Trials@uspto.gov
571-272-7822

Paper 37
Entered: October 11, 2023

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CEREBRAS SYSTEMS INC.,
Petitioner,

v.

REX COMPUTING, INC,
Patent Owner.

IPR2022-00741
Patent 10,355,975 B2

Before MICHAEL R. ZECHER, MICHAEL T. CYGAN, and JULIET MITCHELL DIRBA, *Administrative Patent Judges*.

CYGAN, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision
Determining No Challenged Claims Unpatentable
35 U.S.C. § 318(a)

I. BACKGROUND

A. Background and Summary

Cerebras Systems Inc. ("Petitioner") filed a Petition requesting *inter partes* review ("IPR") of claims 1–19 of U.S. Patent No. 10,355,975 B2 (Ex. 1001, "the '975 patent"). Paper 1 ("Pet."). The Petition was supported by a declaration by Dr. Thomas Conte. Ex. 1002 ("Conte Declaration"). Rex Computing, Inc. ("Patent Owner") filed a Preliminary Response. Paper 7. On October 14, 2022, we instituted trial on all claims and all of the asserted grounds of unpatentability. Paper 8 ("Dec. on Inst."). Patent Owner filed a request for rehearing, which we denied. Papers 10, 11. Patent Owner filed a Response, supported by a declaration by Dr. Nader Bagherzadeh, and accompanied by a transcript of a deposition of Dr. Conte. Paper 19 ("Resp."); Ex. 2001 ("Bagherzadeh Declaration"); Ex. 2002 ("Conte Deposition"). Petitioner filed a Reply, supported by a second declaration by Dr. Conte, and accompanied by a transcript of a deposition of Dr. Bagherzadeh. Paper 25 ("Pet. Reply"); Ex. 1025 ("2nd Conte Declaration"); Ex. 1024 ("Bagherzadeh Deposition"). Patent Owner filed a Sur-reply. Paper 31 ("Sur-reply"). Oral argument was held on July 17, 2023, and a transcript of the oral argument is included in the record. Paper 36 ("Tr.").

We have jurisdiction under 35 U.S.C. § 6. This decision is a Final Written Decision under 35 U.S.C. § 318(a) as to the patentability of claims 1–19 of the '975 patent. For the reasons we identify below, we determine that Petitioner *has not shown* based on a preponderance of evidence that any of claims 1–19 are unpatentable.

B. Real Party in Interest

The parties identify themselves as the real parties in interest. Pet. 1; Paper 6, 1.

C. Related Matters

The parties indicate that the '975 patent is involved in *Rex Computing, Inc. v. Cerebras Systems Inc.*, No. 1:21-cv-00525-MN (D. Del.), filed April 13, 2021. Pet. 2; Paper 6, 2.

D. The '975 Patent

The '975 patent generally relates to architectures for multi-core computer processors. Ex. 1001, 1:19–37. The '975 patent identifies a need for architectures that are "faster, smaller, less power consuming, more reliable, and easier to use." *Id.* at 1:37–39. The '975 patent describes a system in which each processing core corresponds to a different router of a set of routers, and through adjacent routers of the set, communicates to other processing cores. *Id.* at 1:43–53. The system may be arranged as a set of tiles in a grid configuration, with each tile including a processing core and its corresponding router. *Id.* at 1:63–67.

Each router receives data packets from the adjacent routers based on a physical destination address of the data packet. *Id.* at 1:53–57. Each router sends data packets to adjacent routers, or to the processing cores connected thereto, wherein each router is operable to retain a data packet in the event of a traffic condition. *Id.* at 1:57–61. Each router implements a deterministic, i.e., static, routing policy. *Id.* at 1:61–62. Execution of an application may be optimized by an optimization module, which assigns functions of the application to tiles. *Id.* at 15:51–61.

E. Challenged Claims

Petitioner challenges claims 1–19 of the '975 patent. Claims 1 and 13 are independent. Claim 1 is reproduced below (with bracketed identifying labels added):

1. [1pre] A system comprising:

[1a] a set of processor cores;

[1b] a set of routers each including a set of input ports and a set of output ports, wherein:

[1c] each processor core of the set of processor cores corresponds to a different router of the set of routers, wherein each processor core and corresponding router form a tile,

[1d] each processor core is communicatively coupled with a corresponding router via the router's set of input ports and set of output ports,

[1e] each router is communicatively coupled with one or more adjacent routers via the set of input ports and the set of output ports, and wherein each processor core is communicatively coupled with the other processor cores via the set of routers,

[1f] each router is operable to receive one or more data packets from the one or more adjacent routers or the processor core corresponding to the router,

[1g] based on a physical destination address of a data packet, each router is operable to send one or more data packets to the one or more adjacent routers or the processor core corresponding to the router, wherein each router is operable to retain a data packet in the event of a traffic condition, and

[1h] each router implements a static priority routing policy; and

[1i] an optimization module configured to:

[1j] determine optimal function assignment configurations for groups of tiles, and

> [1k] assign two or more functions, which communicate at least unilaterally more frequently with one another than with other functions, to groups of adjacent tiles based on an optimal function assignment configuration determination, wherein the two or more functions are assigned to groups of tiles communicatively coupled in square configurations when the function executes optimally when executed by the groups of tiles communicatively coupled in the square configurations and are assigned to groups of tiles communicatively coupled in linear configurations when the function executes optimally when executed by the groups of tiles communicatively coupled in the linear configurations.

Ex. 1001, 28:24–67.

F. Prior Art Relied Upon

Petitioner relies upon the prior art references listed below. Pet. 4.

U.S. Patent No. 8,531,943 B2 to Olofsson, filed October 29, 2009, and issued on September 10, 2013 (Ex. 1004);

U.S. Patent Application Publication No. 2016/0179728 A1 to Kaul et al. ("Kaul"), filed December 17, 2014, and published on June 23, 2016 (Ex. 1005);

"The Raw Microprocessor: A Computational Fabric for Software Circuits and General-Purpose Programs," Taylor et al., IEEE Micro 25, March-April 2002 ("MIT2002"; Ex. 1006);

"Scalar Operand Networks: On-Chip Interconnect for ILP in Partitioned Architectures," Taylor et al., ("MIT2003"; Ex. 1007);

"Evaluation of the Raw Microprocessor: An Exposed-Wire-Delay Architecture for ILP and Streams," Taylor et al., ("MIT2004"; Ex. 1008);

"A Modified NoC Router Architecture with Fixed Priority Arbiter," Ansari et al. ("Ansari"), International Journal of Science and Research, Vol. 4, n.10, 923, October 2015 (Ex. 1009).

Petitioner provides a showing that these references are prior art patents or printed publications under 35 U.S.C. § 102. *See* Pet. 3–4; Ex. 1005, codes (22), (43); Ex. 1006, codes (22), (45); Ex. 1007, codes (22), (63). With respect to MIT2002, MIT2003, and MIT2004, Petitioner relies upon a declaration of Dr. Sylvia Hall-Ellis to establish the provenance of these references. Pet. 10. Dr. Hall-Ellis states that MIT2002 was "available on March 1, 2002, in the *ACM Digital Library* and on August 7, 2002, in the *IEEE Xplore* database." Ex. 1016 ¶ 51. Dr. Hall-Ellis further states that the IEEE Micro article publication of MIT2002 was publicly available on, or shortly after, April 17, 2002. *Id.* ¶ 53.

With respect to MIT2003, Petitioner asserts that it is referenced in MIT2004; we understand the corresponding citation to be the same-titled paper in "2003 HPCA" at pages 341–353. Pet. 10; Ex. 1008, 12. Dr. Hall-Ellis states that MIT2003 is "a copy of a conference paper published in the volume titled *HPCA-9 2003: the Ninth International Symposium on High-Performance Computer Architecture: Proceedings* published by the IEEE Computer Society with a 2003 copyright date." Ex. 1016 ¶ 58. Dr. Hall-Ellis states that this conference took place on February 8–12, 2003. *Id.* Dr. Hall-Ellis further states that MIT2003 "was added to the *IEEE Xplore* database on February 28, 2003, and to the *ACM Digital Library* on February 8, 2003." *Id.* ¶ 59. Dr. Hall-Ellis further states "that the book titled *HPCA-9 2003: the Ninth International Symposium on High-Performance Computer Architecture: Proceedings*, including the MIT2003 reference, was publicly available in the Linda Hall Library as of March 5, 2003." *Id.* ¶ 61.

With respect to MIT2004, Dr. Hall-Ellis states that it "is a copy of a conference paper published in the volume titled *Proceedings of the 31st*

Annual International Symposium on Computer Architecture published by the IEEE Computer Society with a 2004 copyright date." *Id.* ¶ 65. Dr. Hall-Ellis states that this conference took place on June 19–23, 2004. *Id.* Dr. Hall-Ellis further states that MIT2004 was added to the IEEE Xplore database on July 12, 2004. *Id.* ¶ 66. Dr. Hall-Ellis states "that the book titled *Proceedings of the 31st Annual International Symposium on Computer Architecture* was publicly available in the Linda Hall Library as of July 16, 2004." *Id.* ¶ 68.

