Trials@uspto.gov
571-272-7822

Paper 53
Entered: March 1, 2022

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION and XILINX, INC.,[1]
Petitioner,

v.

FG SRC LLC,
Patent Owner.

IPR2020-01449
Patent 7,149,867 B2

Before KALYAN K. DESHPANDE, GREGG I. ANDERSON, and KARA L. SZPONDOWSKI, *Administrative Patent Judges*.

SZPONDOWSKI, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision
Determining All Challenged Claims Unpatentable
Denying Patent Owner's Motion to Amend
35 U.S.C. § 318(a)

[1] Xilinx, Inc. filed a motion for joinder and a petition in IPR2021-00633, which were granted, and, therefore, Xilinx, Inc. has been joined as petitioner in this proceeding.

I. INTRODUCTION

We instituted an *inter partes* review of claims 1–19 of U.S. Patent 7,149,867 B2 (Ex. 1001, "the '867 patent"), in response to a Petition (Paper 1, "Pet") filed by Intel Corporation ("Petitioner"). Paper 13 ("Dec."). During the trial, FG SRC LLC ("Patent Owner") filed a Response (Paper 34, "PO Resp."), Petitioner filed a Reply (Paper 40, "Reply"), and Patent Owner filed a Sur-reply (Paper 44, "Sur-reply").

Patent Owner also filed a Motion to Amend the claims of the '867 patent. Paper 26. After considering Petitioner's Opposition to the Motion to Amend (Paper 36), we issued Preliminary Guidance on Patent Owner's Motion (Paper 38). Patent Owner subsequently filed a Revised Motion to Amend the claims of the '867 patent that includes proposed substitute claims 20–38. Paper 41 ("Mot. Amend"). Petitioner opposed Patent Owner's Revised Motion to Amend (Paper 45, "Opp. Amend"), Patent Owner replied (Paper 49, "Reply Amend"), and Petitioner filed a Sur-reply (Paper 50, "Sur-reply Amend").

An oral hearing was held on January 6, 2022, and a copy of the transcript was entered into the record. Paper 52 ("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 the claims on which we instituted trial. Based on the complete record, Petitioner has shown, by a preponderance of the evidence, that claims 1–19 of the '867 patent are unpatentable. We also deny Patent Owner's Revised Motion to Amend, because Patent Owner has not met its burden in asserting that proposed substitute claims 20–38 have written description support in the original application that issued as the '867 patent.

II. BACKGROUND

A. Real Parties in Interest

Petitioner identifies Intel Corporation as the sole real party in interest. Pet. 2. Patent Owner identifies FG SRC LLC as the sole real party in interest. Paper 4, 2.

B. Related Matters

The parties advise that the '867 patent is the subject of the following district court litigations:

FG SRC LLC v. Intel Corporation, 6:20-cv-00315-ADA (W.D. Tex.), filed April 24, 2020 ("the co-pending district court litigation");

FG SRC LLC v. Xilinx, Inc., 1:20-cv-00601-LPS (D. Del.), filed April 30, 2020; and

SRC Labs, LLC et al. v. Amazon Web Services, Inc., et al., 2:18-cv-00317-JLR (W.D. Wash.), filed February 26, 2018.

Pet. 2; Paper 4, 2. Petitioner also advises that the '867 patent was the subject of IPR2019-00103 (institution denied on May 10, 2019). Pet. 2.

C. The '867 Patent (Ex. 1001)

The '867 patent issued from Application No. 10/869,200 filed June 16, 2004, and claims the benefit of Provisional Application No. 60/479,339, filed June 18, 2003. Ex. 1001, codes (21), (22), (60). The '867 patent is titled "System and Method of Enhancing Efficiency and Utilization of Memory Bandwidth in Reconfigurable Hardware" and is generally directed to "enhancing the efficiency and utilization of memory bandwidth in reconfigurable hardware" and "implementing explicit memory hierarchies in reconfigurable processors that make efficient use of off-board, on-board, on-chip storage and available algorithm locality." *Id.* at code (57), 1:15–24.

According to the '867 patent, there was a growing need to develop improved memory hierarchies that limited overhead of a memory hierarchy without also reducing bandwidth efficiency and utilization. Ex. 1001, 3:57–60. The '867 patent describes a system including a memory hierarchy and a reconfigurable processor that includes a data prefetch unit. *Id.* at 4:4–10, 5:60–62, 6:9–13, 7:34–48. The '867 patent states that a "Reconfigurable Processor" is "a computing device that contains reconfigurable components such as FPGAs [(field programmable gate arrays)] and can, through reconfiguration, instantiate an algorithm as hardware." *Id.* at 5:26–29. The '867 patent states that a "Data prefetch Unit" is "a functional unit [a set of logic that performs a specific operation] that moves data between members of a memory hierarchy [a collection of memories]," where such "movement may be as simple as a copy, or as complex as an indirect indexed strided copy into a unit stride memory." *Id.* at 5:34–43.

Figure 1 of the '867 patent, reproduced below, shows a reconfigurable processor (RP) 100 of the claimed invention. *Id.* at 4:38–40.

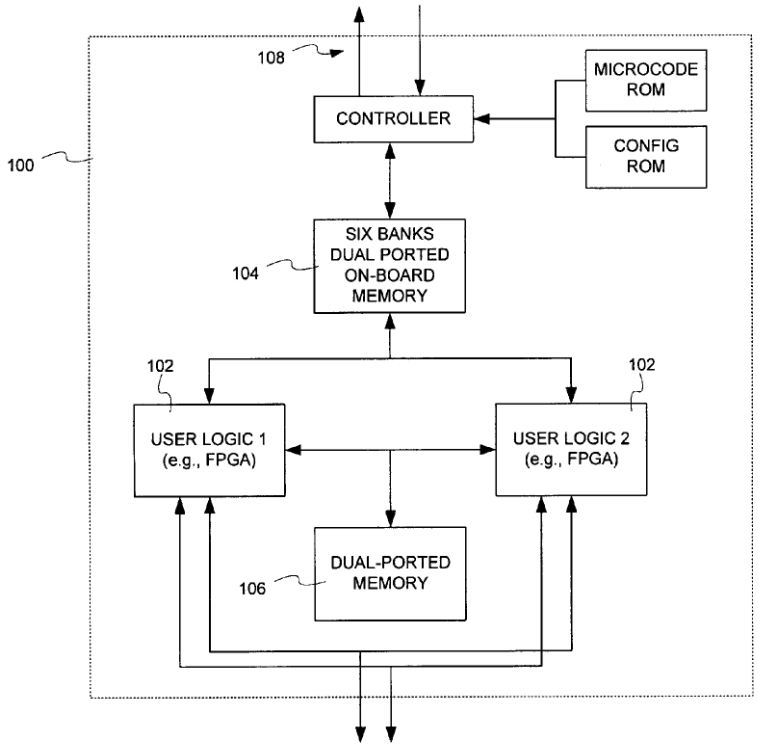

FIG. 1

Figure 1 depicts a reconfigurable processor (RP) 100. *Id.* at 4:38–40.

Figure 1 depicts reconfigurable processor 100, which "may be implemented using field programmable gate arrays (FPGAs) or other reconfigurable logic devices, that can be configured and reconfigured to contain functional units and interconnecting circuits, and a memory hierarchy comprising on-board memory banks 104, on-chip block RAM 106, registers wires, and a connection 108 to external memory." *Id.* at 6:5–11. In addition, "[o]n-chip reconfigurable components 102 create memory structures such as registers, FIFOs, wires and arrays using block RAM." *Id.* at 6:11–14. "Dual-ported memory 106 is shared between on-chip reconfigurable components 102." *Id.* at 6:14–15. "The reconfigurable processor 100 also implements user-defined computational logic . . . constructed by programming an FPGA to implement a particular interconnection of computational functional units." *Id.* at 6:15–19. "In a

particular implementation, a number of RPs 100 are implemented within a memory subsystem of a conventional computer, such as on devices that are physically installed in dual inline memory module (DIMM) sockets of a computer." *Id.* at 6:19–23. "In this manner the RPs 100 can be accessed by memory operations and so coexist well with a more conventional hardware platform." *Id.* at 6:23–25. The '867 patent explains that "[u]nlike conventional static hardware platforms . . . the memory hierarchy provided in a RP 100 is reconfigurable" and "through the use of data access units and associated memory hierarchy components, computational demands and memory bandwidth can be matched." *Id.* at 7:17–22.

