

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,

Petitioner,

v.

LIONRA TECHNOLOGIES LIMITED,

Patent Owner.

Case IPR2023-01047

Patent 7,921,323

PATENT OWNER'S RESPONSE

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

TABLE OF CONTENTS

	Page
I. INTRODUCTION.....	1
II. THE '323 PATENT.....	1
III. PERSON OF ORDINARY SKILL IN THE ART	4
IV. ASSERTED GROUNDS.....	4
A. Ground 1: Hartsell does not render obvious claims 27-28, 31, or 33	4
1. Processing packets and communicating packets are not packet routing.....	6
2. The topology of Hartsell's content delivery system is inconsistent with each processor module including a <i>packet router</i>	17
B. Ground 2: The combination Hartsell and Hauck does not render obvious claims 27-28, 31, or 33	22
1. No motivation to combine Hartsell with Hauck	22
2. The Hartsell/Hauck combination would not render any challenged claim obvious.....	25
3. Petitioner has not shown that Hauck is a prior art printed publication.....	28
C. Ground 3: The combination Hartsell and Raj does not render obvious claims 27-28, 31, or 33	29
D. Ground 4: The combination Hartsell, Hauck, and Raj does not render obvious claims 27-28, 31, or 33	29
V. CONCLUSION	30

TABLE OF AUTHORITIES

Cases

<i>Google LLC f/k/a Google Inc. v. Singular Computing LLC</i> , IPR2021-00155, Paper 62 (PTAB May 23, 2022)	27
<i>Hulu, LLC v. Sound View Innovations, LLC</i> , IPR2018-01039, Paper 29 (PTAB Dec. 20, 2019) (precedential)	33
<i>In re Gordon</i> , 733 F.2d 900 (Fed. Cir. 1984)	27
<i>In re Wyer</i> , 655 F.2d 221 (C.C.P.A. 1981)	33
<i>Medtronic, Inc. v. Barry</i> , 891 F.3d 1368 (Fed. Cir. 2018)	33
<i>Nobel Biocare Servs. AG v. Intradent USA, Inc.</i> , 903 F.3d 1365 (Fed. Cir. 2018)	33
<i>Trivascular, Inc. v. Samuels</i> , 812 F.3d 1056 (Fed. Cir. 2016)	27

PATENT OWNER’S EXHIBIT LIST

Exhibit No.	Description
2001	Order, <i>Lionra Technologies Limited v. Fortinet, Inc.</i> , No. 2:22-cv-00322-JRG-RSP (Lead Case), Dkt. No. 21 (E.D. Tex. Oct. 24, 2022)
2002	First Amended Docket Control Order, <i>Lionra Technologies Limited v. Fortinet, Inc.</i> , No. 2:22-cv-00322-JRG-RSP (Lead Case), Dkt. No. 138 (E.D. Tex. August 23, 2023)
2003	Docket Sheet, <i>Lionra Technologies Limited v. Fortinet, Inc.</i> , No. 2:22-cv-00322-JRG-RSP (Lead Case) (E.D. Tex. October 27, 2023)
2004	Order, <i>MyPort, Inc. v. Samsung Electronics Co., Ltd., et al.</i> , No. 2:22-cv-00114-JRG (E.D. Tex. June 12, 2023)
2005	Defendants’ Invalidity Contentions, <i>Lionra Technologies Limited v. Fortinet, Inc.</i> , No. 2:22-cv-00322-JRG-RSP (Lead Case) (E.D. Tex. Feb. 10, 2023)
2006	U.S. District Courts—Federal Court Management Statistics—Profiles—During the 12-Month Periods Ending June 30, 2018 Through 2023 available at https://www.uscourts.gov/statistics/table/na/federal-court-management-statistics/2023/06/30-1
2007	Declaration of Dr. Eric Cole
2008	<i>Curriculum Vitae</i> of Dr. Eric Cole
2009	Excerpts of <i>Introduction to Networking</i> , Reed, Kenneth D.
2010	Excerpts of <i>Computer Networks, Architecture, Protocols, and Software</i> , Hsu, John Y.