For the above described reasons, we determine that Petitioner has shown that the cited references are available as prior art to the '975 patent.

G. Asserted Grounds of Unpatentability

Petitioner challenges the patentability of claims 1–19 of the '975 patent on the following grounds (Pet. 4):

Claim(s) Challenged	35 U.S.C. §[1]	Reference(s)
1, 3, 4, 7–9, 11–13, 15–19	103	MIT2002, MIT2003, MIT2004
1–4, 7–19	103	MIT2002, MIT2003, MIT2004, Ansari
5, 6	103	MIT2002, MIT2003, MIT2004, Kaul
5, 6	103	MIT2002, MIT2003, MIT2004, Ansari, Kaul
1, 3–9, 11–13, 15–19	103	Olofsson, Kaul
1–19	103	Olofsson, Kaul, Ansari

[1] The Leahy-Smith America Invents Act ("AIA"), Pub. L. No. 112-29, 125 Stat. 284, 287–88 (2011), revised 35 U.S.C. §§ 102 and 103, effective March 16, 2013. Because the application from which the '975 patent issued was not filed before this date, and does not claim the benefit of an application that was filed before this date, the post-AIA versions of §§ 102 and 103 apply.

II. DISCUSSION

A. Legal Standards

A claim is unpatentable under § 103 if the differences between the claimed subject matter 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 said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art and (4) when in evidence, objective indicia of non-obviousness (i.e., secondary considerations).[2] *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). We analyze the asserted grounds based on obviousness with these principles in mind.

B. Level of Skill in the Art

Petitioner offers an assessment as to the level of skill in the art as of the time the '975 patent was filed. Pet. 6. Relying upon the Conte Declaration, Petitioner asserts a person of ordinary skill in the art would have had "a Bachelor of Science degree in electrical engineering, computer engineering, or a related subject matter, plus at least two years of professional experience in the field of microprocessor design or a similar field . . . and more experience would compensate for less formal education, and vice versa." *Id.* (citing Ex. 1002 ¶¶ 27–30). Patent Owner applies this

[2] Patent Owner does not identify any secondary considerations to be considered.

assessment in its Response. Resp. 8. Based on the fully developed record, we accept the assessment offered by Petitioner because it is supported by the Conte Declaration and is consistent with the '975 patent and the asserted prior art.

C. Claim Construction

In this proceeding, we apply the same standard that is used to construe a patent claim in a civil action under 35 U.S.C. § 282(b), and accordingly, we construe each claim in accordance with the ordinary and customary meaning of such claim, as would have been understood by one of ordinary skill in the art in the context of the patent, and the prosecution history pertaining to the patent, as articulated in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). 37 C.F.R. § 42.100(b) (2021).

Petitioner submits that no constructions are necessary because the prior art renders the claims obvious under either the Petitioner's or the Patent Owner's proposed litigation constructions. Pet. 6. In particular, Petitioner points to Patent Owner's proposed litigation construction of "static priority routing policy," and "destination address." *Id.* at 6–7. With respect to "optimization module," Petitioner asserts that it is a means-plus-function limitation, but that Patent Owner characterizes it as not a means-plus-function limitation because the "'linear' and 'square' configurations of tiles [are] the definite, structural way to perform the optimization," and that Patent Owner points to corresponding structure in the form of "a multi-core microprocessor having linear and square configurations of tiles." *Id.* at 7 (citing Ex. 1018, 8, 10). Petitioner states that it adopts Patent Owner's construction of this term. *Id.* Petitioner further adopts Patent Owner's construction in the litigation of "optimal" and "optimally" as the comparison

and binary choice between tiles in a square or linear configuration. *Id.* at 8. Petitioner asserts that "data packet" should be given its plain and ordinary meaning of "a unit of data sent over a network." *Id.* (citing Ex. 1002 ¶¶ 65, 157–158).

Patent Owner asserts that, based on intrinsic and extrinsic evidence, "data packet" would have been understood by a person having ordinary skill in the art to mean "a single stage of data that includes a payload and address information used for transmitting the data packet from one location to another in a network." Resp. 9–11. Petitioner agrees, so long as a "single stage" has the meaning "payload data" ascribed to it by Patent Owner's declarant, Dr. Bagherzadeh. Pet. Reply. 1 (citing Ex. 1024, 109:9–19).[3]

"[W]e need only construe terms 'that are in controversy, and only to the extent necessary to resolve the controversy.'" *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng'g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999)). We determine that no constructions are required for the above-disputed terms.

Claim 13 sets forth the term "scratchpad memory." Neither party identifies the claim term "scratchpad memory" as one that requires construction, but the parties differ as to whether the term applies to a

[3] During the hearing, Patent Owner proposed a modified claim construction, which would require data to be "broken up into different stages, transmitted separately, and then . . . reassembled at its destination." Tr. 48:10–12; *see also id.* at 50:2–11. This argument was not presented in Patent Owner's Response (or even its Sur-reply) and, therefore, it is forfeited. *See* Paper 9 (Scheduling Order), 9 ("Patent Owner is cautioned that any arguments not raised in the response may be deemed waived.").

memory in the form of a data or instruction cache, as appearing in the asserted MIT2002 reference. Pet. 32; Resp. 43–44. Consequently, we address the scope and meaning of this term below as it pertains to the parties' arguments.

D. Asserted Obviousness of Claims over MIT2002, MIT2003, and MIT2004 (Ground 1)

Petitioner contends that claims 1, 3, 4, 7–9, 11–13, and 15–19 of the '975 patent are unpatentable as obvious under 35 U.S.C. § 103, over the combined teachings of MIT2002, MIT2003, and MIT2004. Pet. 11. We begin our analysis with a brief overview of MIT2002, MIT2003, and MIT2004.

1. *Overview of MIT2002*

MIT2002 relates to a "scalable instruction set architecture [] to attack the emerging wire-delay problem by providing a parallel, software interface to the gate, wire, and pin resources of the chip." Ex. 1006, 1.[4] The "Raw" processor exhibiting such architecture "divides the usable silicon area [of a chip] into 16 identical, programmable tiles." *Id.* at 2. Each tile contains a processor, static router, two dynamic routers, floating point unit, data cache, and software-managed instruction cache. *Id.* at 2–3. Each tile connects only to its four neighbors, thereby reducing the length of the wires to nearest-neighbor distances, with consequent reduction to wire-delay. *Id.* at 3. The wiring resources of the processor are directly programmed by the software.

[4] We cite to the numbering of the Exhibit provided by the Petitioner in the bottom right-hand corner of each page; any use of the numbering of the journal pages by either party has been converted to the corresponding page number of the Exhibit.

Id. The "Raw" operating system "allows both space and time multiplexing of processes," and can run multiple independent processes simultaneously and "switch them in and out as on a conventional processor." *Id.* at 5. The operating system "allocates a rectangular-shaped number of tiles (corresponding to physical threads that can themselves be virtualized) proportional to the amount of computation that is required by that process." *Id.* In switching, the operating system "finds a contiguous region of tiles that corresponds to the dimension of the process, and resumes the execution of the physical threads." *Id.*

2. *Overview of MIT2003*

MIT2003 relates to scalar operand networks, defined as "interconnects optimized for scalar data transport, whether centralized or distributed," including the algorithms that manage them. Ex. 1007, 1. The described scalar operand networks may be implemented in the "Raw" microprocessor. *Id.* at 2. Scalar operand networks, which can be designed to have short wire lengths, can scale with increasing transistor counts and provide transport for multiple forms of data. *Id.* at 1. Implemented in the "Raw" processor, the scalar operand network "addresses the *bandwidth scalability* challenge by replacing buses with a point-to-point mesh interconnect" that "is programmed to route operands only to those tiles that need them," through the combination of a "static network, an intelligent compiler, and bypass path integrated network interfaces." *Id.* at 8. Further, "[f]or each word sent between tiles on the static network, there is a corresponding instruction in the instruction memory of each router that the word will travel through." *Id.* The static routers "collectively reconfigure

the entire communication pattern of the network on a cycle-by-cycle basis." *Id.*

3. *Overview of MIT2004*

MIT2004 is an evaluation of the "Raw" microprocessor. Ex. 1008, 1. MIT2004 provides an overview of its architecture, its implementation, results of testing, and analysis of the results. *Id.* at 1–11.

