by optical communication links without any intermediary processing device.

167. Similarly, Raj's Figs. 1 and 2 show no other processing device between the electrical circuits of the target channel adapters connected through the optical switch 124. *See* Ex. 1007, Figs. 1, 2. Raj further teaches the “fiber may be utilized with a fabric switch 124 or a passive star network.” Ex. 1007, 3:5-6. Accordingly, a POSITA would have recognized that under Raj's network structure, circuits are interconnected by the switch fabric without any intermediary processing device. Alternatively, it would have been obvious for each of Hartsell's processor modules (*ASIC devices*) to have an optical link to each other processor module in a “passive star network” as suggested by Raj. Since such a network arrange is described as

“passive,” it would have been obvious to a POSITA that the passive star network embodiment would not include any processing elements between the processor modules and the optical links.



Ex. 1007, Fig. 1 (partial view, annotated), Fig. 2 (annotated).

168. Therefore, Hartsell in view of Raj renders obvious this limitation.

2. Claim 28

- a. [28.1] *The communications infrastructure of claim 27, wherein said interconnection further comprises a high bandwidth interconnection medium.***

169. Hartsell in combination with Raj renders obvious that its *intervening high speed serial optical link* would comprise a *high bandwidth interconnection medium* (to the extent there are differences between the two terms). Claim 28 recites “said interconnection,” but claim 27 does not recite an “interconnection.” *See also* Ex. 1002, 41 (amending claim 27 to remove “an interconnection comprising a high speed serial optical link” and to add “an intervening high speed serial optical link”). The following analysis interprets the recited “interconnection” to refer to the “intervening high speed serial optical link” from claim 27.

170. As noted in limitation [27.3], Hartsell’s interconnect 1080 includes node-to-node path links that connect one processing module to another processing module and transmit data bit-by-bit serially over “data aligned serial link interfaces.” Ex. 1005, [0143]-[0144]. Hartsell explains that interconnect 1080 is an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode. Ex. 1005, [0144]. Such multi-gigabit communication capabilities are retained in the combinations with Raj. Ex. 1007, 1:63-64. A POSITA would have understood that Hartsell’s interconnect 1080 is a *high bandwidth interconnection medium* due to data transmission capabilities of at least 1.77 Gbps. Ex. 1005, [0144]. This is confirmed

Ex. 1003

by the '323 Patent, which explains that “[e]xamples of suitable high bandwidth interconnection mediums include...any interconnection mediums...having a data transmission capability of greater than or equal to about 1 gigabit per second (Gbps).” Ex. 1001, 6:2-7, 13:25-28 (describing an embodiment supporting “data transmission rates between 622 Mbps and 3.125 Gbps”).

171. Because Hartsell describes its interconnect 1080 as an “8x3.54 Gbps fabric” in eight port mode, and a “16x1.77 Gbps fabric” in sixteen port mode—each supporting a data transmission capability greater than 1 Gbps—Hartsell renders obvious this limitation. Ex. 1005, [0144]; Ex. 1001, 6:2-7.

172. Additionally, Raj describes using optical communications technology in which “data rates may be in excess of 2.5 gigabits per second.” Ex. 1007, 1:63-64. Thus, it would have been obvious to a POSITA for Hartsell’s interconnect 1080, when implemented using the optical communications technology of Raj, to support *high bandwidth* communications greater than or equal to at least 1 gigabit per second. Raj further describes transmitting data over fiber optic cables, such as “dispersion shifted fibers (DSF) or dispersion compensated fibers (DCF),” both of which “support **high data rates** with low attenuation.” Ex. 1007, 3:2-4. Such optical fibers were well-known as communications *media*[a].

