

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Inter Partes Review

Cisco Systems, Inc.,
Petitioner,
v.

Lionra Technologies Ltd.,
Patent Owner.

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

DECLARATION OF DR. ERIC COLE

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
II.	QUALIFICATIONS.....	2
III.	SCOPE OF WORK	9
IV.	LEVEL OF ORDINARY SKILL IN THE ART.....	10
V.	LEGAL PRINCIPLES.....	12
	A. Anticipation.....	12
	B. Obviousness	14
VI.	ASSERTED PRIOR ART	17
VII.	HARTSELL DOES NOT DISCLOSE “ASIC DEVICES THAT EACH ITSELF INCLUDES A PACKET ROUTER”	18
VIII.	THE COMBINATION OF HARTSELL AND HAUCK DOES NOT DISCLOSE “ASIC DEVICES THAT EACH ITSELF INCLUDES A PACKET ROUTER”	37

I, Eric Cole, of Ashburn, Virginia, declare that:

I. INTRODUCTION

1. Lionra Technologies, Ltd. has retained me as an independent expert in connection with IPR2023-1047 to provide testimony regarding the technology involved in U.S. Patent No. 7,921,323 (the “’323 Patent”) and how the inventions described and claimed therein compare with the asserted prior art. I have been asked to provide my technical review, analysis, insights, and opinions regarding the references set forth in the above-referenced IPR by Petitioner Cisco Systems.

2. I am being compensated for my work in this matter at my standard hourly rate for consulting services at \$700 per hour. My compensation is not dependent on the outcome of the case and does not affect the substance of my statements in this Declaration. I have no financial interest in the matter or the patents at issue.

3. I have reviewed the specification and the claims of the ’323 Patent, as well as the prosecution histories for the ’323 Patent. I am familiar with the technology involved therein, having practical, hands-on experience in the field for over 30 years. In addition, in preparing this Declaration, I have reviewed and

considered the Petition, Dr. Blumenthal's declaration submitted in this proceeding, and each document cited in my declaration.

II. QUALIFICATIONS

4. My academic credentials include a master's degree in computer science and a doctorate in information security, and I have worked in the networking, cyber and technical information security industry for over twenty-five years. The details of my education and experience are summarized in my curriculum vitae ("CV"), which is attached as Exhibit 2008. My CV also includes a list of my publications over the past ten years.

5. I have worked in networks and computer security for over twenty-five years. To briefly summarize my career, I have: 1) been responsible for the security of sensitive government projects; 2) managed networks and network security at major private sector companies; 3) consulted for major private sector companies; 4) been actively involved in designing, building, and deploying intrusion detection systems; 5) worked on protecting and securing critical information including trade secrets; 6) continually taught and written about cybersecurity.

6. I am the founder and president of Secure Anchor Consulting, LLC, located in Ashburn, Virginia. Secure Anchor Consulting provides cybersecurity and

network consulting services and lead research and development initiatives to advance information systems security. Many of the clients I work with are responsible for protecting sensitive client and/or consumer information and implementing effective security. Throughout my career I have designed, built, and implemented many networks for both Fortune 500 companies and the federal government.

7. At Secure Anchor Consulting, I also actively work on incidents, and I have a detailed understanding of measures that organizations can take to both prevent and detect breaches. In addition, I have worked with organizations to design, build, and implement security measures that would be classified as both fixes and compensating controls, and I have detailed knowledge of the effort that is required to implement various remediation measures.

8. I worked for the U.S. government for eight years as an employee and have held various contracting jobs with government agencies, which involved working with classified information. I have worked for a wide range of government organizations, including the Federal Bureau of Investigation (FBI), the National Security Agency (NSA), the Central Intelligence Agency (CIA), the Department of Energy, the Department of Defense (DOD), the Treasury Department, the Secret

Service, and the Nuclear Regulatory Commission (NRC). I have held various top-secret security clearances with the DOD, the CIA, and the NRC. I also was involved in funded research from the Defense Advanced Research Projects Agency (DARPA), an agency of the DOD responsible for the development of new intrusion detection technologies for use by the military. During this time, I have reviewed, edited, and created numerous security policies. On several occasions, I have also reviewed the reasonableness of measures intended to protect and secure Confidential Information through various types of documentation to include policies, procedures, standards, and baselines.

9. While serving as a Senior Officer for the CIA as Program Manager / Technical Director for the Internet Program Team with the Office of Technical Services, I implemented the Internet Program Team that designs, develops, tests, and deploys internet security products in three- to six-month intervals. In this role, I received a letter of appreciation from the Director of Central Intelligence and six Exceptional Performance Awards.

10. As a member of the Information Security Assessment Team with the Office of Security at the CIA, I evaluated and performed security assessments of network operating systems to identify potential vulnerabilities and solutions. I also

designed a large-scale auditing system with an automated review capability and worked on several virus investigations for the Office of Security.

11. In my role as Chief Information Officer for the American Institutes for Research, I repaired and developed IT infrastructures for various organizations. I was also responsible for building out the entire data classification program to ensure that critical and proprietary data was properly protected and secured. I was also one of the principal leads on research funding from DARPA. In this role, I was responsible for understanding the current state of the art in intrusion detection and designing and building a new cybersecurity system that could catch attacks that were targeting organizations.

12. As Chief Scientist and Senior Fellow for Lockheed Martin, I performed research and development in information systems security. I also specialized in evaluating and designing secure network design, perimeter defense, vulnerability discovery, penetration testing, and intrusion detection systems. At Lockheed Martin, I served as a technical advisor in high-profile security projects for government clients including the DOD, the FBI Sentinel case management systems, Department of Homeland Security Enterprise Acquisition Gateway for Leading Edge solutions,

Jet Propulsion Labs, Hanford Labs, and FBI Information Assurance Technology Infusion programs.