2011	Excerpts of <i>Networking Fundamentals</i> , Roberts, Richard M.
2012	<i>The Crosspoint-Queued Switch</i> , Kanizo, Y., Hay, D., Keslassy, I., IEEE Infocom 2009, pp. 729-737

I. INTRODUCTION

Patent Owner Lionra Technologies Ltd. (“Patent Owner”) submits this response to Petitioner Cisco Systems Inc.’s (“Petitioner”) Petition for *inter partes* review of U.S. Patent No. 7,921,323 (“the ’323 patent”).

II. THE ’323 PATENT

The ’323 Patent is generally directed to the use of certain interconnects (such as high-speed serial optical links) among multiple ASIC devices. *See, e.g.*, Ex. 1001, Abstract (“Reconfigurable communications infrastructures may be implemented to interconnect ASIC devices (e.g., FPGAs) and other computing and input/output devices using high bandwidth interconnection mediums. The computing and input/output devices may be positioned in locations that are physically segregated from each other, and/or may be provided to project a reconfigurable network across a wide area. The reconfigurable communications infrastructures may be implemented to allow such computing and input/output devices to be used in different arrangements and applications, e.g., for use in any application where a large array of ASIC devices may be usefully employed such as supercomputing, etc.”)

The ’323 Patent describes a communication infrastructure where multiple signal processing circuits are physically separate and with each signal processing circuit consisting of one or more Application Integrated Circuits (“ASICs”). *See, e.g.*, Ex. 1001, 20:8-17 (“In the practice of the disclosed systems and methods, any

two or more computing devices that are physically segregated from each other may be interconnected as a part of a reconfigurable communications infrastructure to form a reconfigurable network. For example, FIG. 14 illustrates a reconfigurable communications infrastructure 1300 having high bandwidth interconnection medium 1350 interconnecting two signal processing circuits 1314 and 1316 that are physically segregated from each other as illustrated by a physical segregation point 1410 that is positioned therebetween.”).

The '323 Patent further describes an embodiment with multiple signal processing circuits where each signal processing circuit is comprised of multiple ASICs, in each of signal processing circuit are interconnected and are fiber optically connected to the ASICs in another signal processing circuit without any intervening processing devices. *See, e.g., Ex. 1001, 21:17-27* (“FIG. 15 illustrates one exemplary embodiment of a reconfigurable communications infrastructure 1300 having reconfigurable signal processing circuits 1310, 1314 and 1316 that are interconnected by high bandwidth interconnection medium 1350, which in one exemplary embodiment may be a fiber optic transmission medium configured for transporting WDM or DWDM optical data communication. As shown in FIG. 15, reconfigurable signal processing circuit 1314 includes four FPGA devices that may be configured to communicate with each other in a manner as described for FPGA devices 102,104,106 and 108 of FIGS. 1-12 herein.”).




FIG. 15

Ex. 1001, Fig. 15.

With further reference to FIG. 15, the FPGA ASIC devices also include packet routers for routing packets from a source FPGA to a destination FPGA:

As shown in FIG. 15, reconfigurable signal processing circuit 1314 includes four FPGA devices that may be configured to communicate with each other in a manner as described for FPGA devices 102,104,106 and 108 of FIGS. 1-12 herein, i.e., 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”) to route packets between each of the individual FPGA devices via respective duplex serial data communication links 1520 as shown, as well as for intra-FPGA communications.

Ex. 1001, 21:23-35; *see also id.*, 14:31-34 (“[P]acket header ... may contain a routing code that allows the sender to determine the packets destination within a

PRISM infrastructure or matrix.’). Routing may generally be accomplished directly between a source FPGA and a destination FPGA through the direct duplex serial data communication link between the two FPGAs, *see id.*, 15:24-29, or alternatively if bandwidth on the direct duplex serial data communication link is insufficient, by relaying through intermediate FPGAs, *see id.*, 15:49-59.

III. PERSON OF ORDINARY SKILL IN THE ART


Around November 16, 2006 (the filing date of patent application 11/600,934 as a continuation-in-part of patent application 10/843,226), 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. Ex. 2007, ¶28.

IV. ASSERTED GROUNDS

A. Ground 1: Hartsell does not render obvious claims 27-28, 31, or 33

In Ground 1, Petitioner alleges that claims 27, 28, 31, and 33 are rendered obvious by Exhibit 1005 (“Hartsell”). Petitioner has failed to meet its burden. Each of independent claims 27 and 31 recites *each 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). Claim 28 depends from claim 27, and claim 33 depends from claim 31. As explained below, Petitioner has not shown that Hartsell discloses or renders obvious this limitation.