173. In summary, the prior art describes transmitting data between processing modules via optical fibers at data rates in excess of 1 gigabit per second, which renders obvious “*high bandwidth interconnection medium.*”

3. Claim 31

a. [31.1]-[31.4] (elements recited in claim 31)

174. As noted above in Ground 1, the analysis for elements [31.1]-[31.4] of claim 31 are substantially similar to elements [27.1]-[27.4] of claim 27. And as discussed in further detail below in [31.5], the Hartsell-Raj combination renders obvious a method involving transferring data packets between the signal processing circuits. *See infra*, [31.5]. Thus, these elements are obvious for the same reasons discussed above for claim 27 in Ground 3.

b. [31.5] *selectively transferring at least one data packet from each said packet router of each one of said ASIC devices of each given one of said respective two or more signal processing circuits to each respective packet router of said at ASIC devices of each of the other of said two or more signal processing circuits through said first and second common interfaces and said intervening high speed serial optical link without routing said data through any other intervening processing device between each respective ASIC device and the high speed serial optical link.*

175. As discussed in [27.3], the prior art renders obvious that data may be transmitted between one *packet router* in a processor module 1001 or 1003 (*ASIC device*) of a transport processing engine 1050 (*signal processing circuit*) to another

processor module 1001 or 1003. And as discussed in [27.2], a POSITA would have found it obvious for each of Hartsell's processor modules 1001 or 1003, in view of Raj, to operate as (or to include) a *packet router* because in the combination, each transport processing module would route request packets to their intended destinations over peer-to-peer optical links available through interconnect 1080. Ex. 1005, [0045] ("high speed distributive interconnect... that allows peer-to-peer communication between individual subsystem modules"), [0107] ("Routing of requests... may be controlled by transport and/or interface processing modules"), [0239]. Further, as discussed in [27.3], Hartsell's processor modules 1001 and 1003, configured as transport processing engines 1050, are coupled through a first interface 34 (*first common interface*), interconnection 1080 (*high bandwidth interconnection medium*), and a second interface 34 (*second common interface*). See Ex. 1005, [0144], Figs. 1C, 4.

176. Further, Hartsell teaches *selectively transferring at least one data packet* by disclosing the data communicated is in the form of data packets. See *supra*, [27.2]. Hartsell discloses that transport processing engines 1050 (*signal processing circuits*) have routing functions to route data between processor modules. "Routing of requests...may be accomplished using a distributive interconnect and may be controlled by transport and/or interface processing modules." Ex. 1005, [0107]. It

would have been obvious to a POSITA that routing requests among processor modules includes *selectively transferring at least one data packet*. Hartsell further explains that “information management flow may advantageously occur between any two or more processing engines that may be present on a network and interconnected via a virtual distributed inter connect backplane.” Ex. 1005, [0276]. In the combination with Raj, Hartsell’s distributed interconnect backplane becomes a collection of node-to-node high-speed serial optical links using a packet-based communication protocol such as Infiniband. See *supra*, [27.3].

177. Hartsell’s transport processing engine 1050 performs “network transport protocol sub-tasks, such as processing content requests received from network interface engine 1030.” Ex. 1005, [0072]. A POSITA would have understood that the received requests would include packets, because network interface engine 1030 “process[es] high speed streams of network packets,” Ex. 1005, [0066], requests are passed to network transport processing engine 1050, and transport processing engine 1050 performs “packet header generation” as part of its tasks. Ex. 1005, [0072]-[0073]; *see also* Ex. 1005, [0127] (“each processing engine” manipulates data packet headers by, e.g., setting priority levels), [0138] (processing engines reduce “the packet size transmitted across the distributed interconnect” through protocol layer manipulation).

178. Further, Hartsell and Raj both discuss using Infiniband, as noted above. *See supra*, [27.2]-[27.3]; Ex. 1005, [0231], [0275]; Ex. 1007, 1:5-26, 1:65-67. A POSITA would have understood that the Infiniband protocol transmits data in the form of packets. *See* Ex. 1013, Abstract (discussing switching Infiniband data packets); Ex. 1014, Abstract (discussing converting ATM cells to Infiniband packets). Accordingly, a POSITA would have found it obvious for Hartsell's transport processing engine 1050, in the Hartsell-Raj combination, to *selectively transfer[] at least one data packet*.

179. Therefore, the prior art renders obvious this limitation.

4. Claim 33

- a. [33.1] *The method of claim 31, wherein said interconnection further comprises a high band width interconnection medium.***

180. See analysis at [28.1] in Ground 3. Like claim 28, claim 33 recites "said interconnection," though claim 31 does not recite an "interconnection." The analysis presented above at [28.1], which interprets the recited "interconnection" to refer to the "intervening high speed serial optical link," similarly applies to claim 33.

XII. CONCLUSION

181. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and that these statements were made with 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.

Dated: June 9, 2023



Daniel Blumenthal