13. As Chief Technical Officer for McAfee, I executed the strategy for technology platforms, partnerships, and external relationships to establish product vision and achieve McAfee's goals and business strategies. In this capacity, I worked closely with groups tasked with the development of intellectual property. While at McAfee, I worked with many business customers responsible for protecting clients' information, helping them implement effective security solutions.

14. I was a fellow and instructor for twenty years with the SANS Institute, a research and education organization consisting of information security professionals. SANS is the largest source for information security training and security certifications in the world. I am an author of several security courses, such as SEC401-Security Essentials and SEC501-Enterprise Defender. SEC401-Security Essentials is a required course for agents of the FBI and the NSA. These courses have been taken by over 30,000 people, and I have taught more than 50% of these students via live and online training. The courses that I have written have multiple modules and lessons on security policies, employee agreements, and protection of

intellectual property. As one of the top-rated instructors at SANS, I have delivered keynote addresses around the world.

15. I am a member of the European InfoSec Hall of Fame, a professional membership awarded by nomination and election by a panel of industry experts.

16. I hold a Certified Information System Security Professional (CISSP) certification and was the creator of the Global Information Assurance Certification (GIAC) Security Essentials Certification (GSEC).

17. I am a contributing author of “Securing Cyberspace for the 44th President” and served as a Commissioner on cybersecurity for President Obama. My Nine books on cybersecurity include “Online Danger,” “Network Security Bible - 2nd Edition,” “Advanced Persistent Threat,” and “Insider Threat,” which have become recognized as industry-standard sources. They are used in many classrooms and colleges around the world as authoritative texts on cybersecurity. I have also published several articles in the field.

18. The focus of my work over the last ten years has been on data protection for commercial organizations. I am very familiar with best practices in cybersecurity and have developed many of the standard practices.

19. Today, my consulting efforts help organizations properly protect client information and relate to various alternative methods of security that can be implemented. This includes hardening of systems, securing servers, proper configuration, patching, and compensating controls. In cases where a fix cannot be made because it would negatively impact the business, I have extensive experience in understanding various compensating controls as an alternative to minimize the risk to an acceptable level and how to implement those controls.

20. I routinely help clients develop and implement company-wide policies, training, and best practices to prevent potential cybersecurity incidents. My work addresses threats from both internal and external sources, including accidental and intentional incidents.

21. I have led, and been involved in working on, responses to many cybersecurity incidents. This includes performing root cause analysis, verifying the reasonableness of security measures, and validating the implementation of additional remediation measures.

22. In my consulting work and in the various positions that I have held, I have been responsible for protecting organizations to minimize and prevent data breaches, in addition to responding to breaches post detection. I have a very detailed

understanding of what is and is not a reasonable level of security based on the hundreds of organizations with which I have worked and the various security task forces with which I have been involved.

III. SCOPE OF WORK

23. I have been asked to consider and provide my opinions on the claims of the '323 Patent, including the technical subject matter and the application of various references that predate the priority date of the '323 Patent. In particular, I have been asked to consider what one of ordinary skill in the art would have understood from the claims and specification of the '323 Patent, whether certain references disclose or suggest the features recited in the claims of the '323 Patent, and whether one of ordinary skill in the art would have had reason to combine certain references to arrive at the claimed subject matter.

24. As I noted previously, my compensation is in no way contingent on the nature of my findings, the presentation of my findings in testimony, or the outcome of any proceeding.

25. My opinions set forth below are based on my education, training, experience, and the content of the references considered.

IV. LEVEL OF ORDINARY SKILL IN THE ART

26. My opinions are provided based on what a person of ordinary skill in the art in the technical field of the inventions of the '323 Patent would have understood at the time of filing of the '323 Patent.

27. I have been informed and understand that the content of a patent and prior art should be interpreted the way a person of ordinary skill in the art would have interpreted those references at the time of filing of the patent. I understand that the person of ordinary skill is a hypothetical concept referring to one who understands the technology and associated conventional wisdom at the time.

28. I was asked to provide an opinion as to the level of one of ordinary skill in the art pertinent to the subject matter set forth in the '323 Patent at the time of the priority date. Around November 16, 2006 (the filing date of patent application 11/600,934 as a continuation-in-part of patent application 10/843,226), it is my opinion that a person of ordinary skill in the art in the field of the '323 Patent would have been a person with a Bachelor's degree in computer science, computer engineering, or electrical engineering and at least two years of experience in computer networking. Throughout this declaration, I use the term "POSITA" to refer to a person with these or similar qualifications. To have similar qualifications,

an individual with additional education or additional industrial experience could still be of ordinary skill in the art if that additional education or experience compensates for a deficit in one of the other aspects of the requirements stated above. The Petition and Dr. Blumenthal assert that a POSITA “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.” Petition at 11-12; Ex. 1003, ¶ 22. While I disagree with the amount of experience they assert is required, I note that, applying their asserted level of skill, my opinions contained herein do not change.

29. The basis for my familiarity with the level of ordinary skill is my interaction with large numbers of workers in computer networking who were at this level of skill. In reaching this opinion as to the hypothetical person of ordinary skill in the art, I have considered the types of problems encountered in the art, the prior art solutions to those problems, the rapidity with which innovations are made, the sophistication of the technology, and the educational level and professional

capabilities of workers in the field. I further note that, as of November 16, 2006, I was at least at this level of skill.

V. LEGAL PRINCIPLES

A. Anticipation

30. I understand and have been informed that in order for an inventor to be entitled to a patent, the invention must be “new,” and the inventor must not have lost his or her right to a patent by delaying filing of a patent application directed to the invention. There are a number of ways that an invention can be “anticipated”—that is, not new.