One or more data prefetch units are used to improve the memory hierarchy and bandwidth efficiency and utilization. *Id.* at 3:58–60, 8:62–65. Fig. 4 of the '867 patent, reproduced below, depicts a logic block 300 with an addition of a data prefetch unit 401. *Id.* at 4:44–46.

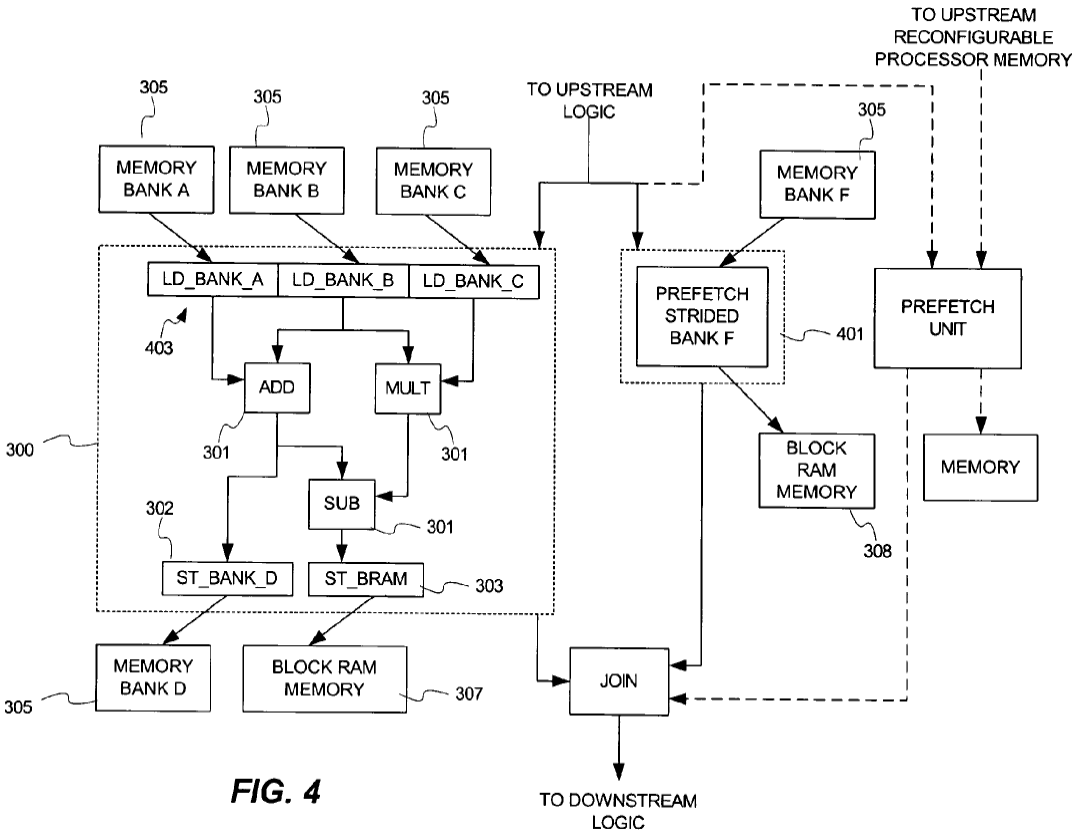

FIG. 4

Figure 4 illustrates a logic block 300 (a block composed of computational functional units capable of taking data and producing results with each clock pulse) with the addition of a data prefetch unit 401. *Id.* at 7:6–8, 7:34–35.

Logic block 300 includes computational functional units (computational logic) 301, 302, and 303, a control, and data access functional units 403 that present data to computational logic 301, 302, and 303. *Id.* at 7:25–48, Fig. 4. Data prefetch unit 401 moves data from one member of the memory hierarchy 305 to another 308 (a block RAM memory). *Id.* at 7:34–37, Fig. 4. Data prefetch unit 401 operates "independently of other functional units 301, 302, and 303 and can therefore operate prior to, in parallel with, or after computational logic." *Id.* at 7:37–40. In addition, data prefetch unit 401 may be "operated independently of logic block 300 that uses prefetched data." *Id.* at 7:45–48. Data prefetch

unit 401 deposits data into the memory hierarchy, where computational logic 301, 302, and 303 can access it through data access units. *Id.* at 7:42–44.

The '867 patent explains:

> An important feature of the present invention is that many types of data prefetch units can be defined so that the prefetch hardware can be configured to conform to the needs of the algorithms currently implemented by the computational logic. The specific characteristics of the prefetch can be matched with the needs of the computational logic and the format and location of data in the memory hierarchy.

Id. at 7:49–55. The '867 patent provides examples of configuring a data prefetch unit depending on the needs of the computational logic. *Id.* at 7:52–62, 8:3–21, Figs. 9A–9B.

D. Illustrative Claims

Among the challenged claims, claims 1, 9, and 13 are independent. Independent claims 1, 9, and 13 are reproduced below, with brackets noting Petitioner's identifiers.

> 1. [preamble] A reconfigurable processor that instantiates an algorithm as hardware comprising:
>
> [1(a)] a first memory having a first characteristic memory bandwidth and/or memory utilization; and
>
> [1(b)] a data prefetch unit coupled to the first memory, [1(c)] wherein the data prefetch unit retrieves only computational data required by the algorithm from a second memory of second characteristic memory bandwidth and/or memory utilization and places the retrieved computational data in the first memory [1(d)] wherein the data prefetch unit operates independent of and in parallel with logic blocks using the computional [sic] data, and [1(e)] wherein at least the first memory and data prefetch unit are configured to conform to needs of the algorithm, and [1(f)] the data prefetch unit is configured to match format and location of data in the second memory.

9. [preamble] A reconfigurable hardware system, comprising:

[9(a)] a common memory; and

[9(b)] one or more reconfigurable processors that can instantiate an algorithm as hardware coupled to the common memory, [9(c)] wherein at least one of the reconfigurable processors includes a data prefetch unit to read and write only data required for computations by the algorithm between the data prefetch unit and the common memory [9(d)] wherein the data prefetch unit operates independent of and in parallel with logic blocks using the computational data, and [9(e)] wherein the data prefetch unit is configured to conform to needs of the algorithm and [9(f)] match format and location of data in the common memory.

13. [preamble] A method of transferring data comprising:

[13(a)] transferring data between a memory and a data prefetch unit in a reconfigurable processor; and

[13(b)] transferring the data between a computational unit and a data access unit, [13(c)] wherein the computational unit and the data access unit, and the data prefetch unit are configured to conform to needs of an algorithm implemented on the computational unit and transfer only data necessary for computations by the computational unit, and [13(d)] wherein the prefetch unit operates independent of and in parallel with the computational unit.

Ex. 1001, 12:39–54; 13:13–26; 14:1–11.

E. Evidence

Petitioner relies on the following references (*see* Pet. 4–5).

Reference	Exhibit	Patent/Printed Publication
Zhang	1003	Xingbin Zhang et al., *Architectural Adaptation of Application-Specific Locality Optimizations*, published in the Proceedings of the International Conference on Computer Design - VLSI in Computers and Processors (IEEE, October 12–

		15, 1997), 150–156
Gupta	1004	Rajesh Gupta, *Architectural Adaptation in AMRM Machines*, Proceedings of the IEEE Computer Society Workshop on VLSI 2000 (IEEE, April 27–28, 2000), 75–79
Chien	1005	Andrew A. Chien et al., *MORPH: A System Architecture for Robust High Performance Using Customization (An NSF 100 TeraOps Point Design Study)*, Proceedings of Frontiers '96 – The Sixth Symposium on the Frontiers of Massively Parallel Computing (IEEE, October 27–31, 1996), 336–345

In addition, Petitioner relies on the Declarations of Rajesh K. Gupta, Ph.D. (Exs. 1010, 1030), Declarations of Jacob Robert Munford (Exs. 1012, 1031), Declaration of Gordon MacPherson (Ex. 1027), Declaration of Eileen D. McCarrier (Ex. 1028), Declaration of Austin Schnell (Ex. 1029), and Declaration of Dr. Stanley Shanfield (Ex. 1006).