4. *Analysis*

Upon review of Petitioner's explanations and supporting evidence, and the arguments presented in Patent Owner's Response and Sur-reply, we are persuaded that Petitioner has not demonstrated, by a preponderance of the evidence, that any of the challenged claims would have been obvious over the combined teachings of MIT2002, MIT2003, and MIT2004. In particular, we determine that Petitioner has not shown limitation [1g] to be taught by this combination, and we determine that Petitioner has not shown the combination teaches limitations [13a] or [13d] of independent claim 13. We begin with claim 1, addressing limitations [1pre]–[1f] because of their relevance to the terms in limitation [1g].

a) *Claim 1*

(1) [1pre]

The preamble of claim [1pre] recites, "[a] system comprising." Petitioner contends this claim language, if limiting, is met by teachings of the combination. Petitioner points to MIT2002's description of a "Raw" microprocessor. Pet. 12 (citing Ex. 1006, 1–2; Ex. 1002 ¶ 49). We agree with Petitioner that the preamble, to the extent limiting, has been shown to be taught by MIT2002.

(2) [1a]

Limitation [1a] recites, "a set of processor cores." Petitioner points to the "Compute resources" of the "Raw" microprocessor described in MIT2002. Pet. 12–13 (citing Ex. 1006, Fig. 1). Petitioner also points to the computer processor in each tile of the processor described in MIT2003. *Id.* at 13 (citing Ex. 1007 § 3; Ex. 1002 ¶ 51). Consequently, we determine that Petitioner has identified teachings in the asserted combination of references for each element of limitation [1a].

(3) [1b], [1c], [1d], [1e]

Limitation [1b] recites, "a set of routers each including a set of input ports and a set of output ports, wherein." Limitation [1c] recites, "each processor core of the set of processor cores corresponds to a different router of the set of routers, wherein each processor core and corresponding router form a tile." Limitation [1d] recites, "each processor core is communicatively coupled with a corresponding router via the router's set of input ports and set of output ports." Limitation [1e] recites, "each router is communicatively coupled with one or more adjacent routers via the set of input ports and the set of output ports, and wherein each processor core is communicatively coupled with the other processor cores via the set of routers."

Petitioner asserts that the "Raw" system describes each tile in the set of 16 tiles having a dynamic router, with each router having input and output ports. Pet. 14 (citing Ex. 1006, Fig. 1; Ex. 1008 § 6). Petitioner provides an annotated version of Figure 1 illustrating this assertion, reproduced below:

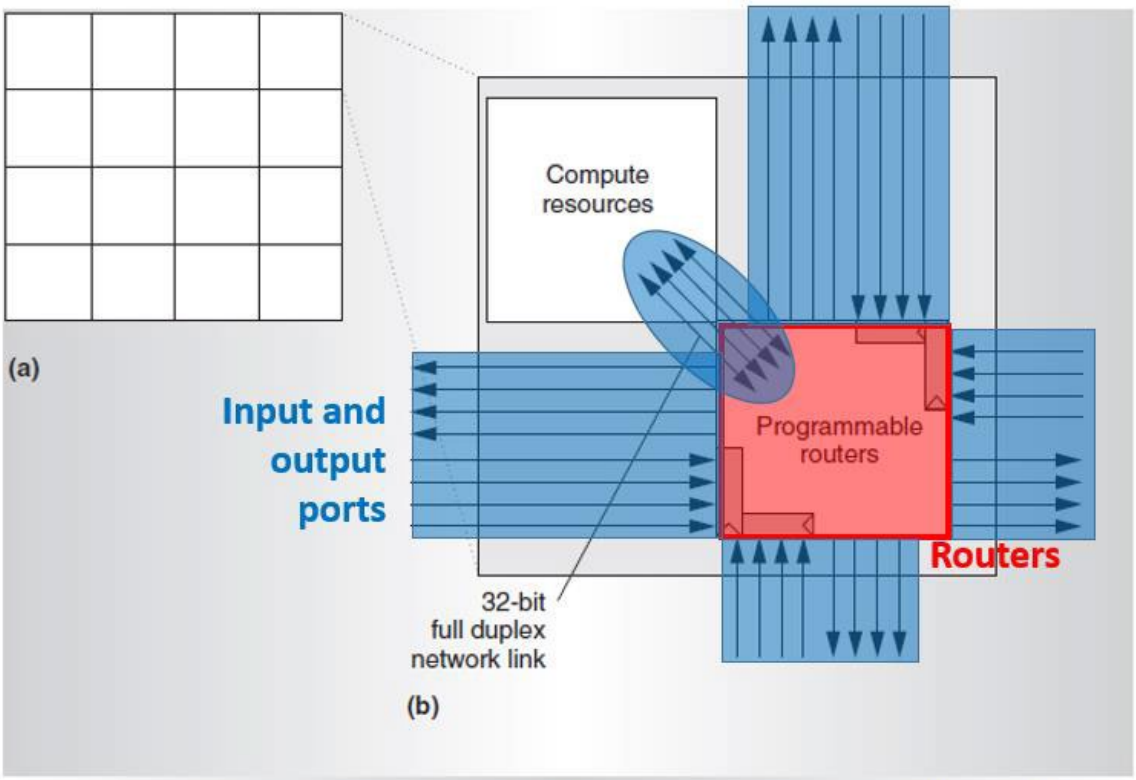

Figure 1. On-chip interconnects in Raw. The Raw microprocessor comprises 16 tiles (a). Each tile (b) has computational resources and four networks, each with eight point-to-point 32-bit buses to neighbor tiles.

Id. Annotated Figure 1 depicts a set of 16 tiles (1a), wherein each tile contains "Compute resources" connected, by arrows described as a "32-bit full duplex network link," to a square box labeled "Programmable routers," wherein the routers have multiple double-headed arrows extending from each side of the square, which Petitioner annotates as "Input and output ports." *Id.*

Patent Owner argues that MIT2002 teaches that each processor core corresponds to multiple routers, and not to a single router so as to form unique core-router pairs. Resp. 13. Patent Owner argues that this results in significant operational differences. *Id.* at 17. Patent Owner annotates Figure 2 of the '975 patent to illustrate the distinction, reproduced below:

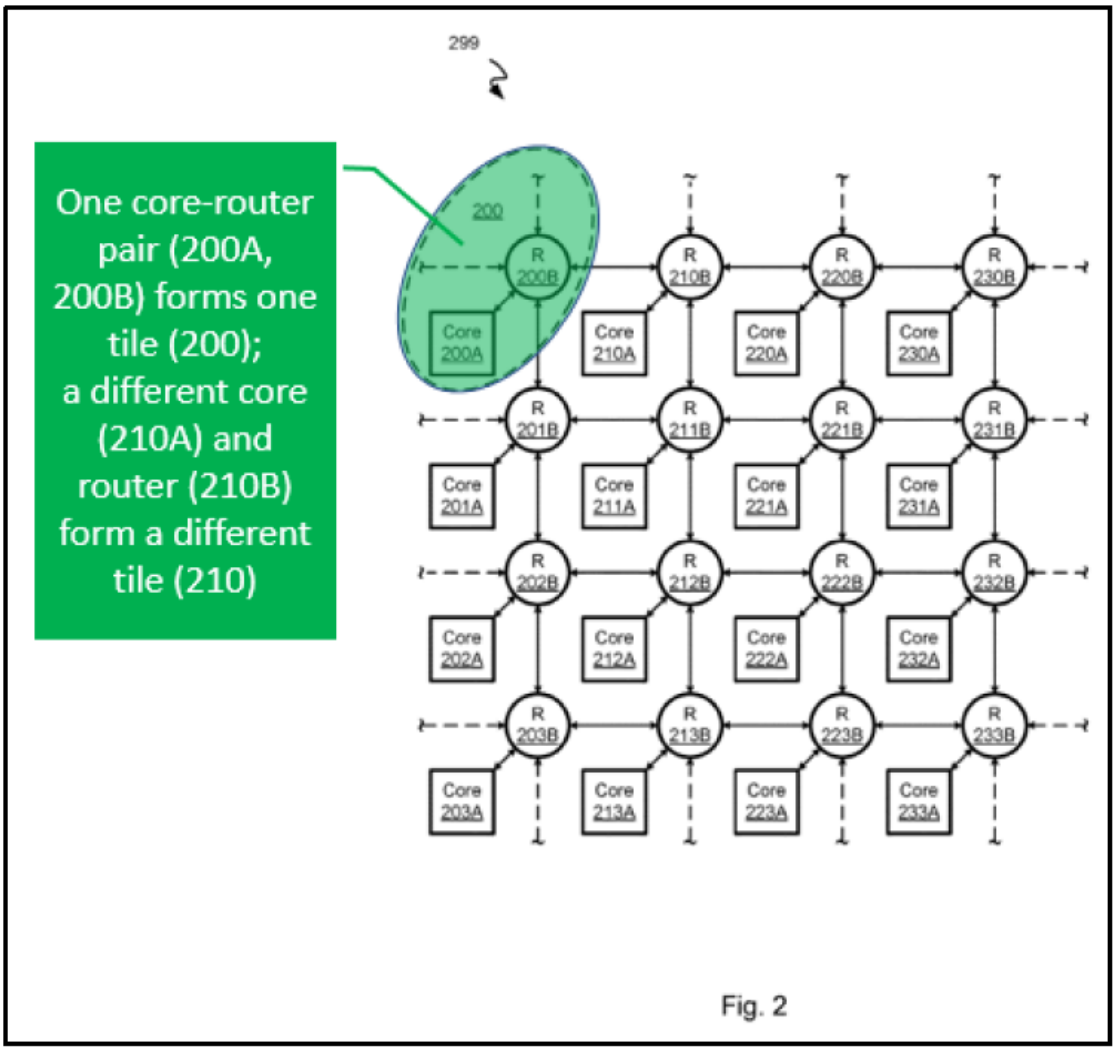

Fig. 2

Id. at 14. Figure 2 illustrates core-router pairs in which each core is connected by a double-headed arrow to a single connected router, and each router is connected by a double-headed arrow to four connected routers that are themselves connected to a core.