31. One such way is when the invention was already patented or described in a printed publication anywhere in the world before the date of invention by the inventor or more than one year before the earliest effective filing date of the patent application. Additionally, a patented invention may not be new if it was on sale or publicly used more than one year before the filing of an application for patent. A patented invention is also not new if it was “known” or used by others in the United States before it was invented by the patent applicants. To invalidate a patent claim based on knowledge, I understand that the knowledge must be accessible to the

public. If the subject matter covered by a patent claim is found within the “four corners” of a prior art reference, the claim is said to be anticipated by that reference.

32. A patent claim can be said to be anticipated if each and every limitation of the patent claim is found either expressly or inherently in a single item of prior art. While a prior art reference need not use the same words as a patent claim to anticipate the claim, the prior art reference must describe the requirements of the claim with sufficient clarity such that a POSITA would have been able to make and use the claimed invention based on the reference and his or her knowledge in the applicable technical field.

33. The disclosure of a feature of a claimed invention can be either “express,” meaning that the text, figures, or other content of a reference actually teaches the aspects of the patent claim, or the disclosure can be “inherent.” An express disclosure does not need to use the same words as the patent claim being considered and does not need to be depicted in exactly the same way, but the disclosure must actually convey the claimed invention to a POSITA using some combination of words, figures, data, or other portions of the patent. I have been informed and understand that in order to establish that an element of a claim is “inherent” in the disclosure in a prior art reference, it must be clear to a person skilled

in the art that the missing element is an inevitable part of what is explicitly described in the reference, and that it would have been recognized as necessarily present by a person skilled in the art. Inherency cannot be established by probabilities or possibilities, and the mere fact that something may result from a given set of circumstances is not enough to establish inherency.

34. Finally, I further understand that to anticipate a claim, the elements of the claim must be found in the prior art reference “arranged as in the claim.” I have been informed that this means that the elements of the prior art reference may not be recombined in a way that is not taught in the prior art to arrive at the invention; instead, the prior art must teach the claimed subject matter in the arrangement required by the claims.

35. I understand that any disclosure that anticipates a claim also renders that claim obvious.

B. Obviousness

36. I have been informed and understand that subject matter claimed in a patent is obvious under 35 U.S.C. § 103 if a person of ordinary skill in the art at the time the alleged invention was made would have had reason to combine or modify the disclosures of one or more prior art references to arrive at the claimed subject

matter. A claim may be unpatentable even if each and every claim limitation is not present or disclosed in a single prior art item.

37. I have been informed and understand that, under the doctrine of obviousness, a claim is unpatentable if the differences between the 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 person having ordinary skill in the art to which the subject matter pertains. A person of ordinary skill in the art is presumed to have knowledge of the relevant prior art at the time of the claimed invention.

38. I have been informed and understand that obviousness is based on the scope and content of the prior art, the differences between the prior art and the claim, the level of ordinary skill in the art, and secondary indicia of obviousness and non-obviousness to the extent such indicia exist.

39. I have been informed and understand the determination of whether the asserted claims would have been obvious to a person of ordinary skill in the art and, therefore, invalid is not governed by any rigid test or formula. A determination that a claim is obvious is, instead, based on a common-sense determination that the claimed invention is merely a combination of known limitations to achieve predictable results. Any of the following rationales are acceptable justifications to

conclude that a claim would have been obvious: (1) the claimed invention is simply a combination of known prior art methods to yield predictable results; (2) the claimed invention is a simple substitution of one known element for another to obtain predictable results; (3) the claimed invention uses known techniques to improve similar devices (methods or products) in the same way; (4) the claimed invention applies a known technique to a known device (method or product) ready for improvement to yield predictable results; (5) the claimed invention was “obvious to try”—choosing from a finite number of identified, predictable solutions with a reasonable expectation of success; (6) there is known work in one field of endeavor that may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations would have been predictable to one of ordinary skill in the art; or (7) there is some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill in the art to modify the prior art reference to combine prior art teachings to arrive at the claimed inventions.

40. I have been informed and understand that a claim may be obvious in light of a single reference, without the need to combine references, if the elements of the claim that are not found in the reference can be supplied by the common sense

or knowledge of one of ordinary skill in the art or taught in different areas of the single reference.

41. I have been informed and understand that an analysis of whether a claimed invention is obvious must not rely on a hindsight combination of prior art. The analysis must proceed in the context of the time of the invention or claimed priority date and consider whether the invention as a whole would have been obvious to a person of ordinary skill in the art, taking into consideration any interrelated teachings of the prior art, the effects of demands known to the design community or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art, all in order to determine whether there was an apparent reason to combine any known elements in the fashion claimed by the patent at issue.

VI. ASSERTED PRIOR ART

42. The Petitioner and its expert, Dr. Daniel Blumenthal, have asserted that certain claims of the '323 Patent are obvious over the combination of two or more of the following prior art references: (1) U.S. Published Patent Application No. 2002/0059274 ("Hartsell"); (2) "The Roles of FPGA's in Reprogrammable Systems" by Scott Hauck ("Hauck"); and (3) U.S. Patent No. 6,496,291 ("Raj").

43. The particular combinations that the Petitioner and Dr. Blumenthal rely on are as follows: (1) in Ground 1, independent claims 27 and 31 of the '323 Patent, as well as dependent claims 28 and 33, are obvious over Hartsell; (2) in Ground 2, independent claims 27 and 31 of the '323 Patent, as well as dependent claims 28 and 33, are obvious over the combination of Hartsell and Hauck; (3) in Ground 3, independent claims 27 and 31 of the '323 Patent, as well as dependent claims 28 and 33, are obvious over the combination of Hartsell and Raj; and (4) in Ground 4, independent claims 27 and 31 of the '323 Patent, as well as dependent claims 28 and 33, are obvious over the combination of Hartsell, Hauck and Raj.