Patent Owner relies on the Declaration of Dr. William Mangione-Smith (Ex. 2028) and Declaration of Ryan Kastner, Ph.D. (Ex. 2010).

Deposition transcripts have been entered into the record for Dr. Gupta (Ex. 1039), Mr. MacPherson (Ex. 1040), Dr. Shanfield (Exs. 1043, 2029), and Dr. Mangione-Smith (Ex. 1044).

F. Prior Art and Asserted Grounds

Petitioner asserts that claims 1–19 are unpatentable on the following grounds (Pet. 5):

Claims Challenged	35 U.S.C. §[2]	References
1, 2, 4–8, 13–19	103	Zhang, Gupta
3, 9–12	103	Zhang, Gupta, Chien

III. ANALYSIS

A. Legal Standards

A claim is unpatentable under 35 U.S.C. § 103(a) if "the differences between the subject matter sought to be patented 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) objective evidence of nonobviousness, i.e., secondary considerations. *See Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

A patent claim "is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art." *KSR*, 550

[2] The Leahy-Smith America Invents Act ("AIA"), Pub. L. No. 112-29, 125 Stat. 284 (2011), amended 35 U.S.C. § 103, effective March 16, 2013. Because the application from which the '867 patent issued was filed before this date, the pre-AIA version of § 103 applies.

U.S. at 418. An obviousness determination requires finding "both 'that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.'" *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367–68 (Fed. Cir. 2016) (citation omitted); *see KSR*, 550 U.S. at 418. Further, an assertion of obviousness "cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness." *KSR*, 550 U.S. at 418; *In re NuVasive, Inc.*, 842 F.3d 1376, 1383 (Fed. Cir. 2016) (a finding of a motivation to combine "must be supported by a 'reasoned explanation'" (citation omitted)).

B. Level of Ordinary Skill in the Art

Petitioner asserts a person of ordinary skill in the art "would have had an undergraduate degree in electrical engineering or related field with at least three years of experience in computer processor architecture and FPGAs, a master's degree with two or more years of experience in those fields, or an equivalent combination of education and experience." Pet. 14 (citing Ex. 1006 ¶ 67).

Patent Owner does not dispute Petitioner's proposed level of skill. *See generally* PO Resp.

We find Petitioner's proposal is consistent with the level of ordinary skill in the art reflected by the prior art of record, and, therefore, adopt Petitioner's proposed level of ordinary skill in the art for purposes of this Decision. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001).

C. Claim Construction

We construe each claim "in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent," the same standard used to construe the claim in a civil action. 37 C.F.R. § 42.100(b) (2020).

1. Agreed Constructions

The parties agree to the construction of the following terms:

Claim term	**Proposed Construction**
Reconfigurable Processor	a computing device that contains reconfigurable components such as FPGAs and can, through reconfiguration, instantiate an algorithm as hardware
Data Prefetch Unit	a functional unit that moves data between members of a memory hierarchy. The movement may be as simple as a copy, or as complex as an indirect indexed strided copy into a unit stride memory
Data Access Unit	a functional unit that accesses a component of a memory hierarchy, and delivers data directly to the computational logic
Functional Unit[3]	a set of logic that performs a specific operation. The operation may for example be arithmetic, logical, control, or data movement. Functional units are used as building blocks of reconfigurable logic
Memory Hierarchy[4]	a collection of memories

[3] Petitioner states that although not directly recited in the claims, this term is used in the '867 patent's definition of "data prefetch unit" and "data access unit." Pet. 15; Ex. 1001, 5:40–46.

[4] Petitioner states although not directly recited in the claims, this term is used in the '867 patent's definition of "data prefetch unit" and "data access unit." Pet. 15; Ex. 1001, 5:40–46.

Pet. 14–16 (citing Ex. 1001; Ex. 1006); PO Resp. 29–30. We adopt these agreed constructions for purposes of this Decision.

Patent Owner also states that in the co-pending district court litigation, the parties have agreed to the following constructions:

Claim term	Agreed Construction
(Preamble) A reconfigurable processor that instantiates an algorithm as hardware	Preamble is limiting
Common Memory	an external memory shared by processors in a multiprocessor system
Computational Unit	a functional unit of a reconfigurable processor that performs a computation
the Data Prefetch Unit Receives Processed Data	the data prefetch unit receives the results of the algorithm
Configured To Conform to Needs of The Algorithm	configured in reconfigurable logic to conform to the needs of the algorithm
Reconfigurable Logic	reconfigurable logic is composed of an interconnection of functional units, control, and storage that implements an algorithm and can be loaded into a Reconfigurable Processor

PO Resp. 30. We further adopt these agreed constructions for purposes of this Decision.

2. Patent Owner's Proposed Constructions

Patent Owner proposes construction for two terms.

> *a) "retrieves only computational data required by the algorithm from a second memory . . . and places the retrieved computational data in the first memory" (limitation 1(c))*

Patent Owner argues that this limitation should be construed as "retrieves from a second memory that computational data which is required by the algorithm ***and no other computational data*** … and places the

retrieved computational data in the first memory." PO Resp. 31–32 (emphasis added).

Patent Owner contends that "[t]he plain meaning requires that no superfluous computational data is transferred to the first memory." *Id.* at 31. In support of its proposed construction, Patent Owner identifies column 9, lines 1–5 of the '867 patent, which states "an important feature of the present invention is the ability to implement various kinds or styles of prefetch units to meet the needs of a particular algorithm being implemented by computational elements." *Id.* Patent Owner also relies on column 9, lines 8–10, which states "in most cases the function being implemented by components 301 would change and therefore alter the decision as to which prefetch strategy is most appropriate." *Id.* at 31–32. Patent Owner also relies on column 9, lines 35–40, which states "[g]ains are made by delivering only requested data from transfer buffer 1305 (not the remainder of a data block as in cache line oriented systems) by eliminating the need to transfer an index array either to the processor or to the memory controller." *Id.* at 32.

Petitioner contends that Patent Owner does not rely on this construction to overcome the prior art, and Patent Owner's expert does not rely on (or even recite) this construction. Reply 2 (citing PO Resp. 40–44; Ex. 1044, 28:18–29:5, 32:10–14, 107:19–108:7). Petitioner also argues that Patent Owner's proposed construction is unsupported by intrinsic evidence. *Id.* at 3–4.

Patent Owner has not explained the impact of this proposed claim construction on its arguments supporting patentability of the claims. We disagree with Patent Owner that its proposed construction is the "plain and ordinary meaning," because the claim language recites "retrieves ***only***

computational data ***required by the algorithm*** from a second memory," and does not require the additional limitation "and no other computational data." Under these circumstances, we decline to rewrite the claim language to include additional words that are not present in the claim, e.g., "and no other computational data." *See Hoganas AB v. Dresser Indus., Inc.*, 9 F.3d 948, 950 (Fed. Cir. 1993) (quoting *E.I. Du Pont de Nemours & Co. v. Phillips Petroleum Co.*, 849 F.2d 1430, 1433 (Fed. Cir. 1988)) ("It is improper for a court to add 'extraneous' limitations to a claim, that is, limitations added 'wholly apart from any need to interpret what the patentee meant by particular words or phrases in the claim.'").

Moreover, in light of the parties' arguments, we determine that we need not expressly construe this term. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) ("[W]e need only construe terms 'that are in controversy, and only to the extent necessary to resolve the controversy.'" (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng'g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999))). We agree with Petitioner that Patent Owner does not rely on this construction to overcome the prior art; for example, Dr. Mangione-Smith testified that he had not been given or applied any particular claim construction. *See* Ex. 1044, 28:18–32:14; 107:19–109:4; *see* PO Resp. 40–44; Ex. 2028 ¶ 74.

b) "read and write only data required for computations by the algorithm between the data prefetch unit and the common memory" (claim 9)

Patent Owner proposes that this term should be construed as "read, using the data prefetch unit, only data required for computations by the algorithm from common memory and write, using the data prefetch unit, only data required for computations by the algorithm." PO Resp. 33. In

support of its construction, Patent Owner states that "[t]his term should be accorded its plain and ordinary meaning." *Id.*

Petitioner contends that Patent Owner does not rely on this construction to overcome the prior art, and Patent Owner's expert does not rely on (or even recite) this construction. Reply 2 (citing PO Resp. 40–44; Ex. 1044, 28:18–29:5, 32:10–14, 107:19–108:7). Petitioner also argues that Patent Owner's proposed construction is unsupported by intrinsic evidence. *Id.* at 3–4.