We agree with Petitioner that MIT2002 teaches a set of routers ("Programmable routers") each including a set (Figure 1a) of input ports and a set of output ports (indicated by the ingoing and outgoing arrows at each router in Figure 1b), wherein each processor core ("Compute resources") of the set of processor cores corresponds to a different router of the set of

routers, wherein each processor core and corresponding router form a tile (Figure 1b). We further agree that each processor core is communicatively coupled with a corresponding router via the router's set of input ports and set of output ports (double arrows therebetween in Figure 1b), and each router is communicatively coupled with one or more adjacent routers via the set of input ports and the set of output ports (Figure 1b and Figure 1 description). We further agree that each processor core is communicatively coupled with the other processor cores via the set of routers (Figure 1 description), as illustrated in its Figure 1. We do not discern any distinction that would distinguish over limitation [1c]. In both Figure 1 of MIT2002 and in Figure 2 of the '975 patent, each core is directly connected to a corresponding router, and each router is bidirectionally connected to four other routers that are part of core-router pairs. The fact that the MIT2002 system may offer further data paths does not alter the fact that its router-core configurations meet the limitations set forth in [1c]. Consequently, we agree with Petitioner that the limitations of [1b]–[1e] are taught by the applied combination of references.

(4) [1f]

Limitation [1f] recites, "each router is operable to receive one or more data packets from the one or more adjacent routers or the processor core corresponding to the router." Petitioner points to MIT2002's description of dynamic routers using data packets, including a destination address and data payload. Pet. 18–19 (citing Ex. 1006, 8 ("To send a message on one of these networks, the user injects a single header word that specifies the destination tile (or I/O port), a user field, and the length of the message. The user then sends up to 31 data words.")).

In addition, Petitioner asserts that MIT2003 states that the "Raw" system employs a data packet in the form of an "operand message," which consists of "the message header with the addressing information, an operand tag, and the operand itself." *Id.* at 19 (quoting Ex. 1007 § 6). Petitioner also points to the "Raw" system's "low-latency scalar operand network." *Id.* (citing Ex. 1008 § 7). Relying on the declaration testimony of Dr. Conte, Petitioner asserts that the description of the "Raw" system in the three applied references renders obvious the "data packet" limitation. *Id.* (citing Ex. 1002 ¶ 64). Petitioner asserts that the "Raw" system's operand messages are data packets "because they are 'a unit of data sent over a network,' and they also contain fields including a destination address and other data fields," which is "the precise way the '975 patent specification discusses the format of the claimed data packets." *Id.* at 20 (citing Ex. 1002 ¶ 65; Ex. 1001, 4:50–54, 11:20–22, 14:50–52).

Patent Owner argues that MIT2002 teaches operands, not "data packets." Resp. 18; Sur-reply 6–9. Patent Owner argues that operands do not carry payloads, but are merely values. Resp. 18 (citing Ex. 2001 ¶ 60); Sur-reply 7. Patent Owner argues that MIT2002 focuses on maximizing the networks' utility by integrating the network into the bypass paths of the processor pipeline, which indicates a need for maximum efficiency that would have been achieved by using low overhead operands in the bypass networks, and a low latency characteristic which indicates, to a person having ordinary skill in the art, that only operands are transmitted. Resp. 18–19 (citing Ex 1006, 30; Ex. 2001 ¶ 60), 21 (citing Ex. 1002 ¶ 64; Ex. 1006, 31); Sur-reply 9. Patent Owner further argues against Petitioner's assertion that an operand can be part of a data packet, because Petitioner has

not explained why a person having ordinary skill in the art would have had a reason to modify MIT2002's operands to be encapsulated in data packets. Resp. 20 (citing Ex. 1002 ¶ 63); Sur-reply 8–9. Patent Owner further argues that, although MIT2002 teaches "words," which would have been understood to be data that could be structured as at least an operand or a data packet, a person having ordinary skill in the art would have understood them to be operands and not data packets. Resp. 21 (citing Ex. 2001 ¶ 63); Sur-reply 7.

We agree with Petitioner that MIT2003 teaches that the "operand messages" in the Raw system include "the message header with the addressing information, an operand tag, and the operand itself." Pet. 19 (quoting Ex. 1007, 11). We agree that this meets Patent Owner's construction of data packet as "a single stage of data that includes a payload and address information used for transmitting the data packet from one location to another in a network," which we have adopted. *Supra* at II(C). We further agree with Petitioner that MIT2003 describes the properties of the Raw system described in MIT2002. Ex. 1007, 2 (MIT2003 describes "the details of the actual 16-way issue scalar operand network designed and implemented in the Raw microprocessor"). Therefore, a person having ordinary skill in the art reading MIT2003 would have understood that its descriptions of the Raw system would obviously apply to the Raw system as discussed in MIT2002. *See* Ex. 1002 ¶¶ 61–65.

Although Patent Owner argues that an operand, by itself, is a value, not a data packet, Petitioner relies on the entirety of the operand message as teaching the data packet. S*ee* Resp. 18–22; Sur-reply 6–9; Pet. 19–20. Although Dr. Bagherzadeh opines that the combined MIT references teach

"only . . . operand components" rather than payloads (Ex. 2001 ¶ 62), Dr. Bagherzadeh states that "[d]ata payloads include information, data, and instructions." Ex. 2001 ¶¶ 52, 62. Both declarants confirm that an operand is a form of data. Ex. 2002, 61:17–21 ("a data packet may comprise in part an operand . . . since it's data"); Ex. 2001 ¶ 63 ("In my opinion, a [person of ordinary skill in the art] would have understood that a word is a 32-bit data that can be part of an operand or data packet."); Ex. 1024, 110:10–111:3.

Similarly, Patent Owner's argument that operand messages are not "*encapsulated in* data packets" (Sur-reply 9 (emphasis added)) is misplaced. Patent Owner's argument appears to be that operand messages are not encapsulated *within* a data packet, but this fails to address Petitioner's allegations, which contend that the operand messages *are* "data packets." Petitioner has shown that the operand message described by the combined MIT references includes a data payload in the form of an operand that, combined with the addressing information and the operand tag, is a single stage of data that includes a payload and address information used for transmitting across a network, as required by Patent Owner's construction, which we have adopted. Nor does this construction require any encapsulation. Consequently, we are persuaded that the asserted combination of references teaches limitation [1f].

(5) [1g]

Limitation [1g] recites, "based on a physical destination address of a data packet, each router is operable to send one or more data packets to the one or more adjacent routers or the processor core corresponding to the router, wherein each router is operable to retain a data packet in the event of a traffic condition." Petitioner asserts that the claimed "traffic condition"

encompasses both mild congestion (where traffic is delayed) and severe congestion like deadlocks (where traffic cannot proceed unless some action is taken). Pet. Reply 8.

Petitioner asserts that the "Raw" system uses a single header word to specify a destination tile; i.e., a particular physical tile. Pet. 20 (citing Ex. 1002 ¶ 67). Petitioner asserts that the claimed retention of a data packet in the event of a traffic condition is taught by the "Raw" system's retention of data in the event of "network congestion or 'deadlock.'" *Id.* at 20–21 (citing Ex. 1006, 8–9).

Petitioner further points to retention of data by the "Raw" system's buffering system, and how the system detects "if words have waited too long on the input." *Id.* at 21 (citing Ex. 1002 ¶ 68). Petitioner asserts that, "while the counter is incrementing and before a deadlock has been detected, the data is retained—for multiple clock cycles—on the router's input buffer. This is all that is required for the limitation to be met." Pet. Reply 11.