44. For purposes of this declaration, I restrict my review and analysis to Grounds 1 and 2 and the independent claims of the '323 Patent with respect to the Hartsell and Hauck references.

VII. HARTSELL DOES NOT DISCLOSE “ASIC DEVICES THAT EACH ITSELF INCLUDES A PACKET ROUTER”

45. Each of independent claims 27 and 31 of the '323 Patent recites “each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router.” *See* Ex. 1001, 31:39-41 (Cl. 27); Ex. 1001, 32:15-17 (Cl. 31). The Petition and Dr. Blumenthal assert that Hartsell renders this limitation obvious. I disagree. Hartsell does not describe that Hartsell’s processor

modules 1001 and 1003, of which transport processing engine 1050 is made, operate as a packet router, either on their own or when the network interface processing engine 1030 is incorporated.¹ And a POSITA would not have found it obvious to modify Hartsell such that “each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router.”

46. First, the Petition and Dr. Blumenthal assert that “Hartsell’s processor modules (ASIC devices) each itself includes a packet router because each processor module controls packet routing.” Petition at 35; Ex. 1003, ¶ 91. I disagree. Notably, the Petition and Dr. Blumenthal do not assert that the processor module is a “packet router,” but that it “controls packet routing.” The Petition and Dr. Blumenthal further assert that the processor modules “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.” Petition at 36; Ex. 1003, ¶ 93. To support this proposition, the Petition and Dr. Blumenthal contend that “Hartsell explains that its transport processing engines 1050 (which are implemented via multiple processor

¹ As Dr. Blumenthal did, *see* Ex. 1003, ¶ 56, I have interpreted the term “packet router” under its plain and ordinary meaning to a POSITA.

modules 1001 or 1003) ‘may perform time out checks for each network connection, session management, data reordering and retransmission, data queuing and flow control, packet header generation, etc.’ Ex. 1005, [0073].” Petition at 36; Ex. 1003, ¶ 93.

47. Neither the Petition nor Dr. Blumenthal explain how performance of these recited functions, namely “time out checks for each network connection, session management, data reordering and retransmission, data queuing and flow control, packet header generation,” make transport processing engines 1050, processors modules 1001/1003, or the FPGA within these processor modules into a “packet router.” Furthermore, neither the Petition nor Dr. Blumenthal explain how the ’323 Patent describes these functions as defining a “packet router.” As explained further below, a POSITA would not have understood time out checks for each network connection, session management, data reordering and retransmission, data queuing and flow control, or packet header generation to be indicative or suggestive of Hartsell’s transport processing engine, processor modules 1001/1003, or the FPGA identified within these processor modules including or functioning as a packet router.

48. A POSITA would have understood that a “packet router” is a device that operates in the network layer, the third layer of the Open Systems Interconnection (“OSI”) model, and that a packet router functions to receive packets on an input, examine network addresses of a packet, e.g., the IP address in the Internet Protocol (IP), to determine the appropriate output path for the packet from among multiple output paths and forward the packet along that path. *See, e.g.,* Ex. 2009, *Introduction to Networking*, Reed, Kenneth D., at 3 (“A router is a layer 3 device, with several ports that can each connect to a network or another router. A router examines the logical network address of each packet, then uses its internal routing table to forward the packet to the routing port associated with the best path to the packet’s destination.”), Ex. 2010, *Computer Networks, Architecture, Protocols, and Software*, Hsu, John Y., at 10-12 (“[A] router is the switching node to support many different ports as shown in Figure 5.21(c). Its DL [data link] layer receives frames from one port and passes the inner packet to its N [network] layer to determine the route.”).

49. Consistent with the understanding of a POSITA, the ’323 Patent describes Field Programmable Gate Array (“FPGA”) based PRISM packet routers that can receive packets on multiple “reader” input interfaces and determine the best

path for the packet through the network of FPGA-based packet routers, what the '323 Patent calls a "PRISM matrix," and select the appropriate one of multiple "writer" output interfaces to forward the packet on. For example, the '323 Patent discloses how each "writer" output interface in a PRISM packet router can forward packets from any one of the "reader" input interfaces:

As illustrated in the exemplary embodiment of FIG.9, each packet destination or writer interface 832 of PRISM router 702 may be configured with a connect multiplexer ("mux connect") 900 that is capable of accessing each available PRISM reader interface 830a to 830m (e.g., each reader interface available in the local FPGA or other local ASIC device). Within a PRISM router 702 there may be multiple such packet destinations or writer interfaces 832 with a respective connect multiplexer 900 assigned to each corresponding packet destination. Using such an exemplary configuration, a packet transfer request may be asserted at a given reader interface 830 of the PRISM router 702 in combination with packet destination information that corresponds to the connect multiplexer 900 of the desired packet destination or writer interface 832.

Ex. 1001, 15:62-16:9.




FIG. 9

Id., Fig. 9.

50. Figure 10 of the '323 Patent further illustrates how the multiplexer connect module of Fig. 9 is replicated for so that the “writer” output interfaces, 832a to 832n, match the “reader” input interfaces, 830a to 830n:




FIG. 10

Id., Fig. 10 (showing PRISM multiplex matrix module combining $x+1$ multiplex connect modules for the “writer” output interfaces, 832a to 832n).