Petitioner identifies Hartsell's content delivery system 1010 as a *communications infrastructure* and Hartsell's transport processing engines 1050 as signal processing circuits. Pet., 16-18. Alternatively, Petitioner alleges that the combination of Hartsell's network interface processing engine 1030 and transport processing engine 1050 renders obvious a *signal processing circuit*. Pet., 20.



Ex. 1005, Fig. 1A

Petitioner further alleges that Hartsell's processor modules 1001 and 1003, configured to operate as transport processing engines are *ASIC devices*. Pet., 35. Petitioner acknowledges that Hartsell's processor modules 1001 and 1003 are not themselves FPGA's but argues that a POSITA would have recognized that they could have been implemented using a "single integrated chip FPGA device." Pet. 27, 32-33. As explained below, Hartsell's content delivery system does not disclose

or render obvious *each of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router*, so regardless of whether the claimed *FPGA device* would be understood as a single integrated chip FPGA or a circuit that includes an FPGA, Petitioner's arguments in Ground 1 fail.

1. Processing packets and communicating packets are not packet routing

Petitioner alleges that Hartsell's processor modules 1001 and 1003 "*each itself includes a packet router* because each processor controls packet routing." Pet. 35. As evidence of this "packet routing" allegedly performed by each of Hartsell's processor modules, Petitioner argues that modules 1001 and 1003 in Hartsell "are used to process packets and communicate packets to other processor modules 1001 and 1003." Pet. 36. Petitioner further contends that "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 retransmission, data queuing and flow control, packet header generation, etc.'." Pet., 36. Tellingly, neither Petitioner nor its expert take the position that any of these identified functions is packet routing. Pet. 36-37; Ex. 1003, ¶93. As Dr. Cole, confirms, they are not: "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.” Ex. 2007, ¶47; *see also id.*, ¶¶48-61.

Petitioner then argues that “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.” Pet., 37. Although Petitioner never explains what these “routing functions” are, Dr. Cole’s analysis confirms that processor modules 1001 and 1003 **do not** perform the functions of a packet router. Ex. 2007, ¶62. Petitioner appears to equate the “transfer of packets” to packet routing, which is not how a POSITA would understand it. Ex. 2007, ¶¶48-54, 62.

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. Ex. 2007, ¶48; *see, e.g.*, Ex. 2009, *Introduction to Networking*, at 436 (“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*, at 288 (“[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.”).

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. Ex. 2007, ¶49. 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.

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

Ex. 1001, Fig. 10 (showing PRISM multiplex matrix module combining $x+1$ multiplex connect modules for the “writer” output interfaces, 832a to 832n).

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

Ex. 1001, 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; Ex. 2007, ¶¶49-51.




FIG. 12

Ex. 1001, Fig. 12 (showing multiple mgt_connect modules for duplex links to “Other FPGAs and/or External Interfaces”); *see also id.*, 18:53-19:3.

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 a intermediate second FPGA and then to the destination FPGA through the PRISM router/s of one or more other FPGAs),” *id.*, 15:49-59.

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.

Based on the ordinary meaning of a *packet router*, a POSITA would readily recognize that facilitating the transfer of packets between engines by Hartsell's processor modules 1001 and 1003 is not packet routing nor does it render obvious a packet router in processor modules 1001 and 1003. Ex. 2007, ¶¶55-56.

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. Ex. 2007, ¶¶52-54; *see, e.g., id.*, 14:41-48, 19:4-14, 22:63-23:1. Petitioner alleges that "[t]he '323 Patent explains that any routing scheme, including TCP/IP and Infiniband, would involve a packet router," and argues that Hartsell's disclosure of "Infiniband" means that "processor modules 1001 or 1003 perform[] routing of Infiniband packets." Petition at 40, Ex. 1003, ¶97. This is incorrect. Dr. Cole explains that "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." Accordingly, disclosure by Hartsell of the Infiniband protocol does not show that processor modules 1001 and 1003 perform **any** routing functions or render obvious a *packet router* in each processor module, or that Hartsell's processor modules "would have operated as a packet router in the same way as the FPGAs described in the '323 Patent." Pet., 36; Ex. 2007, ¶¶54-55. In addition, as detailed further below, the topology of Hartsell's content delivery system 1010, and

specifically the topology of transport processing engine 1050 and processor modules 1001 and 1003 with respect to interconnection 1080 is not consistent with Hartsell's processor modules operating as packet routers.