Patent Owner argues that it is the Petitioner's burden to show that it is the routers that are "operable to retain a data packet," and that Petitioner has not sufficiently shown how MIT2002, in combination with MIT2003 and MIT2004, teaches this limitation. Resp. 22; Sur-reply 11–12 ("Petitioner has failed to prove where MIT2002 teaches local retention of a data packet"). Further, Patent Owner argues that deadlock cannot be the recited "traffic condition" because the '975 patent is premised on mitigation of deadlock due to its static priority routing policy. Resp. 24–25. Patent Owner argues that this static routing policy is assigned before data packets are transmitted, and it manages traffic without determining how long a packet waits. *Id.* at 31.

Even if deadlock would be a "traffic condition," Patent Owner argues that the Raw system responds to detection of deadlock by draining the network into the dynamic random access memory ("DRAM"), not holding the data on the routers. *Id.* at 27–28. Moreover, Patent Owner's declarant, Dr. Bagherzadeh, states that "a [person of ordinary skill in the art] would not have understood words waiting at an input for some amount of time less than 'too long' in MIT2002 as teaching the recited 'traffic condition.'" *Id.* at 30 (citing Ex. 2001 ¶ 75).

We agree with Patent Owner that Petitioner has not shown sufficient support for its assertion that the router of MIT's Raw system retains a data packet in the event of a traffic condition. The Petition's entire discussion of that limitation is provided below:

> ***Wherein each router is operable to retain a data packet in the event of a traffic condition:*** MIT2002 describes how data may be retained in the event of network congestion or "deadlock." MIT2002 at 32-33. MIT2002 also describes how the system detects "if words have waited too long on the input." *Id.* at 33. This buffering system "retain[s] a data packet in the event of a traffic condition." [Ex. 1002] ¶ 68.

Pet. 20–21 (citing Ex. 1002 ¶ 68). The most straightforward reading of this discussion is that it is an assertion that network congestion, which is described as "deadlock" in MIT2002, is a traffic condition that occurs when "words have waited too long on the input," which triggers retention of data packets in MIT2002's buffering system. However, such an assertion does not meet the requirement that the "router be operable to retain a data packet," because when MIT2002 detects "if words have waited too long on the input," an interrupt causes the network to be "drained into DRAM." Ex. 1006, 8–9. Because MIT2002's DRAM is not a router as specified by

limitations 1[b]–1[f] (*see* Resp. 27–28; Ex. 2001 ¶ 73), Petitioner's assertions under this straightforward reading are not persuasive.

However, in our Institution Decision, we recognized that the alternative language presented in the Petition (i.e., "MIT2002 also describes") could be read in another manner. We viewed MIT2002's statement that the processor detects "if words have waited too long on the input" as supporting Petitioner's assertions by teaching that data packets are retained on the input of the processor. Dec. on Inst. 19 (citing Ex. 1006, 9;[5] Pet. 21). We noted that MIT2002 describes a processor having input and output passing through a corresponding router. *Id.* (citing Ex. 1006, Fig. 1). We concluded that "the current record appears to support Petitioner's assertion that data packets are retained by a buffer corresponding to the router; i.e., the router associated with the input of that particular processor." *Id.* We clarified, however, that the final written decision would take into account the additional arguments and evidence presented during trial, and we would evaluate Petitioner's position based on the fully developed record using a preponderance of evidence standard. Dec. on Inst. 19; Reh'g Dec. 10–11.

On the full trial record, we agree with Patent Owner that Petitioner has not shown by a preponderance of the evidence that MIT2002 teaches that these buffers are part of the router such that the router retains data. Resp. 22–29; Sur-reply 10–12. The Petition focused on words waiting on the "input." Pet. 21. Petitioner states that, in the Raw system, retention of

[5] We noted that Exhibit 1006 footnotes this quotation with a reference not currently of record. Ex. 1006 n.6.

data "occurs in the *router's* input buffer" (Pet. Reply 10 (emphasis added)), but Petitioner does not support this assertion with analysis or evidence. To the contrary, Petitioner has clarified that the "input" on which words wait is "the input of the processor," not the router. Pet. Reply 12. Claim 1 requires that a router must be operable to retain such data (limitation [1g]), and that the processor cores are distinct from the routers (limitation [1a]). Thus, to the extent that MIT2002 teaches data waiting at the input of the processor, Petitioner has not pointed to evidence or reasoning persuasively showing that the claimed *router* is operable to retain a data packet, even if a buffer at the input of the processor is able to do so.

Instead, Petitioner primarily relies on our preliminary determination, and on Patent Owner's lack of argument to the contrary. *Id.* at 11–12. Neither of these act to satisfy Petitioner's burden to show that MIT2002, in combination with the other applied references, teaches a router operable to retain data packets in the event of a traffic condition. Consequently, we are not persuaded that retention of data in the Raw system occurs at the router's input buffer.

For the above described reasons, we determine that Petitioner has not shown by a preponderance of evidence that the combined teachings of the applied references of this ground teach or suggest the "each router is operable to retain a data packet in the event of a traffic condition" limitation of [1g]. Accordingly, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of claim 1 would have been obvious.

b) *Claim 13*

Petitioner relies upon the same teachings as for claim 1 to account for the limitations of claim 13. Pet. 31–33. Petitioner additionally cites to a

processor containing a data cache and a software-managed instruction cache. *Id.* at 32.

Claim 13 requires, *inter alia*, a router including a memory element operable to store data packets for subsequent sending to a router of an adjacent tile. Ex. 1001, 30:1–5. Petitioner addresses this limitation (denoted [13d]), through reference to its reasoning for limitation [1b] and [1f] of claim 1, and additionally for claim 4. Pet. 33.

We agree with Petitioner that MIT2002 teaches a set of routers each including a set of input ports and a set of output ports, as set forth in limitation [1b], and that each router is operable to receive one or more data packets from the one or more adjacent routers, as set forth in limitation [1f]. *Supra* at II(D)(4)(a)(3)–(4). For claim 4, Petitioner points to the input and output FIFO buffers shown in Figure 4 of MIT2002. Pet. 27 (citing Ex. 1002 ¶ 34). Petitioner asserts that these buffers store data elements in transit among neighboring tiles and routers. *Id.* at 27–28 (citing Ex. 1006, 6 ("For example, a read from register 24 will pull an element from an input FIFO buffer, while a write to register 24 will send the data word out onto that network.")).

However, as for claim 1, Petitioner does not show how MIT2002's FIFO buffers are part of its router. Petitioner merely draws a box around certain elements in Figure 4 and labels that box "Input FIFO buffers," without explaining why those elements are part of the router, despite Figure 4's label describing itself as the "Raw compute processor pipeline." Figure 4 of MIT2002, as annotated by Petitioner, is reproduced below:

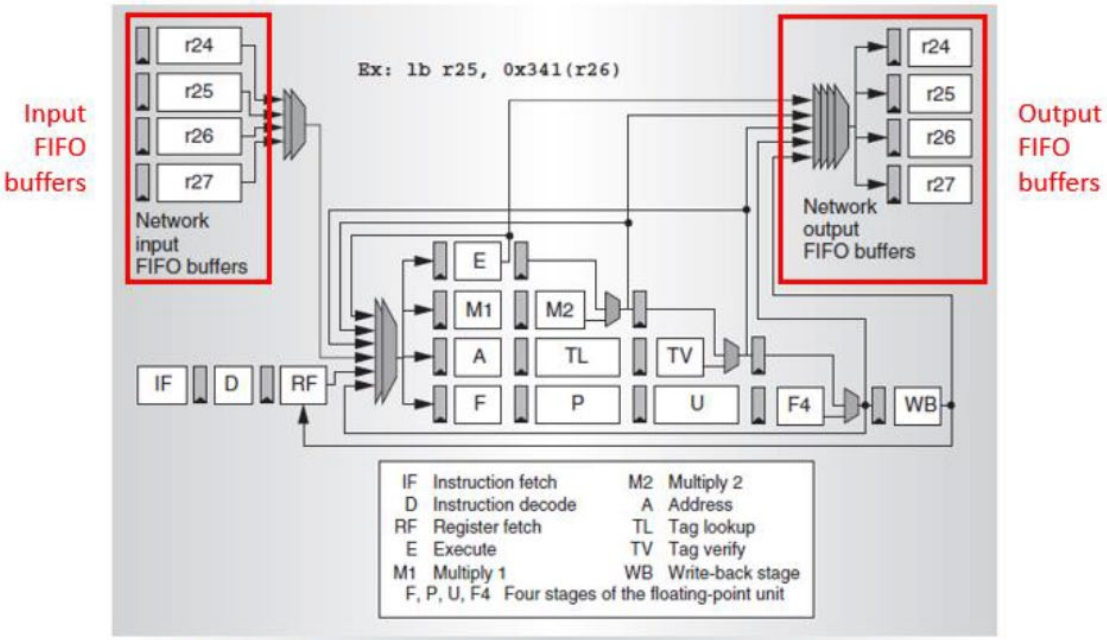

Figure 4. Raw compute processor pipeline.