51. Figures 9 and 10 illustrate the multiplex connect and multiplex matrix module features of the PRISM packet router as having parallel bus inputs and outputs, e.g., 952 and 962, respectively. The '323 Patent describes that the PRISM router FPGAs are interconnected with duplex serial communication links. *See, e.g., id.*, 21:23-35 (“FPGA devices 102,104,106 and 108 ... with duplex serial data communication links 1520 provided between each given two of FPGA devices 102,104,106 and 108 (e.g., via high speed serial I/O connections in the form of MGTs), and with each FPGA being provided with a packet router interface switch matrix (‘PRISM’) ...”). Consequently, the '323 Patent further describes with respect to Figs. 11 and 12 the transmit and receive circuitry to connect the serial I/O of the FPGA with the “reader” and “writer” interfaces of the PRISM multiplex matrix module:




FIG. 11

Id., Fig. 11 (showing **mgt_connect** module to interface FPGA serial I/O with complementary “reader” interface 830 and “writer interface” 832); *see also id.*, 18:6-52.




FIG. 12

Id., Fig. 12 (showing multiple mgt_connect modules for duplex links to “Other FPGAs and/or External Interfaces”); *see also id.*, 18:53-19:3.

52. The '323 Patent describes that the best path is generally for “a packet [to] cross no more than one duplex data communication link to reach any destination in the PRISM matrix,” *id.*, 15:24-29, although “additional routing capability” may be added so that “when insufficient bandwidth is provided by all available direct duplex data communication links between two devices (e.g., between source and destination FPGA devices), additional bandwidth may be obtained by relaying all or a portion of packets through more than one duplex data communication link (e.g., by relaying packets from a source FPGA to an intermediate second FPGA and then to the destination FPGA through the PRISM router/s of one or more other FPGAs),” *id.*, 15:49-59.

53. In the exemplary embodiment of the '323 Patent, the PRSIM routing determinations are made based on the destination FPGA:

As illustrated for the exemplary embodiment of FIG. 6, packet destination information may include, for example, both an FPGA destination code and a “local destination code within the destination FPGA. In one example, the packet destination information may be written so that either the FPGA destination code matches the code for the current FPGA, or in the event that it does not, so that a specific MGT destination interface is provided. Thus, if the FPGA destination code does not match the code for the current FPGA, then one of the

MGT interfaces will match an available MGT and the packet will be routed off the FPGA to the specified MGT destination.

Id., 16:10-21.

54. The '323 Patent also makes clear that other packet protocols that have their own addressing schemes may be implemented in lieu of the PRISM routing scheme, such as TCP/IP with source and destination IP addresses. *See, e.g., id.*, 14:41-48, 19:4-14, 22:63-23:1. The Petition and Dr. Blumenthal state that “[t]he '323 Patent explains that any routing scheme, including TCP/IP and Infiniband, would involve a packet router,” and argue that Hartsell’s disclosure of “Infiniband” means that “processor modules 1001 or 1003 perform[] routing of Infiniband packets.” Petition at 40, Ex. 1003, ¶ 97. I do not agree. As a POSITA would readily recognize, not **all devices** operating according to the TCP/IP or Infiniband protocols must be a “packet router” and that there are often, if not always, many points in a packet network that are not packet routers. For example, where a personal computer (PC) is connected to a local area network (LAN), the network interface card (NIC) in the PC may connect to the LAN and transmit and receive packets to/from the LAN using the NIC. The NIC is an endpoint in the LAN and is not, nor does it function as, a “packet router.” By the same token, a POSITA would not understand the processor on modules 1001 and 1003, including the transport processing engine

1050, to be a “packet router” and would not find it obvious to modify these processors to perform the functions of a packet router.


55. I disagree that Hartsell’s processor modules “would have operated as a packet router in the same way as the FPGAs described in the ’323 Patent,” as the Petition and Dr. Blumenthal assert.

56. Hartsell does not describe the transport processing engine 1050 as a “packet router.” To maximize the processing capability of the transport processing engine 1050, Hartsell limits its networking processing requirements. *See* Ex. 1005, [0079] (“[T]he transport processing engine 1050 does not have to process the header information of every packet of a connection. Rather, the transport interface engine merely receives a contextually meaningful identifier or tag that identifies the previous processing results for that connection.”); *see also id.*, [0080] (“In one embodiment, the data link, network, transport and session layers (layers 2-5) of a packet may be replaced by identifier or tag information. For packets related to an established connection the transport processing engine does not have to perform intensive processing with regard to these layers such as hashing, scanning, look up, etc. operations.”). Moreover, the fact that 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.,” *see id.*, [0073], does not indicate that the transport processing engine 1050 contains a packet router.

57. With respect to “time out checks for each network connection,” it is the transport layer, the fourth layer in the OSI model, that is responsible for managing connections between network devices using protocols like the Transport Control Protocol (TCP) or User Datagram Protocol (UDP). *See, e.g.*, Ex. 2011, *Networking Fundamentals*, Roberts, Richard M., at 3-5; Ex. 2010, *Computer Networks, Architecture, Protocols, and Software*, Hsu, John Y., at 3-5, 7. With respect to “session management,” it is the session layer, the fifth layer in the OSI model, that is responsible for managing sessions. *See, e.g.*, Ex. 2011, *Networking Fundamentals*, Roberts, Richard M., at 3-5; Ex. 2010, *Computer Networks, Architecture, Protocols, and Software*, Hsu, John Y., at 3-5, 7. With respect to “data reordering and retransmission, data queuing and flow control,” it is the transport layer, the fourth layer in the OSI model, that is responsible for managing TCP and UDP connections. *See, e.g.*, Ex. 2011, *Networking Fundamentals*, Roberts, Richard M., at 3-5; Ex. 2010, *Computer Networks, Architecture, Protocols, and Software*, Hsu, John Y., at 3-5, 7.