Petitioner's contention that transport processing engine 1050 performs packet routing functions is flawed for additional reasons as well. As Dr. Cole explains, "Hartsell does not describe the transport processing engine 1050 as a 'packet router.'" Ex. 2007, ¶56. In fact, Hartsell expressly discloses limiting the network processing requirements of transport engine 1050. *Id.*; *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. Ex. 2007, ¶56. In fact, Hartsell’s disclosure of how transport processing engine functions confirms that it does not perform the functions of a packet router.

As to “time out checks for each network connection,” Dr. Cole explains that “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).” Ex. 2007, ¶57; *see, e.g.*, Ex. 2011, *Networking Fundamentals*, at 54; Ex. 2010, *Computer Networks, Architecture, Protocols, and Software*, at 26. As to “session management,” Dr. Cole confirms that “it is the session layer, the fifth layer in the OSI model, that is responsible for managing sessions.” Ex. 2007, ¶57; *see, e.g.*, Ex. 2011, 54; Ex. 2010, 26. And “data reordering and retransmission, data queuing and flow control,” it is the transport layer, the fourth layer in the OSI model (the transport layer) that is responsible for managing TCP and UDP connections. Ex. 2007, 57; *see, e.g.*, Ex. 2011, 54; Ex. 2010, 26.

As Dr. Cole explains, “[p]acket header generation’ . . . does not have any effect on the path out of the transport processing engine 1050 and is therefore not a packet routing function.” Ex. 2007, 58. The transport processing engine 1050, which has only a single connection to the distributive interconnection 1080, does not make a routing, or path, determination for packets transmitted from the transport

processing engine 1050. *Id.* Furthermore, the FPGA on each of the transport processing engines 1050 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. 2007, ¶58.



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

More specifically, 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.

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.

Ex. 2007, ¶¶59-60.

2. The topology of Hartsell’s content delivery system is inconsistent with each processor module including a *packet router*


Petitioner’s contention that modules 1001 and 1003 would be configured to support “routing functions” is entirely unsupported by Hartsell. Indeed, Hartsell illustrates just the opposite. Hartsell’s content delivery system selectively transfers

packets between engines of the content delivery system (i.e., determines the path of those packets to their destination), as Hartsell plainly discloses, using distributive interconnection 1080, which is entirely separate from processor modules 1001/1003 and engines 1030 and 1050. Hartsell discloses that distributive interconnection 1080 is “any multi-node I/O interconnection hardware or hardware/software system suitable for distributing functionality by selectively interconnecting two or more content delivery engines of a content delivery system.” Ex. 1005, [0058] (emphasis added). Petitioner’s analysis ignores that it is actually Hartsell’s distributive interconnect 1080 that is responsible for switching packets from one engine/module to another. Ex. 2007, ¶60; 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.”); *see also* Ex. 2007, ¶74-76. This functionality of interconnect 1080, which is entirely separate from processing modules 1001 and 1003, is further evidence that these processor modules do not include or render obvious a *packet router*.

Petitioner makes an alternative argument regarding network interface processing engine 1030 that also fails. Petitioner says “the incorporation of the network interface processing engine 1030 into the transport processing engine 1050 further renders obvious the claimed *packet router*.” Pet., 39. This is again

unsupported by Hartsell's disclosure.

As an initial matter, Hartsell does **not** describe incorporating network interface processing engine 1030 **into** transport processing engine 1050. Instead, Hartsell says “the transport engine 1050 and the network interface processing engine 1030 may be combined into a single engine.” EX.1005, [0084]. This combination is illustrated, for example, in Figure 2.



Ex. 1005, Fig. 2. Hartsell explains that “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.” Ex. 1005, [0162]. As Dr. Cole explains, “[w]ith 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.’” Ex. 2007, ¶65 (citing Ex. 1005, [0162]). Thus, this embodiment, “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,” which is no different than how the transport processing engine works in the embodiment where engines 1030 and 1050 are separated. Ex. 2007, ¶66.