Patent Owner has not explained the impact of this proposed claim construction on its arguments supporting patentability of the claims. Moreover, Patent Owner does not explain why we should adopt this proposed construction, or provide any supporting intrinsic evidence. Moreover, given that Patent Owner relies on the same arguments for claim 9 as it presented for claim 1 (which does not recite this limitation), we agree with Petitioner that Patent Owner does not rely on this construction in order to overcome the prior art. *See* PO Resp. 53–55; *see, e.g.,* Ex. 1044, 28:18–32:14 (Dr. Mangione-Smith testifying that he had not been given or applied any particular claim construction). Accordingly, we determine that we need not expressly construe this term to resolve the parties' dispute.

D. The Asserted Prior Art References

1. Zhang (Ex. 1003)

Zhang is a paper published by the Institute of Electrical and Electronics Engineers, Inc. (hereafter "IEEE") as part of the *Proceedings of the International Conference on Computer Design - VLSI in Computers and Processors*. Ex. 1003, 1–2, 4.[5] Zhang describes "a machine architecture

[5] Citations to Exhibit 1003 are to the page numbering provided by Petitioner.

that integrates programmable logic into key components of the system with the goal of customizing architectural mechanisms and policies to match an application," using application-specific hardware assists. *Id.* at 12 (Abstract). Zhang "demonstrate[s] that application-specific hardware assists and policies can provide substantial improvements in performance on a per application basis." *Id.* Zhang's architecture "integrates small blocks of programmable logic into key elements of a baseline architecture, including processing elements, components of the memory hierarchy, and the scalable interconnect, to provide *architectural adaptation*—the customization of architectural mechanisms and policies to match an application." *Id.* at 13. Zhang explains that architectural adaptation provides mechanisms for application-specific hardware assists to overcome rigid architectural choices that do not work well across different applications, as "integration of programmable logic with memory components enables application-specific locality optimizations." *Id.* at 13–14.

Zhang's architecture is depicted in Figure 2 below:

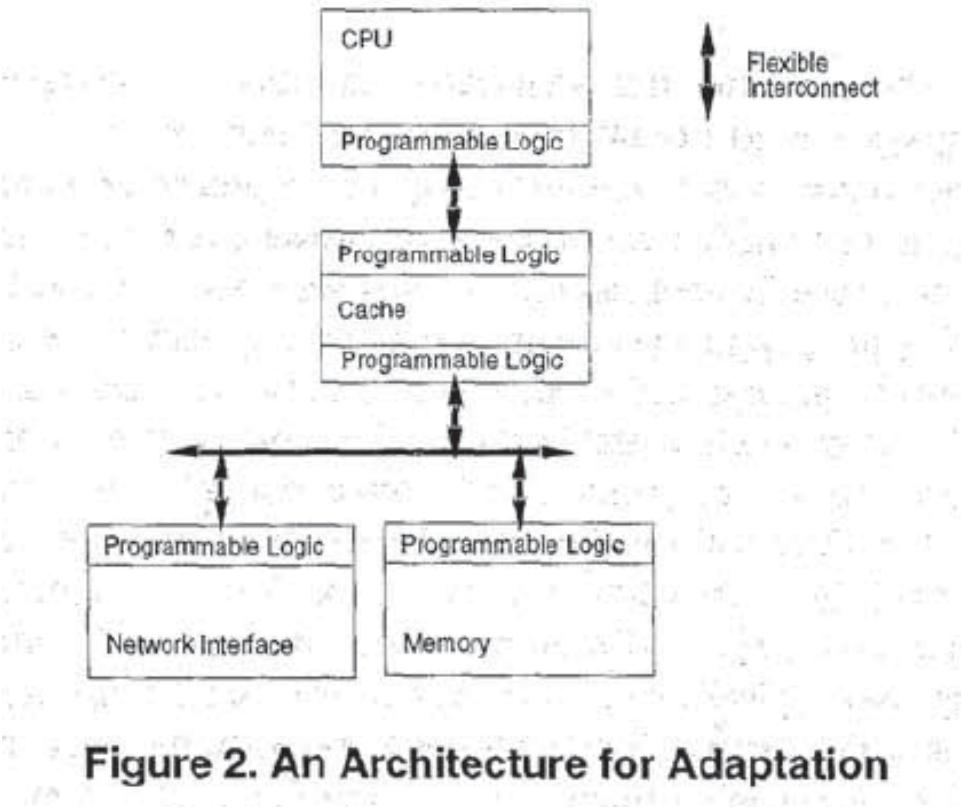

Figure 2. An Architecture for Adaptation

Figure 2 of Zhang depicts programmable logic integrated with CPU, cache, network interface, and memory.

Zhang presents "two case studies of architectural adaption using application-specific knowledge to enhance latency tolerance and efficiently utilize network bisection on multiprocessors." Ex. 1003, 14. The first case study uses architectural adaptation for prefetching and exploits application access pattern information. *Id.* at 15. Figure 4 of Zhang, reproduced below, depicts a prefetcher implementation for Zhang's first case study, using programmable logic integrated with the L1 cache. *Id.*

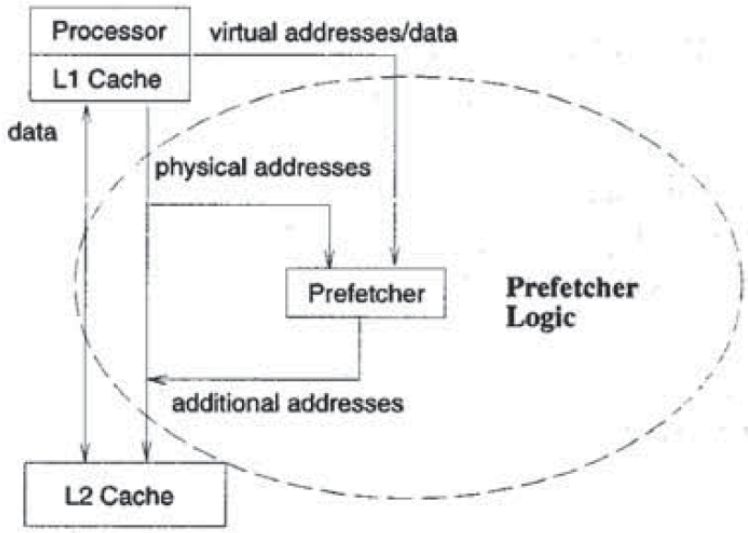

Figure 4. Organizations of Prefetcher Logic

Figure 4 shows a prefetcher implementation using programmable logic integrated with the L1 cache.

The prefetcher in Figure 4 requires two pieces of application-specific information: address ranges and memory layout of the target data structures. *Id.* at 15. The address range, which is application dependent, is needed to indicate memory bounds where prefetching is likely to be useful. *Id.* Prefetching can be enabled or disabled, and is triggered only by read misses. *Id.* Once the prefetcher is enabled, it determines what and when to prefetch by checking virtual addresses of cache lookups to check whether a matrix element is being accessed. *Id.* In one example, records spanning multiple cache lines are targeted to "prefetch[] all fields of a matrix element structure whenever some field of the element is accessed." *Id.* Because each matrix element (which is padded to 64 bytes) spans two cache lines (in a cache with cache line size of 32 bytes), the prefetcher generates an additional L2 cache lookup address from the given physical address that prefetches the other cache line not yet referenced. *Id.* In a second example, particular pointer

fields (those likely to be traversed when their parent structures are accessed) are targeted. *Id.* at 15–16. For example, in a sparse matrix-vector multiply, the record pointed to by the *nextRow* field is accessed close in time with the current matrix element. *Id.* A prefetcher generates an additional address after the initial cache miss is satisfied using the *nextRow* pointer value embedded in the data returned by the L2 cache. *Id.* at 16.

Zhang's second case study uses a sparse matrix-matrix multiply routine to show architectural adaptation that improves data reuse and reduces data traffic between the memory unit and the processor. *Id.* at 15–16. "The architectural customization aims to send only used fields of matrix elements during a given computation to reduce bandwidth requirement using dynamic scatter and gather." *Id.* at 16. "The two main ideas are prefetching of whole rows or columns using pointer chasing in the memory module and packing/gathering of only the used fields of the matrix element structure." *Id.* Figure 5 of Zhang, reproduced below, illustrates an architecture including a cache and a main memory module, and containing two units of logic, an address translation logic and a gather logic.