58. “Packet header generation,” e.g., generation of the header for an IP packet, does not have any effect on the path out of the transport processing engine 1050 and is therefore not a packet routing function. The transport processing engine 1050, which both the Petition and Dr. Blumenthal identify as a “signal processing circuit,” *see* Petition at 23, Ex. 1003, ¶ 74, have only a single connection to the distributive interconnection 1080 and therefore do not make a routing, or path, determination for packets transmitted from the transport processing engine 1050. Moreover, the FPGA on each of the transport processing engines 1050 through which the single connection is made, e.g., 1050A or 1050B in Fig. 1D” reproduced below, has no networking functionality. The FPGA is only responsible for converting the connection formats used by the processor on one side and the IBM UADSL circuit on the other side.



Ex. 1005, Fig. 1D'' (excerpted).

59. More specifically, with respect to Fig. 4, reproduced below, Hartsell discloses:

The interface converter circuits 1082 convert the data aligned serial link interface (DASL) to a UTOPIA (Universal Test and Operations PHY Interface for ATM) parallel interface. FPGAs (field programmable gate array) may be utilized in the processor modules as a fabric interface on the processor modules as shown in FIG. 1C. These fabric interfaces provide a 64/66 Mhz. PCI interface to the interface converter circuits 1082. FIG. 4 illustrates a functional block diagram of such a fabric interface 34. As explained below, the interface 34 provides an interface between the processor module bus and the UDASL switch interface converter circuit 1082. As shown in FIG. 4, at the Switch fabric side, a physical connection interface 41 provides connectivity at the physical level to the switch fabric. An example of interface 41 is a parallel bus interface complying with the UTOPIA standard. In the example of FIG. 4, interface 41 is a UTOPIA 3 interface providing a 32-bit 110 Mhz connection.

Id., [0144].




FIG. 4

Id., Fig. 4.

60. In addition to performing queueing of messages transmitted and received by the transport processing engine 1050, the FPGA performs segmentation and reassembly, and associated buffering, to convert between the 64-bit, 66 MHz PCI bus on the side of the transport processing engine 1050 and the 32-bit, 110 MHz UTOPIA bus on the side of the UDASL switch 1082. Because there is only a single path on each side, Hartsell does not teach, nor would a POSITA find it obvious, that the FPGA makes any routing, or path determination, for data received on either side. In other words, data being transmitted from the transport processing engine 1050 will be received by the FPGA on the 64-bit, 66 MHz PCI bus and, after required buffering, segmentation and/or reassembly, will be transmitted on the 32-bit, 110 MHz UTOPIA bus, or vice-versa, without the need for a routing determination. I further note, as described below in ¶¶ 74-76, Hartsell describes the use of an IBM PRIZMA-E crossbar switch fabric as the distributive interconnect 1080 in the content delivery system. It is the crossbar switch fabric that is responsible for switching packets from one engine/module to another engine/module in Hartsell's content delivery system. *See* Ex. 1005, [0061] ("Among the advantages offered by such a switch fabric interconnection ... are throughput, scalability and fast and

efficient communication between individual discrete content delivery engines of content delivery system 1010.”).

61. For these reasons, I disagree with the Petition and Dr. Blumenthal that Hartsell describes or renders obvious that the FPGA of the transport processing engine 1050 contains a “packet router.”


62. The Petition and Dr. Blumenthal further suggest that “it would have been obvious for the processor ... 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.” Petition at 38, Ex. 1003, ¶ 95. I disagree for the same reasons set forth above. Furthermore, the processor, a Pentium III on the 1001 modules or a PowerPC on the 1003 modules, see Ex. 1005, Figs. 1C’ and 1C’’, does not act as a packet router. The processor is not responsible for the path determination to the destination. A POSITA would not take the processor on modules 1001 and 1003, including the transport processing engine 1050, to be a “packet router,” nor would a POSITA find it obvious to modify the processor on modules 1001 and 1003 to include a “packet router.” The processor modules 1001 and 1003 are endpoints with a single point of interconnection through the FPGA to the distributive interconnection 1080 and

therefore do not make a routing, or path, determination for transmitted packets transmitted.

63. The Petition and Dr. Blumenthal assert that, “[a]lternatively, the incorporation of the network interface processing engine 1030 into the transport processing engine 1050 further renders obvious the claimed packet router.” Petition at 39, Ex. 1003, ¶ 96. I disagree. Other than noting that the “interface engine 1030 may process[] all or some of the TCP/IP protocol stack,” Ex.1005, [0075], neither the Petition nor Dr. Blumenthal explain how the combination would exhibit any different behavior that would qualify the combination as a packet router. Dr. Blumenthal cites to Hartsell [0063], *see* Ex. 1003, ¶ 96, stating that the network interface processing engine 1030 “may be suitable” for various network functions, but does not analyze the extent to which those functions are performed in the engine. Hartsell notes that the engine “may be employed to”: (1) “process or terminate one or more layers of the network protocol stack”; (2) “perform look-up intensive operations”; (3) “load balance among other content delivery processing engines”; and (4) perform “security processing.” Ex. 1005, [0063]. This paragraph does not indicate the network processing engine is responsible for any packet routing in Hartsell’s content delivery system. Dr. Blumenthal similarly cites to Hartsell [0066],

see Ex. 1003, ¶ 96, stating that the 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,’” but this paragraph in Hartsell only indicates the “conventional application” and “conventional operation” of the “specialized” “network processor” used in the network interface processing engine 1030, not that the network interface processing engine 1030 performs this function. See Ex. 1005, [0066].

64. Fig. 2 of Hartsell, reproduced below, shows an embodiment where network interface processing engine 1030 and transport processing engine 1050 are combined into a single networking subsystem 205:



Id., Fig. 2; *see also id.*, [0162] (“FIG. 2 illustrates a single networking subsystem processor module 205 that is provided to perform the combined functions of network interface processing engine 1030 and transport processing engine 1050 of FIG. 1A.”)