Figure 4, titled "RAW compute processor pipeline," depicts "Network input FIFO buffers" and "Network output FIFO buffers," each comprising registers r24–r27 connected to different portions of an eight stage processing unit, including four stages of a floating point unit. Ex. 1006, 3, 6–7.

However, in its assertion for limitation [1g] of claim 1, Petitioner presents a contrary characterization, that the "input" on which words wait is "the input of the processor," not the router. Pet. Reply 12. MIT2004 appears to confirm Petitioner's [1g] characterization in its Figure 1, depicted below:

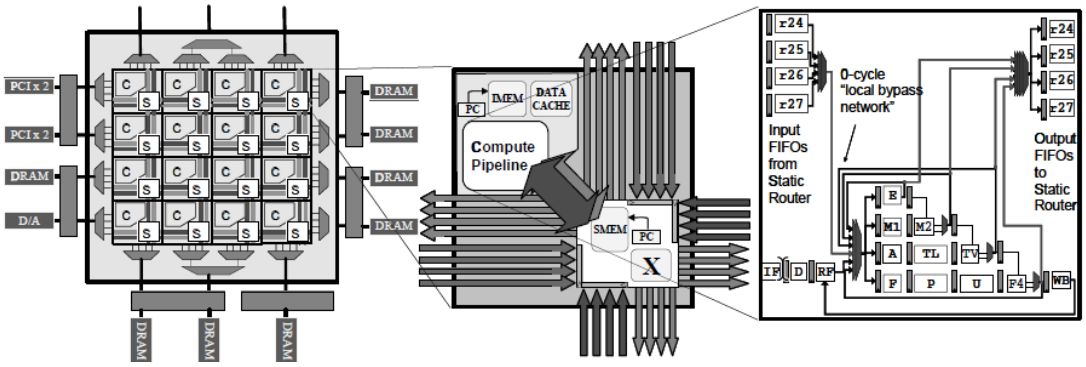

Figure 1 of MIT2004 depicts three images; on the left is an image of an array of tiles; in the middle, a single tile of that array is expanded to show a compute pipeline connected by a double arrow to a router; on the right, that compute pipeline is expanded to show that it contains "Input FIFOs from Static Router," "Output FIFOs to Static Router," each comprising registers r24–r27, connected to different portions of an eight stage processing unit, including four stages of a floating point unit.

Figure 1 of MIT2004 clearly indicates that the "Input FIFOs" and "Output FIFOs" are part of the compute pipeline. The "Input FIFOs" and "Output FIFOs" are "from Static Router" and "to Static Router," respectively. If the FIFOs were part of the router, one would not expect them to be labeled as "to" and "from" the router. This disclosure of MIT2004 is highly pertinent to Petitioner's assertions for claim 1 because Petitioner states that both MIT2002 and MIT2004 "describe the same system." Pet. 12.

Thus, the evidence of record, including Petitioner's own characterizations, show that the FIFO memory element relied upon by Petitioner is part of MIT2002's processor, not MIT2002's router. Consequently, the evidence of record does not support Petitioner's contention that MIT2002 teaches a router having a memory element

operable to store data packets for subsequent sending to a router of an adjacent tile. Accordingly, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of claim 13 would have been obvious.

Petitioner's reliance on the combination of MIT2002, MIT2003, and MIT2004 to teach all the limitations of claim 13 is deficient for an additional reason. That is, we agree with Patent Owner that MIT2002 does not teach the "set of scratchpad modules" set forth in claim 13. Resp. 43–44. Claim 13 recites, in relevant part, that each processor core "is communicatively coupled with a different scratchpad memory module." Ex. 1001, 29:51–58. To teach the claimed scratchpad memory module, Petitioner points to either the data cache or the instruction cache that is contained in each processor of each tile of MIT2002. Pet. 32. Drawing upon Dr. Conte's testimony, Petitioner asserts that both caches have functions of the scratchpad disclosed in the '975 patent; i.e., that MIT2002's instruction cache loads cached processor instructions, and both of MIT2002's caches are "high-speed local memory used for temporary storage of data during application execution." *Id.* (citing Ex. 1001, 6:21–26, 6:57–59; Ex. 1002 ¶ 99). Petitioner further asserts that each processor is communicatively coupled to its respective data cache. *Id.* at 33 (citing Ex. 1006, Fig. 8; Ex. 1002 ¶ 101). At oral argument, Petitioner clarified that the local cache of MIT2002's main memory, not off-chip memory, is asserted as the scratchpad memory module. Tr. 23:19–26.

Patent Owner argues that the disclosure of scratchpad modules in the '975 patent is limited to cacheless embodiments. Resp. 45–46 (citing Ex. 2001 ¶ 92). Patent Owner points to the discussion in the '975 patent of drawbacks of caches. *Id.* at 46 (citing Ex. 1001, 3:60–64, 4:18–21).

Petitioner argues that the '975 patent has no clear disclaimer of the scope of claim 13 that would limit the full scope of the term "scratchpad." Pet. Reply 19–20. Petitioner argues that "[t]here is no clear statement in the specification that the scratchpad cannot be a cache." *Id.*

Patent Owner responds by arguing that "it is undisputed that the '975 Patent defines and explains 'scratchpad memory' as cache-less." PO Sur-reply 16 (citing Ex. 1001, 2:42–44, 5:37–40, Fig. 3). Specifically, Patent Owner states, "the '975 Patent explains on the one hand, scratchpad memory as part of the invention, as differentiated from 'cache' memory in the prior art on the other." *Id.* at 16–17 (comparing Ex. 1001, 2:42–44, 5:37–40, Fig. 3, with Ex. 1001, Fig. 1, 3:21–22, 3:38–4:35).

We determine that Petitioner has not met its burden to show that MIT2002's data cache or instruction cache that is contained in each processor teaches the claimed "scratchpad memory." As mentioned *supra*, we apply the ordinary and customary meaning of this term as it would have been understood by one of ordinary skill in the art in the context of the patent, and the prosecution history pertaining to the patent. Petitioner does not point to the prosecution history or to any source outside the '975 patent to show how the term would have been understood by one of ordinary skill in the art. The dispute between the parties centers on whether Petitioner sufficiently shows that MIT2002's data cache or instruction cache teaches a scratchpad memory, given the context for that term provided in the specification of the '975 patent.

Petitioner relies primarily on statements of Dr. Conte to show that a person having ordinary skill in the art would have construed "scratchpad memory module" as including a cache memory. Pet. 32 (citing Ex. 1002

¶ 99). For his conclusion, Dr. Conte relies solely on the '975 patent specification for its descriptions of the scratchpad memory loading cached processor instructions and being a high speed local memory used for temporary data storage during application execution. Ex. 1002 ¶ 99 (citing Ex. 1001, 6:21–26, 6:57–59).

There is no dispute that the scratchpad memory functions in the above-described manner to store data. But Dr. Conte's direct testimony does not address other statements in the '975 patent specification that describe the scratchpad memory as an alternative to a cache. *See* Ex. 1001, 8:27–33 ("each of the cores includes a scratchpad memory . . . . As a result, conventional L-level caches may not be used in the architecture 299 because the scratchpad memories . . . provide data storage"); *see also id.* at 3:21–4:35 (describing drawbacks of prior art architectures that rely on caches), 5:14–20 (stating that each core may include a scratchpad memory instead of a cache). As Dr. Bagherzadeh stated, "Dr. Conte never reconciled [his] opinion with the '975 Patent's disclosure, which, as I described, connects the concept of using a scratchpad memory to a cache-less architecture." Ex. 2001 ¶ 92 (citing Ex. 1001 Fig. 3). Dr. Conte's deposition testimony did not clarify this issue, stating that the issue of cacheless architecture was "disconnected" from whether a microprocessor has a scratchpad. Ex. 2002, 40:10–15.

The '975 patent specification's description of a scratchpad memory module as an alternative to cache memory is evidence that an ordinary artisan, upon reading the entirety of the '975 patent, would not have considered the scope of the claimed scratchpad memory to include a cache memory. Dr. Conte has relied upon the '975 patent specification for general properties seemingly common to both caches and scratchpad memory, but

has not explained why the '975 patent specification's statements distinguishing a scratchpad memory from a cache should not limit the scope of the claimed "scratchpad memory." As a result, Dr. Conte's testimony does not significantly assist in construing "scratchpad memory" in any manner relevant to the pertinent issue; i.e., whether it may be taught by a prior art cache memory.

We note that Dr. Conte opines, without further explanation, that MIT2002's data and instruction caches teach cacheless memory because they are "software-managed caches." Ex. 2002, 38:12–40:3. Petitioner does not rely upon this theory in its argument, and we are unconvinced that a "cache," even if software-managed, is fairly characterized as "cacheless."