In addition, other than noting that “interface engine 1030 may process[] all or some of the TCP/IP protocol stack,” Ex.1005, [0075], Petitioner offers no explanation as to how the combination transport engine 1050 and network interface processing engine 1030 would exhibit any different behavior that would qualify the combination as a packet router. And, contrary to Petitioner’s assertion, Hartsell’s disclosure that network interface processing engine 1030 processes “all or some of the TCP/IP protocol stack” and Petitioner’s assertion that network processing engine 1030 “receives TCP/IP packets and strips their header information, replacing it with an identifier or tag” does not demonstrate to a POSITA that Hartsell’s network interface processing engine includes or renders obvious a *packet router*. Ex. 2007, ¶63. Indeed, each of these steps is consistent with the “conventional operation” of a

network interface, which, as Dr. Cole describes, does not indicate that network interface processing engine 1030 is a *packet router*. Ex. 2007, ¶63.

And, once again, Petitioner offers no explanation as to how the topology of Hartsell's content delivery system 1010 could even be compatible with including a *packet router* within 1030 (let alone each processor module 1001 or 1003), since those modules have a single output to distributive interconnection 1080. Ex. 2007, ¶58. Furthermore, even if Petitioner had shown that Hartsell's network processing engine 1030 performs steps associated with packet routing, which it hasn't, this would not demonstrate that Hartsell discloses or renders obvious multiple *ASIC devices that each itself includes a packet router*. Nor has Petitioner shown why a device such as Hartsell's content delivery system would need a packet router within each processing module 1001/1003 of network processing engine 1030 or transport processing engine 1050.

Because Petitioner has not shown the Hartsell discloses or renders obvious *each one of said two or more signal processing circuits including multiple ASIC devices that each itself includes a packet router*, Petitioner also has not and cannot show that Hartsell discloses or renders obvious the elements of claim limitation [31.5], which recites, in part, "*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.”

Because claim 28 depends from claim 27 and claim 33 depends from claim 31, Petitioner’s argument as to each of these dependent claims fails for the same reasons as articulated with respect to independent claims 27 and 31.

B. Ground 2: The combination Hartsell and Hauck does not render obvious claims 27-28, 31, or 33

In Ground 2, Petitioner introduces Exhibit 1030 (“Hauck”) for the limited purpose of refuting an argument that Patent Owner has not made, i.e., “to the extent that Patent Owner argues that the claims’ recitation of a ‘*packet router*’ within each ASIC device implies that each ASIC device has point-to-point interconnections to other ASIC devices.” As an initial matter, Ground 2 should be ignored because, by Petitioner’s express admission, its argument in Ground 2 only applies “to the extent” that Patent Owner proffers a particular argument as to the scope of the challenged claims. Patent Owner has made no such contention.

1. No motivation to combine Hartsell with Hauck

Petitioner’s suggested motivation to combine Hartsell and Hauck is flawed. Petitioner argues that “[i]t would have been obvious to a POSITA to modify Hartsell’s architecture to eliminate the need for switch fabric 1080 by providing more direct communications paths between pairs of processing modules 1001 and

1003.” Pet., 55. This is factually incorrect and therefore would not motivate a POSITA to modify Hartsell as Petitioner suggests. As Dr. Cole explains, Petitioner’s combination of Hartsell and Hauck “would provide less direct communication paths.” Ex. 2007, ¶73-74. Petitioner’s suggestion to replace Hartsell’s switch fabric with a mesh or crossbar interconnection architecture ignores that Hartsell already discloses the use of a crossbar interconnect switch fabric. Ex. 2007, ¶74; Ex. 1005, [0062]. Hartsell teaches that using the crossbar switch fabric for the distributive interconnect 1080 allows the engines/modules of the content delivery system to be directly connected to the other engines/modules of the content delivery system. *See* Ex. 1005, [0060] (“In the case of a switch fabric, latencies are fixed. Further, the connections are ‘direct’, rather than by some undetermined path.”). Petitioner asserts that a POSITA, looking at the teachings of Hartsell, would be motivated to undo the essential nature of Hartsell by replacing the direct connection crossbar switch fabric with the Springbok interconnection scheme disclosed in Hauck, resulting in a limited number of direct connections between engines/modules of the content delivery system. A POSITA would not be so motivated. *See* Ex. 2007, 75-77; *see also* *Trivascular, Inc. v. Samuels*, 812 F.3d 1056, 1068 (Fed. Cir. 2016) (finding no motivation to modify the prior art where doing so “would destroy the basic objective” of the prior art); *In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984) (finding no reason to modify a prior art device where the modification would render the

device “inoperable for its intended purpose”); *Google LLC f/k/a Google Inc. v. Singular Computing LLC*, IPR2021-00155, Paper 62, at 74 (PTAB May 23, 2022) (“Petitioner’s asserted combination ... would ... defeat[] MacMillan’s stated objective of achieving a highly parallel SIMD computer architecture at lower system cost”) (internal quotations omitted).