65. Although the combined engine may reduce traffic across the distributive interconnect 1080, a distributive interconnect is still required for communication between the networking subsystem processor module 205 and the other subsystems, e.g., application RAM subsystem 220, application processing subsystem 224, storage subsystem 210, file system cache subsystem 215, other subsystems 275, shared resources 255 and system monitor 240. *See id.*, Fig. 2. The embodiment of Fig. 2 breaks the distributive interconnect into two functions, an inter-process data movement path and an inter-process communications path. *See id.* Hartsell further teaches that, “[a]s illustrated in FIGS. 1A-1F, the functions of inter-process communications path 230 and inter-process data movement path 235 may be together handled by a single distributive interconnect 1080 (such as a switch fabric, for example).” *Id.*, [0161]. With the proposed combination of the network interface processing engine 1030 and the transport processing engine 1050 into the networking subsystem processor module 205, “content requests may be passed

directly from the combined interface/transport engine to network application processing engine 1070 via distributive interconnect 1080.” *Id.*, [0162].

66. A POSITA would recognize that the operation of the transport processing engine component of the networking subsystem processor module 205 with the remaining subsystems of Fig. 2 of Hartsell would be no different than in the preferred embodiment. The FPGA of the transport processing engine 1050, or modules 1001 and 1003, would still be utilized as a fabric interface to the distributive interconnect 1080 (the Petition’s and Dr. Blumenthal’s identified “high speed serial optical link” of the ’323 Patent claims). For the same reasons set forth above with respect to the FPGA of the transport processing engine 1050, or modules 1001 and 1003, the FPGA of the networking subsystem processor module 205, a POSITA would not consider the FPGA to be a “packet router” and would not find it obvious to modify the FPGA to include a “packet router.”

VIII. THE COMBINATION OF HARTSELL AND HAUCK DOES NOT DISCLOSE “ASIC DEVICES THAT EACH ITSELF INCLUDES A PACKET ROUTER”

67. For Ground 2, the Petition and Dr. Blumenthal assert that combination of Hartsell and Hauck would render the claimed requirement of “each one of said two or more signal processing circuits including multiple ASIC devices that each

itself includes a packet router” obvious. I disagree for at least two reasons. First, the proposed combination does not have a “packet router.” Second, the Petition’s and Dr. Blumenthal’s motivation to combine Hartsell with Hauck, that “it would have been obvious to a POSITA to modify Hartsell’s architecture to provide ... direct, point-to-point interconnections between Hartsell’s processing modules 1001 and 1003,” Petition at 55; *see also* Ex. 1003, ¶ 126, is technically incorrect.

68. Hauck describes systems based on multiple FPGAs used in a variety of interconnected fashions, with these systems being used, for example, in complex math calculations, neural networks, and especially for ASIC logic emulation. Ex. 1030, 623 (“Logic emulation shares many of the advantages (and disadvantages) of both prototyping and software simulation. Like a prototype, the circuit to be evaluated is implemented in hardware so that it can achieve high-performance test cycles. Like software simulation, however, the emulation can easily be observed and altered to help isolate bugs. Logic emulation takes a gate-level description of a logic circuit and maps it onto a multi-FPGA system.”). One of the distinguishing characteristics for these multi-FPGA systems is the way the FPGA’s are interconnected such that signals can be sent from one FPGA to another. *Id.* at 626. Hauck describes that the FPGA’s can be interconnected in mesh and crossbar

fashions, including where some FPGAs in the architecture are employed only for purposes of providing signal interconnects between FPGAs that implement logic functions. *Id.*, 626-627. Hauck refers to these signal interconnection FPGAs as “routing-only” FPGAs. *Id.*

69. One particular multi-FPGA system that Hauck describes, and that the Petition and Dr. Blumenthal rely upon, is the Springbok system:

Springbok is designed to support the rapid prototyping of board-level designs, where numerous premade components must efficiently be embedded in an FPGA-based structure. The FPGA’s in the system are used for both routing and rerouting of signals, as well as for the implementation of random logic (Fig. 25). To allow a specific circuit to be implemented in this structure, the system is comprised of a base plate with sites for daughter cards. The daughter cards are large enough to contain both an arbitrary device on the top and an FPGA on the bottom. Note that the device can be a chip such as a processor or memory, I/O elements such as switches and LCD screen interfaces, or whatever else is necessary to implement the system. If necessary, daughter cards can be built that span several locations on the base plate to handle higher space or bandwidth requirements. The daughter cards plug into the base plate, which handles power, ground, and clock distribution, FPGA programming, and interdaughter card routing.

Id., 632.

70. Based on the description of the Springbok system, the Petition and Dr. Blumenthal assert that “Hauck confirms the practice of directly connecting FPGA devices to route packets between the FPGAs was used ‘for both routing and

rerouting of signals’.” Petition at 57; Ex. 1003, ¶ 131 (citing Ex. 1030, 632). I disagree that the Springbok system describes a “device to route packets,” and therefore disagree that, in describing the Springbok system, Hauck discloses a “packet router.”

71. Hauck generally describes that the Springbok system can make hardwired connection between FPGAs on separate daughter cards using the base plate and other electrical interconnection means such as ribbon cables and dummy daughter cards with only hardwired connections. *See* Ex. 1030, 632-633. Consistent with the explanations of FPGAs being used for signal interconnections, Hauck also describes that signal interconnections between two hardwired connections into a single FPGA can be made in that FPGA. Hauck does not teach that the FPGAs in the Springbok system are used to implement any particular form of computer or networking functionality. In other words, Hauck does not teach the functionality used in the FPGAs of the Springbok system. Consequently, Hauck does not teach that the FPGAs in the Springbok system are used to implement a “packet router” as the Petition and Dr. Blumenthal suggest.