Petitioner's remaining arguments rely on shifting the burden of persuasion away from itself. Petitioner states that the term is "clear on its face," and that "[t]here is no clear statement in the specification that the scratchpad cannot be a cache." Pet. Reply 19–20. However, Petitioner bears the burden to show that a person having ordinary skill would have understood at least one of MIT2002's caches to teach the claimed scratchpad memory. Patent Owner does not bear the burden to show that the '975 patent specification disclaims cache memory instantiations of a scratchpad memory, because Petitioner has not established that a scratchpad memory could, in the view of a person of ordinary skill in the art, be a cache memory.

Consequently, for this additional reason, Petitioner has not shown, by a preponderance of evidence, that claim 13 is obvious over the combination of MIT2002, MIT2003, and MIT2004.

c) *Dependent claims*

By virtue of their dependency, claims 3, 4, 7–9, 11, 12, and 15–19 include the same limitations as at least one of independent claims 1 and 13. Petitioner does not present additional arguments and evidence with respect to these dependent claims that remedy the deficiency in its analysis of the combined teachings of MIT2002, MIT2003, and MIT2004. Pet. 26–31, 34. Accordingly, for the same reasons we identify above, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of these dependent claims would have been obvious.

5. *Conclusion*

Accordingly, we are persuaded that Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of claims 1, 3, 4, 7–9, 11–13, and 15–19 would have been obvious over the combination of MIT2002, MIT2003, and MIT2004.

E. Asserted Obviousness of Claims over MIT2002, MIT2003, MIT2004, and Ansari

Petitioner asserts that claims 1–4 and 7–19 would have been obvious in view of MIT2002, MIT2003, and MIT2004, in view of the teachings of Ansari for limitation [1h]. Pet. 35–36. Ansari is directed to a network on a chip router architecture having an arbiter that detects buffer status of input ports and changes the priorities of the input ports dynamically to enhance router performance. Ex. 1009, 1. An arbiter is used when a number of inputs are requested for the same output port, trapping the source and destination address from the buffer output and generating a control signal so that input data is sent to the output port. *Id.* at 3. Ansari describes a fixed priority arbiter, in which grant is given to the active requester with the

highest priority. *Id.* Ansari describes the fixed priority arbiter as being very common when choosing between just a few requesters. *Id.*

Petitioner adds "Ansari's teaching of static priority routing policies" for limitations [1h] and [13f] to its assertions for Ground 1. Pet. 35–38. However, Petitioner's assertions for the Ansari combination do not provide a teaching of a "router operable to retain a data packet," as found lacking in its assertions for claim 1. Nor do Petitioner's assertions provide a teaching of a router including a "memory element operable to store a data packet" or "scratchpad module" as found lacking in its assertions for claim 13. Thus, the addition of Ansari's teachings to the Ground 1 assertions do not remedy the deficiency of Petitioner's obviousness position regarding the combined teachings of MIT2002, MIT2003, and MIT2004. Accordingly, for the same reasons we identify above, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of these claims would have been obvious.

F. Asserted Obviousness of Claims 5 and 6 over MIT2002, MIT2003, MIT2004, and Kaul

For claims 5 and 6, Petitioner "adds Kaul for specific types of memory hardware" to its assertions for Ground 1. Pet. 39. We begin with an overview of Kaul.

Kaul is directed to a high bandwidth core to network-on-chip interface. Ex. 1005 ¶ 1. Each physical processor includes two cores, and each core include two hardware threads, so as to permit each processor to execute four software threads concurrently. *Id.* at ¶¶ 17–18. An allocator and renamer block provides for reserving resources, such as reorder buffers, to support out-of-order execution. *Id.* at ¶ 21.

Claim 5 adds the further limitation that each FIFO memory element is a single stage memory element. Petitioner asserts that Kaul describes latches; i.e., single stage FIFO memory elements in routers that store data and send packet data to adjacent routers. Pet. 39 (citing Ex. 1005 ¶¶ 40, 42). Petitioner asserts that Kaul's latches would be used to provide temporary storage at each input and output port. *Id.* at 39–40. Petitioner asserts that Kaul's FIFO latches would implement the temporary storage required by MIT2002. *Id.* at 40. Petitioner, relying on Dr. Conte's testimony, asserts that a person having ordinary skill in the art would have had a reasonable expectation of success adding FIFO latches to MIT2002's system because Kaul's latches are standard circuit parts. Petitioner further points to Kaul's latches as teaching the "D flip-flop" limitation of claim 6. *Id.* at 40–41.

However, Petitioner's assertions for the Kaul combination do not provide a teaching of a "router operable to retain a data packet in a traffic condition," as found lacking in its assertions for claim 1. Although Petitioner shows that an additional data storage element may be added to a router, Petitioner does not contend that the additional element alters its analysis for claim 1. That is, Petitioner's claim 1 analysis relies upon "words waiting on the input," which we have determined to be an assertion that data waits on buffers of the processor. *Supra* § II(D)(4)(a)(5). Petitioner's claim 4 analysis of the limitation, "each router uses a set of FIFO memory elements" also relies upon that teaching. Pet. 27 (pointing to the FIFO buffers of Figure 4, which the record shows to be the buffers of a processor, not a router); *Supra* § II(D)(4)(c).

Petitioner's assertion that a buffer may be added to a router does not appear to disturb its claim 1/claim 4 analysis because Petitioner does not

assert that the additional data storage element would, in the combination, be the element upon which words would wait, or how such combination would result in such storage in the event of a traffic condition. If the combination merely replaces Kaul's single stage FIFO for MIT2002's input and output FIFO buffers, these remain on the processor, and do not teach that the router retains a data packet in the event of a traffic condition. If the combination adds Kaul's router, containing single stage FIFO, Petitioner has not satisfactorily explained why the FIFO buffers on MIT2002's processor that are used to store data in the event of a traffic condition would be replaced by the FIFO elements on Kaul's router. Petitioner's reason for adding Kaul's router storage FIFO is to implement the temporary storage system for transferring packets from node to node. Pet. 40. However, MIT2002's processor-based FIFO buffers appear to perform that function, and Petitioner does not explain why that task should instead be given to a newly added set of router-based FIFO storage. Thus, the addition of Kaul's teachings to the Ground 1 assertions do not remedy the deficiency of Petitioner's obviousness position regarding the combined teachings of MIT2002, MIT2003, and MIT2004 for claim 1, from which claims 5 and 6 ultimately depend. Accordingly, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of these dependent claims 5 and 6 would have been obvious.

G. Asserted Obviousness of Claims 5 and 6 over MIT2002, MIT2003, MIT2004, Ansari, and Kaul

For claims 5 and 6, Petitioner adds teachings form both Kaul and Ansari to its assertions for Ground 1. Pet. 41–42. Thus, the addition of both Kaul's and Ansari's teachings to the Ground 1 assertions do not remedy the

deficiency of Petitioner's obviousness position regarding the combined teachings of MIT2002, MIT2003, and MIT2004 for claim 1, from which claims 5 and 6 ultimately depend. Accordingly, for the same reasons we identify above, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of these claims would have been obvious.

H. Asserted Obviousness of Claims over Olofsson and Kaul

Petitioner contends that claims 1, 3–9, 11–13, and 15–19 of the '975 patent are unpatentable as obvious under 35 U.S.C. § 103, over the combined teachings of Olofsson in view of Kaul. Pet. 42.

We begin our analysis with a brief overview of Olofsson. Olofsson is directed to a mesh network in the form of a network on a chip having a three-dimensional grid of elements addressable by dedicated x, y, and z coordinates that reflects physical positions of processors in that grid. Ex. 1004, 2:16–22. Each transaction between processors consists of an address and data. *Id.* at 2:23–24. Olofsson's network employs a selectable routing scheme that can be random, fixed or adaptive. *Id.* at 2:33–34. The routing algorithm determines how data is routed from source to destination, attempting to minimize traffic without adding significant complexity. *Id.* at 1:55–61. Flow control is applied through buffers at each routing node. *Id.* at 1:62–64.

Upon review of Petitioner's explanations and supporting evidence, and the arguments presented in Patent Owner's Response and Sur-reply, we are persuaded that Petitioner has not demonstrated, by a preponderance of the evidence, that any of the challenged claims would have been obvious over the combined teachings of Olofsson and Kaul. In particular, we determine that Petitioner has not shown this combination to teach assigning

two or more functions to square or linear configurations based on optimal execution for that configuration, as set forth in limitation [1k].

Limitation [1i] recites, "an optimization module configured to." Limitation [1j] recites, "determine optimal function assignment configurations for groups of tiles." Limitation [1k] further recites that the optimization module is configured to

> assign two or more functions, which communicate at least unilaterally more frequently with one another than with other functions, to groups of adjacent tiles based on an optimal function assignment configuration determination, wherein the two or more functions are assigned to groups of tiles communicatively coupled in square configurations when the function executes optimally when executed by the groups of tiles communicatively coupled in the square configurations and are assigned to groups of tiles communicatively coupled in linear configurations when the function executes optimally when executed by the groups of tiles communicatively coupled in the linear configurations.