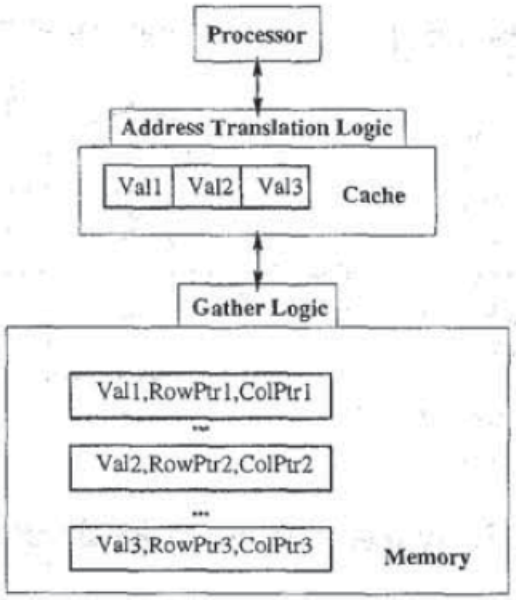

Figure 5. Scatter and Gather Logic

Figure 5 shows a scatter and gather logic using two units of logic, an address translation logic and a gather logic. *Id.* at 16.

2. Gupta (Ex. 1004)

Gupta is a paper published by IEEE as part of the *Proceedings of the IEEE Computer Society Workshop on VLSI 2000*. Ex. 1004, 1, 3, 6.[6] Gupta describes an Adaptive Memory Reconfiguration Management (AMRM) prototype architecture that implements a board-level prototype designed to "simulate a range of memory hierarchies for applications running on a host processor" and "supports configurability of [a] cache memory via an on-board FPGA-based memory controller." *Id.* at 8–10. The goals of Gupta's AMRM prototype are that "it be adaptable to many different memory hierarchy architectures," "be useful for running real time program execution or even memory simulations," and demonstrate "a specific mechanism for

[6] Citations to Exhibit 1004 are to the page numbering provided by Petitioner.

latency management . . . to provide significant performance boost for the class of applications characterized by frequent accesses to linked data structures scattered in the physical memory." *Id.* at 9–10.

Figure 1 of Gupta, reproduced below, illustrates an AMRM prototype board.

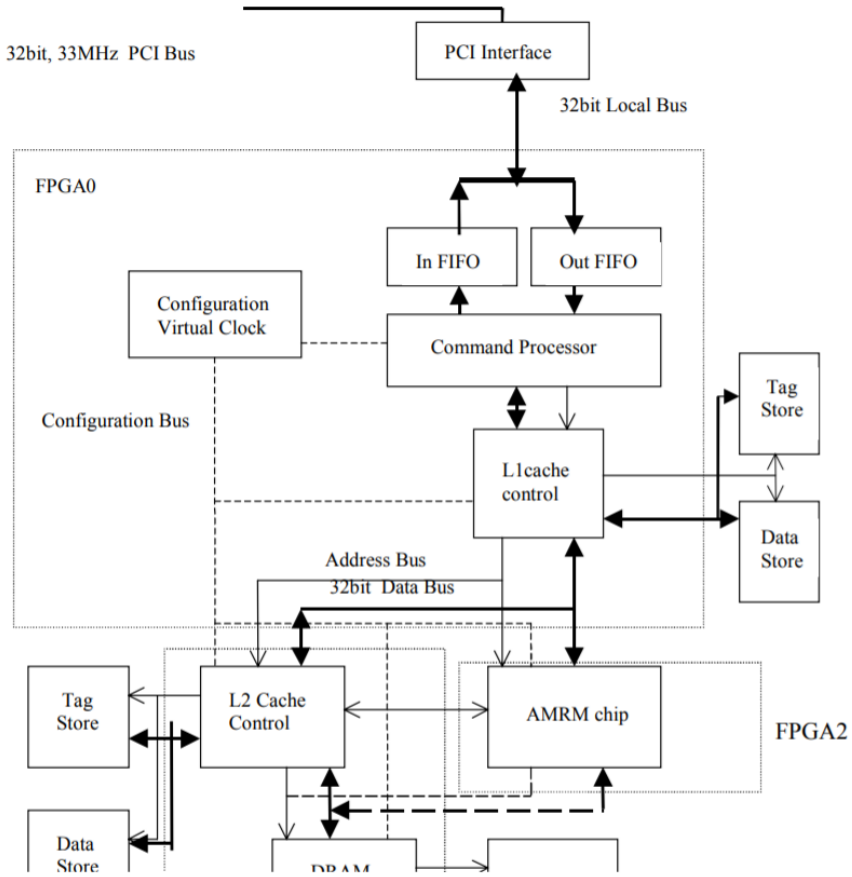

Figure 1: AMRM Phase I Prototype Board

Figure 1 illustrates the main components of an AMRM prototype board.

As shown in Figure 1, an AMRM prototype board includes a general 3-level memory hierarchy plus support for an AMRM ASIC chip implementing architectural assists within a CPU-L1 datapath. *Id.* at 10. The FPGAs on the board contain controllers for the SRAM, DRAM and L1 cache. *Id.* The AMRM chip is positioned between the L1 cache and the rest

of the system and can be accessed in parallel with the L2 cache. *Id.* at 11. The AMRM chip can thus accept and supply data coming from or going to the L1 cache. *Id.* The AMRM chip may contain a write buffer or a prefetch unit to access L2, and also has access to the memory interface and can prefetch from memory. *Id.*

3. Chien (Ex. 1005)

Chien is a paper published by the IEEE Computer Society Press as part of the *Proceedings of Frontiers '96–The Sixth Symposium on the Frontiers of Massively Parallel Computing*. Ex. 1005, 1–2, 5.[7] Chien describes a design and architecture of a MultiprocessOr with Reconfigurable Parallel Hardware (MORPH) that "uses reconfigurable logic blocks integrated with the system core to control policies, interactions, and interconnections" and has "configurability [that] supports component software and interoperabilty frameworks, allowing direct support for application-specified patterns, objects, and structures." *Id.* at 7 (emphasis omitted). Chien explains that "[t]he benefit of configurable logic is that it can be used to customize the machine's behavior to better match that required by the application—in essence a machine can be tuned for each application with little or no performance penalty for this generality." *Id.*

Chien states:

> The key elements of the MORPH architecture include processing elements and memory elements embedded in a scalable interconnect. The scalable interconnect flexibly connects all parts of the system with fast packet routing, efficiently exploiting the wiring resources provided by the system packaging []. The hardware structure allows adaptation of data transport, coordination, association (for granularity), and efficient computation. . . . MORPH could be used to implement either a

[7] Citations to Exhibit 1005 are to the page numbering provided by Petitioner.

> cache-coherent machine, a non-cache coherent machine, or even clusters of cache coherent machines connected by put/get or message passing. Varying the mix of processing and memory elements supports a wide range of machine configurations and balances. Examples of other possible changes include changes in cache block size, branch predictors, or prefetch policies.

Id. at 10.