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. Ex. 2007, ¶74; *see* Ex. 2012, The Crosspoint-Queued Switch, Kanizo, Y., Hay, D., Keslassy, I., *IEEE Infocom 2009*, pp. 729 (“[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...”).

As Dr. Cole explains, “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.” Ex. 2007, ¶75. Furthermore, Hartsell discloses that the IBM PRIZMA-E is available with either 8 or 16 ports. Ex. 2007, ¶76; Ex. 1005, [0062]. If, as Petitioner suggests, Hartsell’s switch fabric were replaced with the Springbok interconnect architecture shown in Hauck’s Figure 25, “any FLGA 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.” Ex. 2007, ¶77. This decrease in direct communication capability would not motivate a POSITA to modify Hartsell using Hauck and would more likely deter such modification.

2. The Hartsell/Hauck combination would not render any challenged claim obvious

The modifications to Hartsell that Petitioner proposes in light of Hauck do not render any challenged claim obvious. First, the proposed combination does not include or suggest a *packet router*, which is recited in each of independent claims 27 and 31. Petitioner alleges that a POSITA would have “implement[ed] 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.” Pet., 56. Petitioner bases this allegation on its contention that “Hauck confirms the practice of directly connecting FPGA devices to route packets between FPGAs was used ‘for both routing and rerouting of signals[.]’ *Id.* (emphasis added). But this is incorrect. Hauck does not disclose routing packets. Ex. 2007, ¶68-70. Hauck instead describes a system of hardwired

connections between FPGAs on separate daughter cards, which Hauck describes as a “Springbok system.” As Dr. Cole explains, “Hauck does not teach that the FPGAs in the Springbok system are used to implement any particular form of computer or network functionality . . . and does not teach that the FPGA’s in the Springbok system are used to implement a “packet router” as the Petition suggests. Ex. 2007, ¶71. As discussed in Ground 1, the ’323 Patent discloses a *packet router* according to the plain meaning of the term, which performs packet routing functionality such as receiving packets on an input, examining 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 forwarding 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. Ex. 2007, ¶72; *see, e.g.*, Ex. 1001, 15:24-29, 15:49-59. Hauck’s hardwired interconnections alone do not make a *packet router*.

In addition, Petitioner’s contention that the modifications it proposes to Hartsell “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)" is unsupported. Pet., 58. It is not enough for Petitioner to simply state that Hartsell's system is "ready for improvement" without articulating any opportunity for improvement to Hartsell's system that would have been recognized by a POSITA. As explained above, Hauck's topology of connecting FPGAs results in a decrease in connectivity compared to Hartsell's express disclosures. And neither Petitioner nor its expert identify any advantages that would be achieved by the modifications it proposes.

Petitioner further contends that Hartsell's processor modules 1001 and 1003 could be implemented as single integrated chip FPGA devices that a POSITA would have recognized would be directly interconnected. And Petitioner identifies the Xilinx Vertex-II Pro FPGAs as exemplary and asserts that these FPGAs were "known to have port modules capable of routing communications to other FPGAs." Pet., 58. But the ability of FPGAs to communicate to other FPGAs does not change the fact that neither Hartsell nor Hauck teaches or renders obvious including a *packet router* within such an FPGA, let alone multiple FPGAs that each include a *packet router*.

Petitioner presents no new arguments in Ground 2 for claims 28 or 33, which fails as to these claims for all the reasons previously provided in Ground 1 and

because of their dependence on claims 27 and 32, respectively.