Petitioner asserts that Olofsson describes assigning particular functions to particular nodes, i.e., tiles. Pet. 59 (citing Ex. 1004, 7:24–27, 7:62–66). Petitioner asserts that, by comparing and choosing between square and linear configurations for the function assignment, Olofsson's system maximizes computational or network efficiency. *Id.* (citing Ex. 1004, 7:24–27, 7:30–33, 8:4–7, 8:13–16; Ex. 1002 ¶ 171). With respect to assigning frequently communicating functions to adjacent tiles, Petitioner asserts that Olofsson teaches this because,

> Olofsson's system assigns two or more processing functions to tiles based on the frequency of communication (including unilaterally) between adjacent tiles. [Ex. 1004,] 7:24-27 ("reduce communication bottlenecks"); 7:42-45 (functions located based on "nearest neighbor communication"); 7:66-8:2 (functional

> assignment made such that only adjacent tiles communicate); 8:4-7. Olofsson's sequencer thus co-locates functions "which communicate at least unilaterally more frequently with one another than with other functions."

Pet. 59–60. Petitioner further cites to Figures 10–13, and column 7, line 34–column 8, line 18, of Olofsson as exemplifying such assignment. *Id.* at 60.

With respect to assigning two or more functions to groups of adjacent tiles, Petitioner asserts that "no node needs to directly communicat[e] with any node that is not adjacent, which is a functional assignment to groups of adjacent tiles." *Id.* (citing Ex. 1004, 7:43–45, 7:66–8:2, 8:4–7).

With respect to assigning two or more functions based on an optimal function assignment configuration determination, Petitioner points to Olofsson's goals of "high computational efficiency," "flexibility and ease of implementation," and to "optimize utilization of network links." *Id.* at 61 (citing Ex. 1004, 7:24–33, 8:4–7, 8:13–16).

Patent Owner argues that no cited portion of Olofsson teaches "an optimal function configuration determination," or the use of such a determination to "assign functions in square or linear configurations. Resp. 59–64. Patent Owner argues that only in Petitioner's annotations is there an indication of linear or square configurations, and that such annotations do not show that Olofsson selects between such configurations based on an optimal configuration. *Id.* at 63 (citing Pet. 61–65). Patent Owner further argues that Petitioner's expert does not set forth any facts that would tend to support such a teaching. *Id.* at 62–64.

Petitioner responds by asserting that Patent Owner "does not appear to dispute that Olofsson does square and linear assignments," and that such assignments were performed to achieve "high computational efficiency" and

"reduce bottlenecks," i.e., to optimize such square and linear assignments. Pet. Reply 24 (citing Ex. 1002 ¶¶ 167–79).

Patent Owner, in its Sur-reply, argues that Petitioner identifies only that Olofsson has a general goal of efficiency when assigning functions. Sur-reply 22.

We agree with Patent Owner that Petitioner has insufficiently explained how its asserted linear and square configurations, shown in its annotations of Figures 10 and 11, are reflected in the description of Olofsson as the result of an optimal function assignment. For example, Petitioner asserts that Figure 11 "demonstrates both linear and square assignment of functions" because "the overall arrangement of rows and columns is square" while exhibiting "a linear workflow proceeding from left to right and from bottom to top." Pet. 63 (citing Ex. 1002 ¶ 177). Although we credit Dr. Conte's characterization of process flows in these figures, Olofsson merely teaches that assignment is made based on "which communicate at least unilaterally more frequently with one another than with other functions." Pet. 59–60.

Limitation [1k] requires that an additional consideration be taken into account when assigning functions; i.e., square configurations when the function executes optimally when executed by the groups of tiles in square configurations and linear configurations when the function executes optimally when executed by the groups of tiles in linear configurations. Petitioner has agreed to adopt Patent Owner's position that the "'linear' and 'square' configurations of tiles [are the] definite, structural way to perform the optimization." Pet. 7. Consistent with the claim language and the parties' positions, the District Court construed the limitation "based on an

optimal function assignment configuration determination" to mean "based on a determination of whether to assign a function to a particular square or linear configuration." Ex. 2003, 1.[6] The Court addressed only that part of limitation [1k], and did not explicitly address the entirety of the limitations in [1k]. *Id.* The plain language of [1k] clearly states that a square configuration is assigned "when" a square configuration provides the optimal execution of the functions. Thus, the claim requires more than determining the optimal execution of the functions based on general considerations—it requires determining whether a square or a linear configuration would provide the optimal execution of the functions, and assigning the functions to that configuration.

Petitioner has not pointed to any teaching in Olofsson that would show the assignment of configurations to include a determination of whether square or linear shape would result in an optimal function, and assignment of tiles based on that determination. Consequently, Petitioner has not shown Olofsson to teach limitation [1k]. Accordingly, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of claim 1 would have been obvious.

---

[6] We have considered the Court's determination that the optimization module of element [1j] should be construed as a means-plus-function term. The parties have addressed both constructions, primarily in sparring footnotes. Pet. 7, 22 n.7; Resp. 36–37 n.7; Reply 16 n.4; Sur-reply 21 n.3. We do not discern any distinction in Petitioner's assertions or Patent Owner's arguments at to element [1k] that turns on whether the [1j] term is construed as means-plus-function, and because we find Petitioner's assertions for [1k] lacking, we need not make such a determination for [1j].

d) *Claim 13*

Petitioner relies upon the same teachings as for claim 1 to account for the limitations of claim 13. Pet. 71–74. Petitioner additionally cites to a processor coupled to a memory. *See, e.g.*, *id.* at 71–73. However, claim 13 contains a limitation ([13i]) commensurate in scope with limitation [1k], which we determined was not taught by the combined teachings of Olofsson and Kaul. In its assertions for claim 13, Petitioner does not present additional arguments and evidence that remedies the deficiency in its analysis of the combined teachings of Olofsson and Kaul. Accordingly, for the same reasons we identify above, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of claim 13 would have been obvious.

e) *Dependent claims*

By virtue of their dependency, claims 3–9, 11, 12, and 15–19 include the same limitations as at least one of independent claims 1 and 13. Petitioner does not present additional arguments and evidence with respect to these dependent claims that remedy the deficiency in its analysis of the combined teachings of Olofsson and Kaul. Accordingly, for the same reasons we identify above, Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of these dependent claims would have been obvious.

I. Asserted Obviousness of claims 1–19 over Olofsson, Kaul, and Ansari

Petitioner relies upon Ansari for its "teaching of its specific types of static priority routing policies" to meet the limitations of element [1h] or the commensurate limitations of element [13f]. Pet. 75–77, 79. Petitioner relies

upon Olofsson and Kaul for all elements of independent claims 1 and 13 other than element [1h]. *Id.* at 75, 79. Because we determined that Olofsson and Kaul did not teach at least limitations [1k] and 13[i], Petitioner has not demonstrated by a preponderance of the evidence that the subject matter of claims 1–19 would have been obvious over Olofsson, Kaul, and Ansari.

III. CONCLUSION

Based on the fully developed record, Petitioner has not demonstrated by a preponderance of the evidence that claims 1–19 are unpatentable under § 103 as being obvious over the cited references. A summary of our conclusions is set forth in the table below.

Claims	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not Shown Unpatentable
1, 3, 4, 7–9, 11–13, 15–19	103	MIT2002, MIT2003, MIT2004		1, 3, 4, 7–9, 11–13, 15–19
1–4, 7–19	103	MIT2002, MIT2003, MIT2004, Ansari		1–4, 7–19
5, 6	103	MIT2002, MIT2003, MIT2004, Kaul		5, 6
5, 6	103	MIT2002, MIT2003, MIT2004, Kaul, Ansari		5, 6
1, 3–9, 11–13, 15–19	103	Olofsson, Kaul		1, 3–9, 11–13, 15–19
1–19	103	Olofsson, Kaul, Ansari		1–19
Overall Outcome				1–19

IV. ORDER

In consideration of the foregoing, it is

ORDERED that claims 1–19 of the '975 patent have not been shown to be unpatentable; and

FURTHER ORDERED that, because this is a Final Written Decision, parties to this proceeding seeking judicial review of our decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

PETITIONER:

James Glass
Nima Hefazi
Yury Kapgan (*pro hac vice*)
Gavin Snyder (*pro hac vice*)
QUINN EMANUEL URQUHART & SULLIVAN, LLP
jimglass@quinnemanuel.com
nimahefazi@quinnemanuel.com
yurykapgan@quinnemanuel.com
gavinsnyder@quinnemanuel.com

PATENT OWNER:

Jessica C. Kaiser
David A. Caine
Christopher T. Marando
ARNOLD & PORTER KAYE SCHOLER LLP
jessica.kaiser@arnoldporter.com
david.caine@arnoldporter.com
chris.marando@arnoldporter.com