4. Printed Publication Status of Zhang, Gupta, and Chien

a) Background

In the Institution Decision, we determined that Petitioner had established a reasonable likelihood that it would prevail in showing the unpatentability of the challenged claims over the asserted prior art. Dec. 34–44. This included a determination that, for purposes of institution and in accordance with the Board's precedential decision in *Hulu, LLC v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29 at 13 (PTAB Dec. 20, 2019) (precedential), Petitioner had established a reasonable likelihood that Zhang, Gupta, and Chien each qualify as prior art printed publications. Dec. 19–23. We stated that "[t]o the extent Patent Owner continues to challenge the printed publication status of these references after institution, the parties are requested to further develop the record on this issue." *Id.* at 44. Following the Institution Decision, we granted Petitioner's Motion to Submit Supplemental Information pursuant to 37 C.F.R. § 42.123(a), allowing Petitioner to submit Exhibits 1027–1031[8] as supplemental

[8] Exhibit 1027 is the Declaration of Gordon MacPherson; Exhibit 1028 is the Declaration of Eileen D. McCarrier; Exhibit 1029 is the Declaration of Austin M. Schnell; Exhibit 1030 is the Supplemental Declaration of Rajesh K. Gupta, Ph.D.; and Exhibit 1031 is the Supplemental Declaration of Jacob Robert Munford.

information, which Petitioner contends confirms the public accessibility of Zhang, Gupta, and Chien references. Paper 27.

b) Petitioner's Contentions

Petitioner contends Zhang was publicly accessible shortly after November 18, 1997, and no later than August 1, 2002, and, therefore, qualifies as prior art under 35 U.S.C. §§ 102(a) and 102(b). Pet. 17. Specifically, Petitioner argues (1) Zhang was published in 1997 by IEEE; (2) Zhang was distributed to conference attendees prior to or during the conference; (3) Zhang has been available on the IEEE Xplore website since at least as early as August 6, 2002; and (4) Zhang was catalogued by several university libraries by as early as November 18, 1997 and no later than August 1, 2002. *See* Pet. 17; Ex. 1010 ¶¶ 21–23; Ex. 1012 ¶¶ 15–20; Reply 4–8; Ex. 1030 ¶¶ 7–8; Ex. 1027 ¶ 11; Ex. 1031 ¶¶ 21–35.

Petitioner contends Gupta was publicly accessible by shortly after May 15, 2000, and therefore qualifies as prior art under 35 U.S.C. §§ 102(a) and 102(b). Pet. 19. Specifically, Petitioner argues (1) Gupta was published in 2000 by IEEE; (2) Gupta was distributed to conference attendees prior to or during the conference; (3) Gupta has been available on the IEEE Xplore website since at least as early as August 6, 2002; and (4) Gupta was catalogued by several university libraries as of May 15, 2000. *See* Pet. 19; Ex. 1010 ¶¶ 24–26; Ex. 1012 ¶¶ 21–26.

Petitioner contends Chien was publicly accessible by shortly after November 18, 1996, and therefore qualifies as prior art under 35 U.S.C. §§ 102(a) and 102(b). Pet. 21. Specifically, Petitioner argues (1) Chien was published in 1996 by IEEE; (2) Chien was distributed to conference attendees prior to or during the conference; (3) Chien has been available on the IEEE Xplore website since at least as early as August 6, 2002; and (4)

Chien was catalogued by several university libraries by shortly after November 18, 1996. *See* Pet. 19; Ex. 1010 ¶¶ 18–20; Ex. 1012 ¶¶ 27–32.

In support of its contentions that Zhang, Gupta, and Chien are printed publications, Petitioner relies on (1) the Declarations of Rajesh K. Gupta, Ph.D., who is one of the authors of Zhang, Gupta, and Chien; (2) the Declarations of Jacob Robert Munford, who has a Master of Library and Information Science degree and over ten years of experience in the library/information science field; (3) the Declaration of Gordon MacPherson, the Director of Board governance & IP Operations at IEEE; (4) the Declaration of Eileen D. McCarrier, Manager of Research Services at Pillsbury Winthrop Shaw Pittman LLP (Patent Owner's counsel); and (5) the Declaration of Austin M. Schnell, an associate at Pillsbury Winthrop Shaw Pittman LLP. Ex. 1010 ¶ 21; Ex. 1030; Ex. 1012 ¶ 2; Ex. 1031; Ex. 1027 ¶ 1; Ex. 1028 ¶ 1; Ex. 1029 ¶ 1.

Dr. Gupta testifies that (1) he presented each of Zhang, Gupta, and Chien at a conference sponsored or organized by IEEE in 1996, 1997, and 2000, respectively; (2) that based on his "experience in attending conferences [sponsored/organized] by the IEEE, and based on the general practice in the scientific and engineering community," he "believe[s] [each of Zhang, Gupta, and Chien] was distributed to the conference attendees prior to or during the conference"; (3) that each of Zhang, Gupta, and Chien "was subsequently published . . . by the IEEE";[9] (4) and he "understand[s]

[9] In his original Declaration, Dr. Gupta states that Gupta was published in 1997. Ex. 1010 ¶¶ 25, 26. Patent Owner points out that Gupta was not authored until 2000. PO Resp. 23. In the Institution Decision, we agreed with Petitioner that this appeared to be a typographical error. Dec. n.7. Dr. Gupta corrected the error in his Supplemental Declaration. Ex. 1030 n. 1.

that [each of Zhang, Gupta, and Chien] has been available from the IEEE Xplore website . . . since at least as early as August 6, 2002." Ex. 1010 ¶¶ 18–26. Dr. Gupta provides additional testimony that, based on his personal knowledge, each of the three references was distributed by the last day of its respective conference. Ex. 1030 ¶¶ 4–5, 7–8, 10–11.

Mr. Munford testifies that "[i]n preparing a material for public availability, a library catalog record describing that material would be created," which is typically written in Machine Readable Catalog ("MARC") code. Ex. 1012 ¶ 12. According to Mr. Munford, "the 008 field of the MARC record is reserved for denoting the date of creation of the library record itself" and "it is my experience that an item's MARC record indicates the date of an item's public availability." *Id.* He testifies that "[i]n my experience, the vast majority of library books cataloged and prepared for public availability in this fashion are made publicly available within 1 to 10 weeks of initial record creation." *Id.* ¶ 13. In addition, he testifies "[w]hen a material is held by multiple libraries, comparing the 008 fields of those records provides a specific window for public availability." *Id.*

For Zhang, Mr. Munford provides the MARC records that he secured from the online catalogs of the University of Cincinnati, Cornell University, and Michigan State University, purporting to show catalog dates of November 18, 1997, August 1, 2002, and November 18, 1997, respectively. *Id.* ¶¶ 16–19. Based on these records, Mr. Munford opines that Zhang "was made available and accessible to the public by shortly after November 18, 1997 and certainly no later than August 1, 2002." *Id.* ¶ 20.

For Gupta, Mr. Munford provides the MARC records that he secured from the online catalogs of Georgia Tech, Notre Dame University, and Linda Hall Library, purporting to show catalog dates of May 15, 2000 for all

three libraries. *Id.* ¶¶ 22–25. Based on these records, Mr. Munford opines that Gupta "was made available and accessible to the public by shortly after May 15, 2000." *Id.* ¶ 26.

For Chien, Mr. Munford provides the MARC records that he secured from the online catalogs of Cornell University, the University of Dayton, and Indiana University, purporting to show catalog dates of November 22, 1996, November 18, 1996, and November 18, 1996, respectively. *Id.* ¶¶ 28–31. Based on these records, Mr. Munford opines that Chien "was made available and accessible to the public by shortly after November 18, 1996." *Id.* ¶ 32.

Mr. Munford also provides additional testimony that he personally retrieved, or directed others at Pillsbury Winthrop Shaw Pittman LLP to retrieve, copies of the three references from various libraries for review, and that he personally reviewed those copies. Ex. 1031 ¶¶ 19–69. Mr. Munford testifies that at the time of his original Declaration, all libraries within his usual travel range had been closed, and the area where he lived was under a travel advisory due to the COVID-19 pandemic. *Id.* ¶ 17.

Mr. MacPherson testifies that (1) according to IEEE's standard practices, copies of proceedings were made available no later than the last day of the conference; (2) each of the three references are currently available for download from the IEEE digital library, IEEE Xplore; and (3) IEEE Xplore populates the information such as the date of publication and additional publication information, using the metadata associated with the publication. Ex. 1027 ¶¶ 1, 11–13.

Ms. McCarrier testifies that that she personally received a pdf copy of Zhang from the Library of the Missouri University of Science and

Technology, Rolla, Missouri, part of the University of Missouri system, and a pdf copy of Gupta from the Georgia Tech Library. Ex. 1028 ¶¶ 1, 6–9.

Mr. Schnell testifies that that he personally retrieved a physical copy of Gupta and Chien from the University of Texas library. Ex. 1029 ¶¶ 1–3.

c) Patent Owner's Contentions

In its Response, Patent Owner disputes the prior art status of Zhang, Gupta, and Chien. PO Resp. 19–29. Specifically, Patent Owner contends that Petitioner has not established that Zhang, Gupta, and Chen are printed publications because Petitioner has not shown that they were publicly accessible. *Id.* Patent Owner presents the identical argument that it presented in the Preliminary Response. *See* Paper 9, 26–35. Patent Owner does not address the printed publication issue in the Sur-reply. *See generally* Sur-reply.