3. Petitioner has not shown that Hauck is a prior art printed publication

Petitioner bears the burden of demonstrating that each reference it relies in qualifies as a “patent or printed publication.” In the Petition, Petitioner merely asserts that Hauck “was published April 1998 in Proceedings of the IEEE, vol. 86, no. 4” but provides no evidence or argument as to why this assertion is supported. Pet., 14. As the Board is aware, “the petition must identify, with particularity, evidence sufficient to establish a reasonable likelihood that the reference was publicly accessible before the critical date of the challenged patent and therefore that there is a reasonable likelihood that it qualifies as a printed publication.” *Hulu, LLC v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29 at 11 (PTAB Dec. 20, 2019) (precedential); *Nobel Biocare Servs. AG v. Intradent USA, Inc.*, 903 F.3d 1365, 1375 (Fed. Cir. 2018) (citing *Medtronic, Inc. v. Barry*, 891 F.3d 1368, 1380 (Fed. Cir. 2018)); *see also In re Wyer*, 655 F.2d 221, 227 (C.C.P.A. 1981) (a party asserting a reference as a prior art printed publication should provide sufficient proof of accessibility). Petitioner has not done so here. And, although *Hulu* recognizes that Petitioners have “limited opportunities” to present additional evidence after an institution decision, that does not relieve Petitioner of its burden to present sufficient evidence in its Petition. Merely placing Patent Owner on notice that it contends a

reference qualifies as a printed publication is not sufficient.

C. Ground 3: The combination Hartsell and Raj does not render obvious claims 27-28, 31, or 33

In Ground 3, Petitioner relies entirely on its prior arguments for limitations [27.2] and [31.2] presented in Grounds 1 and 2. Pet., 65 For the same reasons as articulated above as to these limitations, Ground 3 fails to render obvious this limitation. Because Ground 3 fails to render obvious this limitation, it also fails to render obvious the *packet router* recited in limitations [27.3], [31.3], and [31.5]. Because claim 28 depends from claim 27 and claim 33 depends from claim 31, Petitioner's argument as to each of these dependent claims fails for the same reasons as articulated with respect to independent claims 27 and 31 as to Ground 1.

D. Ground 4: The combination Hartsell, Hauck, and Raj does not render obvious claims 27-28, 31, or 33

In Ground 4, Petitioner relies entirely on its prior arguments for limitations [27.2] and [31.2] presented in Grounds 1 and 2. Pet., 65. For the same reasons as articulated above as to these limitations, Ground 4 fails to render obvious this limitation. Because Ground 3 fails to render obvious this limitation, it also fails to render obvious the *packet router* recited in limitations [27.3], [31.3], and [31.5]. Because claim 28 depends from claim 27 and claim 33 depends from claim 31, Petitioner's argument as to each of these dependent claims fails for the same reasons as articulated with respect to independent claims 27 and 31 as to Grounds 1 and 2.

V. CONCLUSION

Petitioner has not established that any of the challenged claims are unpatentable. Proper analysis of the references relied upon in fact confirms patentability of the challenged claims.

Date: April 17, 2024

Respectfully submitted,

/Brett Cooper/
Brett Cooper, Reg. No. 55,085
Lead counsel for Patent Owner

CERTIFICATE OF WORD COUNT (37 C.F.R. § 42.24(d))

The undersigned hereby certifies that the word count in this Patent Owner Response is 6,052 words, as counted by the word-processing program Microsoft Word that was used to generate this Patent Owner Response, where such word count excludes the table of contents, table of authorities, mandatory notices, certificate of service, appendix of exhibits, claim listing, and this certificate of word count. This Patent Owner Response is in compliance with the 14,000-word limit set forth in 37 C.F.R. § 42.24(b)(1).

Date: April 17, 2024

/Brett Cooper/

BC Law Group, P.C.
200 Madison Avenue, 24th Floor
New York, NY 10016
(212) 951-0100

Brett Cooper
Reg. No. 55,085
Attorney for Patent Owner

CERTIFICATE OF SERVICE (37 C.F.R. § 42.6(e))

The undersigned hereby certifies that the above document was served on April 17, 2024, by filing this document through the Patent Trial and Appeal Case Tracking System as well as delivering a copy via electronic mail upon the following attorneys of record for the Petitioner:

Theodore M. Foster
ipr.theo.foster@haynesboone.com
David L. McCombs
David.mccombs.ipr@haynesboone.com
Brian Kwok
Brian.kwok.ipr@haynesboone.com
Calmann Clements
calmann.clements.ipr@haynesboone.com
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, TX 75201

Date: April 17, 2024

/Brett Cooper/

BC Law Group, P.C.
200 Madison Avenue, 24th Floor
New York, NY 10016
(212) 951-0100

Brett Cooper, Reg. No. 55,085
Attorney for Patent Owner