72. As set forth above in ¶¶ 49-54, the ’323 Patent discloses a packet router consistent with POSITA’s understanding, i.e., a device that operates in the network

layer, the third layer of the Open Systems Interconnection (“OSI”) model, and that a packet router functions to receive packets on an input, examine network addresses of a packet, e.g., the destination FPGA, to determine the appropriate output path for the packet from among multiple output paths and forward the packet along that path. The ’323 Patent does not disclose a “packet router” that operates based on the existence of physical interconnections between two FPGAs, but rather, consistent with the understanding of a POSITA, operates by making an output path determination for a packet based on the destination of that packet and other factors such as bandwidth sufficiency of a given communication link. *See, e.g.*, Ex. 1001, 15:24-29, 15:49-59. Because hardwired interconnections alone do not make a “packet router,” I disagree that Hauck describes anything related to “packet routing,” and I disagree that the combination of Hartsell and Hauck teaches “ASIC devices that each itself includes a packet router,” as required by the challenged claims of the ’323 Patent.


73. Additionally, I disagree with the Petition’s and Dr. Blumenthal’s assertion that “it would have been obvious to a POSITA to modify Hartsell’s architecture to provide ... direct, point-to-point interconnects between the FPGAs on each of Hartsell’s processing modules 1001 and 1003.” Ex. 1003, ¶ 126; *see also*

Petition at 55. The Petition and Dr. Blumenthal claim that the combination of Hartsell with Hauck would provide “more direct communications paths between pairs of [Hartsell’s] processing modules 1001 and 1003.” Petition at 55, Ex. 1003, ¶ 126. This is false because the combination would provide less direct communications paths.

74. The Petition and Dr. Blumenthal assert that 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.” Petition at 55; Ex. 1003, ¶ 126. In doing so, the Petition and Dr. Blumenthal ignore that Hartsell already discloses the use of a crossbar interconnect. More specifically, Hartsell describes that “[o]ne example interconnection system suitable for use as distributive interconnection 1080 is an 8/16 port 28.4 Gbps high speed PRIZMA-E non-blocking switch fabric switch available from IBM.” Ex. 1005, [0062]. The IBM PRIZMA-E is a crossbar-based switch fabric. *See* Ex. 2012, *The Crosspoint-Queued Switch*, Kanizo, Y., Hay, D., Keslassy, I., IEEE Infocom 2009, pp. 729-737, at 1 (“[C]ombined-input-and-crosspoint-queued (CICQ) switch architectures like IBM Prizma ... combine large virtual output queues (VOQs) in the input linecards

with a buffered-crossbar switch fabric.”); *see also* Ex. 1005, [0058] (“Examples of switch fabric architectures include cross-bar switch fabrics...”).

75. Hartsell’s exemplary IBM PRIZMA-E non-blocking switch fabric provides direct packet communications between the devices connected to the ports of the switch fabric. For example:




Ex. 2012, *The Crosspoint-Queued Switch*, Kanizo, Y., Hay, D., Keslassy, I., *IEEE Infocom 2009*, pp. 729-737, at 1, Fig. 1(a) (annotated to illustrate direct communication from input port I_1 to two output ports, O_1 and O_N).

76. Hartsell indicates that the IBM PRIZMA-E is available with either 8-ports or 16-ports. *See* Ex. 1005, [0062]; *see also id.*, Figs. 1C, 1D, 1E, 1F (showing

8-port switch fabric). In the 8-port IBM PRIZMA-E, for example, any module has a direct communication path to any of the up-to-7 other modules connected to the switch fabric. *See, e.g., id.*, [0060] (“In the case of a switch fabric, latencies are fixed. Further, the connections are ‘direct’, rather than by some undetermined path.”). Similarly, the 16-port IBM PRIZMA-E would allow any module to have a direct communication path to any of the up-to-15 other modules.

77. The Petition and Dr. Blumenthal suggest replacing the exemplary IBM PRIZMA-E crossbar switch fabric with the Springbok interconnect architecture shown in Hauck’s Fig. 25. *See* Petition at 57; Ex. 1003, ¶ 132. The Petition and Dr. Blumenthal further assert that this “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.” *See* Petition at 57; Ex. 1003, ¶ 132. I disagree. As shown in Hauck’s Fig. 25, reproduced below, any one FPGA in the Springbok system is only directly connected with four other FPGAs. Because any FPGA in the Hartsell-Hauck combination could only directly communicate with 4 other FPGAs, as opposed to the FPGAs of Hartsell’s original content delivery system design with the distributive interconnect crossbar switch fabric allowing any FPGA to communicate with 7 other FPGAs with the use

of a 8-port IBM PRIZMA-E, or with 15 other FPGAs with the use of a 16-port IBM PRIZMA-E, I disagree that the proposed combination yields an improvement. To the contrary, replacing the distributive interconnect crossbar switch fabric of Hartsell with Hauck's Springbok interconnect system would yield a **diminishment** of direct communication capability.



Ex. 1030, Fig. 25(a) (annotated to illustrate each Daughter Card FPGA interconnected with only up to 4 other Daughter Card FPGAs).

78. For the foregoing reasons, I disagree that the proposed combination of Hartsell with Hauck discloses the requirement that “each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes

a packet router” of the claims of the ’323 Patent, and that a POSITA would be motivated to make the proposed combination.

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, undersection 1001 of Title 18 of the United States Code.

Date: April 17, 2024

A handwritten signature in black ink, appearing to read 'Eric Cole', written over a horizontal line.

Dr. Eric Cole