First, Patent Owner argues that Dr. Gupta does not provide firsthand knowledge as to whether the references were actually distributed to attendees, and Petitioner has provided no evidence that the references were ever circulated at the conferences. PO Resp. 22–23.

Second, Patent Owner argues that Petitioner "provides no indication of how the IEEE Xplore website operated [in 2002, when the references were purportedly uploaded], how information was organized on the website, nor any evidence of the steps a user would have needed to take to find the *Zhang* or *Gupta* references once uploaded." PO Resp. 24. Patent Owner also asserts that "even today, a reasonable search of the website using key concepts from the '867 patent does not identify the asserted references." *Id.* Therefore, according to Patent Owner, Petitioner has not shown that the references were "meaningfully indexed such that an interested artisan exercising reasonable due diligence would have found" the references on the

IEEE Xplore website. *Id.* (citing *Acceleration Bay, LLC v. Activision Blizzard, Inc.*, 908 F.3d 765, 774 (Fed. Cir. 2018); *Samsung Elec. Co. v. Infobridge Pte. Ltd.*, 929 F.3d 1363, 1369 (Fed. Cir. 2019)).

Third, Patent Owner argues that Mr. Munford's testimony is insufficient to show that the references were publicly accessible for two reasons. PO Resp. 27–28. First, Patent Owner argues that the records Mr. Munford relies upon are insufficient because, with the exception of the Chien reference, the records only refer to the name or general subject matter of the conference, as opposed to the name of the author or title of the reference. *Id.* Therefore, according to Patent Owner, there is no way to determine the content or precise subjects covered using reasonable diligence. *Id.* Second, Patent Owner argues that Mr. Munford only shows when the references were catalogued, but the references would not be available until they are shelved and available for distribution. *Id.* at 28. According to Patent Owner, it is not reasonable to presume that the references must have been available shortly after cataloguing. *Id.*

d) Analysis

The determination of whether a document is a "printed publication" under 35 U.S.C. § 102 "involves a case-by-case inquiry into the facts and circumstances surrounding the reference's disclosure to members of the public." *In re Klopfenstein*, 380 F.3d 1345, 1350 (Fed. Cir. 2004). "Because there are many ways in which a reference may be disseminated to the interested public, 'public accessibility' has been called the touchstone in determining whether a reference constitutes a 'printed publication.'" *Jazz Pharm., Inc. v. Amneal Pharm., LLC*, 895 F.3d 1347, 1355 (Fed. Cir. 2018) (quoting *In re Hall*, 781 F.2d 897, 898–99 (Fed. Cir. 1986)).

"Whether a reference qualifies as a printed publication under § 102 is a legal conclusion based on underlying fact findings." *Acceleration Bay*, 908 F.3d at 772. "A reference is considered publicly accessible if it was 'disseminated or otherwise made available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, can locate it.'" *Id*. (citing *Jazz*, 895 F.3d at 1355–1356). "If accessibility is proved, there is no requirement to show that particular members of the public actually received the information." *Constant v. Advanced Micro-Devices, Inc.*, 848 F.2d 1560, 1569 (Fed. Cir. 1988).

After considering the arguments and weighing the evidence presented by the parties, we determine that Petitioner has shown by a preponderance of the evidence that Zhang, Gupta, and Chien were publicly accessible before the priority date, and, therefore qualify as prior art under 35 U.S.C. § 102(b).

Zhang, on its face, is a paper from the *Proceedings of the International Conference on Computer Design, VLSI in Computers and Processors*, which took place on October 12–15, 1997, in Austin, Texas, and was sponsored by the IEEE Computer Society Technical Committee on Design Automation, IEEE Circuits and Systems Society. Ex. 1003, 1. The copyright page indicates a copyright date of 1997 by IEEE, ISBN numbers, IEEE order plan catalog numbers, and addresses in California and New Jersey where additional copies may be ordered. *Id.* at 2. The Table of Contents for the Proceedings indicates over 700 pages of papers associated with various sessions, where Zhang is at pages 150–156. *Id.* at 3–18. At the bottom of the first page of the Zhang paper, it states "1063-6404/97 $10.00 © 1997 IEEE." *Id.* at 12.

Gupta, on its face, is a paper from the *Proceedings of the IEEE Computer Society Workshop on VLSI 2000, System Design for a System-on-Chip Era*, which took place on April 27–28, 2000, in Orlando, Florida, and was sponsored by the IEEE Computer Society Technical Committee on VLSI. Ex. 1004, 3. The copyright page indicates a copyright date of 2000 by IEEE, ISBN numbers, IEEE order plan catalog numbers, and addresses in California and New Jersey where additional copies may be ordered. *Id.* at 4. The Table of Contents for the Proceedings indicates over 150 pages of papers associated with various sessions, where Gupta is at pages 75–80. *Id.* at 5–12. At the bottom of the first page of the Gupta paper, it states "0-7695-0534-1/00 $10.00 © 2000 IEEE." *Id.* at 8.

Chien, on its face, is a paper from the *Proceedings of the Frontiers '96, The Sixth Symposium on the Frontiers of Massively Parallel Computing*, which took place on October 27–31, 1996, in Annapolis, Maryland, and was sponsored by the IEEE Computer Society. Ex. 1005, 1. The copyright page indicates a copyright date of 1996 by IEEE, ISBN numbers, IEEE order plan catalog numbers, and addresses in California and New Jersey where additional copies may be ordered. *Id.* at 2. The Table of Contents for the Proceedings indicates over 350 pages of papers associated with various sessions, where Chien is at pages 336–345. *Id.* at 3–16. At the bottom of the first page of the Chien paper, it states "1088-4955/96 $5.00 © 1996 IEEE." *Id.* at 7.

Although not dispositive, there is relevant evidence that supports our finding that each of Zhang, Gupta, and Chien were publicly accessibility prior to the critical date. Specifically, as set forth above, each bear multiple conventional indicia of publication and such as a copyright date, ISBN number and IEEE order plan catalog number, price, indicia of publication by

an established publisher, IEEE, and instructions for ordering additional copies, all of which are relevant evidence supporting a finding of public accessibility. *See Nobel Biocare Services AG v. Instradent USA, Inc.*, 903 F.3d 1365, 1378 (Fed. Cir. 2018) ("[a]lthough the ABT Catalog's date is not dispositive of the date of public accessibility, its date is relevant evidence"); *see also VidStream LLC v. Twitter, Inc.*, 981 F.3d 1060, 1065 (Fed. Cir. 2020) ("When there is an established publisher there is a presumption of public accessibility as of the publication date.").

In our Institution Decision, we agreed with Patent Owner that Dr. Gupta's testimony as to whether the papers were distributed to participants at the conference was speculative, because he had not provided firsthand knowledge of whether that had actually happened. Dec. 41. In his original Declaration, Dr. Gupta testifies that he "believe[s]" that the papers were distributed to the conference attendees prior to or during the relevant conference, yet in his Supplemental Declaration, Dr. Gupta testifies that he "confirm[s] based on my personal knowledge that [the] paper[s] . . . [were] included in the IEEE printed publication that was distributed to the conference attendees during the conference." Ex. 1010 ¶¶ 18, 19, 21, 22, 24, 25; Ex. 1030 ¶¶ 5, 8, 11. At his deposition, he testified that his testimony changed because "it was a recollection at that time," but later had become certain. Ex. 1039, 11:7–13. We credit Dr. Gupta's deposition testimony which is corroborated based upon Mr. MacPherson's testimony that "[i]n accordance with IEEE's standard practices, copies of the proceedings were made available no later than the last day of the conference." *See* Ex. 1010 ¶¶ 18, 21, 24; Ex. 1030 ¶ 3, 5, 8, 11; Ex. 1027 ¶¶ 11–13. We also find that Dr. Gupta's refreshed memory is supported by his experience in attending conferences sponsored by IEEE. Ex. 1030 ¶ 2. This testimony by Dr. Gupta

and Mr. MacPherson, supporting that the three papers were distributed to participants at the conference, is unrebutted.

We also determine that the MARC records from the various libraries provided by Mr. Munford for Zhang, Gupta, and Chien, and the accompanying testimony by Mr. Munford, including his knowledge as to standard library practices relating to MARC records (*e.g.*, Ex. 1012 ¶¶ 12, 13; Ex. 1031 ¶¶ 13–14), is sufficient evidence to show that Zhang, Gupta, and Chien were cataloged and publicly accessible. *See* Ex. 1012 ¶¶ 12–32; Ex. 1031 ¶¶ 19–69. Mr. Munford provides extensive, unrebutted testimony pertaining to the authenticity of the papers reviewed that were retrieved from the libraries, as well as the MARC records associated with those papers. *Id.* We credit Mr. Munford's testimony that Zhang, Gupta, and Chien, in their entirety, were properly catalogued and publicly available. *Id.* For example, Mr. Munford testifies that all four MARC records for the four copies of Zhang that were retrieved from four different libraries accurately describe its title, publisher, and ISBN, match the number of pages in the various copies he retrieved, and contain "008" fields that "designate the date of record creation" with an entry of "November 18, 1997" or "August 1, 2002." *See* Ex. 1031 ¶¶ 33–34. In addition, Mr. Munford testifies that the copy of Zhang that was retrieved from the University of Missouri library by Ms. McCarrier has a handwritten date of "11-25-97" on the top of the first table of contents page, one week after the date on the MARC record. *Id.* ¶ 35. Mr. Munford provides similar testimony for the Gupta and Chien papers. *Id.* ¶¶ 36–69.

Patent Owner does not dispute that Zhang, Gupta, and Chien were available on the IEEE website as of August 6, 2002. *See* PO Resp. 25 ("All Petitioner has shown, is that the references purportedly were placed on a

website in 2002 . . . ."). Rather, Patent Owner argues that Petitioner has not shown that the references were "meaningfully indexed in a way that they might be located using reasonable diligence." *Id.* at 21; *see also id.* at 31. However, "indexing is not 'a necessary condition for a reference to be publicly accessible'; it is but one among factors that may bear on public accessibility." *Voter Verified, Inc. v. Premier Election Solutions, Inc.*, 698 F.3d 1374, 1380 (Fed. Cir. 2012) (citing *In re Lister*, 583 F.3d 1307, 1312 (Fed. Cir. 2009)). Thus, "while often relevant to public accessibility, evidence of indexing is not an absolute prerequisite to establishing online references . . . as printed publications within the prior art." *Id.* Nonetheless, Mr. Munford provides unrebutted testimony and supporting evidence regarding the indexing provided on the IEEE website in the 2002 timeframe, which we credit. Ex. 1031 ¶¶ 15–16, 70–87.

Accordingly, we find that the totality of the evidence provided by Petitioner, including (1) the indicia of publication and public accessibility on the face of each of the references; (2) IEEE's standard practices to distribute papers of the proceedings no later than the last day of the conference; (3) the availability of the references on the IEEE website; and (4) the cataloguing of the Proceedings for each of the papers in various libraries, demonstrates by a preponderance of the evidence that each of Zhang, Gupta, and Chien were publicly accessible prior to the priority date of June 18, 2003.

E. Ground 1: Obviousness Over Zhang and Gupta

Petitioner contends claims 1, 2, 4–8, and 13–19 would have been obvious over the combination of Zhang and Gupta, and relies on the Declaration of Dr. Stanley Shanfield, Ph.D. (Ex. 1006) in support of its contentions. Pet. 5, 28–87. After reviewing the entire record developed at trial, as explained below, we determine that Petitioner has shown, by a

preponderance of the evidence, that claims 1, 2, 4–8, and 13–19 are unpatentable over the combination of Zhang and Gupta.

1. Patent Owner's General Arguments as to Zhang and Gupta

Patent Owner makes two general arguments as to Zhang and Gupta. First, Patent Owner argues that Zhang and Gupta teach away from the claimed invention. *See* PO Resp. 16 ("*Zhang* explicitly mentions that the main application remains in software (executed on the CPU) and programmable logic is used only for hardware adaptations (that remain static for the duration of a specific application run), thus teaching away from the invention of the '867 patent."); 19 ("Like *Zhang*, in *Gupta*, the main application remains in software (executed on the CPU), and programmable logic is used only for hardware adaptations that remain static for the duration of a specific application run, thus teaching away from the invention of the '867 patent."); *see also* Sur-reply 1–2. In support of its argument, Patent Owner relies on the testimony of Dr. Mangione-Smith. PO Resp. 16, 19 (citing Ex. 2028 ¶ 70).

A reference teaches away from a claimed invention if it "criticize[s], discredit[s], or otherwise discourage[s]" modifying the reference to arrive at the claimed invention. *In re Fulton*, 391 F.3d 1195, 1201 (Fed. Cir. 2004). Neither Patent Owner nor Dr. Mangione-Smith identifies where Zhang and Gupta criticize, discredit, or otherwise discourage the claimed invention. Rather, Patent Owner and Dr. Mangione-Smith argue that the references *do not teach* the disclosure in the '867 patent. For example, Dr. Mangione-Smith provides testimony that Zhang does not "instantiate an algorithm as hardware." Ex. 2028 ¶¶ 68, 70 ("[t]he main application algorithm . . . are not instantiated as hardware"; "the algorithms that comprise the application discussed in *Zhang* also 'remain in software'"). This is not sufficient to

show that the references teach away from the claimed invention. *See Galderma Labs., L.P. v. Tolmar, Inc.*, 737 F.3d 731, 738 (Fed. Cir. 2013) ("A reference does not teach away, . . . if it merely expresses a general preference for an alternative invention but does not criticize, discredit, or otherwise discourage investigation into the invention claimed."). We, therefore, find this argument unavailing.

Second, Patent Owner argues that Zhang is not an enabling prior art reference. *See* PO Resp. 15. However, "[u]nder § 103, . . . a reference need not be enabled; it qualifies as a prior art, regardless, for whatever is disclosed therein." *Amgen Inc. v. Hoechst Marion Roussel, Inc.*, 314 F.3d 1313, 1357 (Fed. Cir. 2003) (citing *Symbol Techs., Inc. v. Opticon, Inc.*, 935 F.2d 1569, 1578 (Fed. Cir. 1991) ("enablement of the prior art is not a requirement to prove invalidity under § 103")); *Reading & Bates Constr. Co. v. Baker Energy*, 748 F.2d 645, 652 (Fed. Cir. 1984). We, therefore, also find this argument is unavailing.

2. Analysis of Independent Claim 1

a) Limitation of the preamble: "A reconfigurable processor that instantiates an algorithm as hardware comprising"

(1) The Parties' Contentions

The parties agree that the preamble limits claim 1. *See* Pet. 28–30; PO Resp. 30, 33. Petitioner asserts that Zhang discloses the preamble. Pet. 28–30. For example, Petitioner relies on Zhang's architecture "that integrates small blocks of programmable logic into key elements of a baseline architecture, including processing elements, components of the memory hierarchy, and the scalable interconnect, to provide *architectural adaptation* – the customization of architectural mechanisms and policies to

match an application." *Id.* at 28 (citing Ex. 1003, 13) (bolding omitted). Petitioner asserts that Zhang's Figure 2 illustrates "a reconfigurable processor that includes a CPU, cache memory, and processor main memory, each integrated with blocks of programmable logic over a flexible interconnect." *Id.* at 28–29 (citing Ex. 1003, 12–15, Figs. 2 and 4). Petitioner contends, with supporting testimony from Dr. Shanfield, that a person of ordinary skill in the art "would understand Zhang's processor architecture, which integrates 'blocks of programmable logic into key elements' customized to 'match an application,' to disclose a computing device that contains reconfigurable components that can, through reconfiguration, instantiate an application's algorithm(s) as hardware." *Id.* at 30 (citing Ex. 1006 ¶ 131).

Patent Owner argues that (1) Zhang's processor is a conventional CPU, not a reconfigurable processor; and (2) Zhang's processor does not "instantiate an algorithm as hardware." PO Resp. 33–37; Sur-Reply 2–5.

According to Patent Owner, Zhang "uses programmable logic (FPGA) as means to deliver data for use by [a] conventional CPU." PO Resp. 34. Patent Owner contends that Figure 2 of Zhang shows the use of a CPU and small pockets of reprogrammable logic blocks, not a reconfigurable processor. *Id.* Patent Owner's annotated Figure 2 is reproduced below:

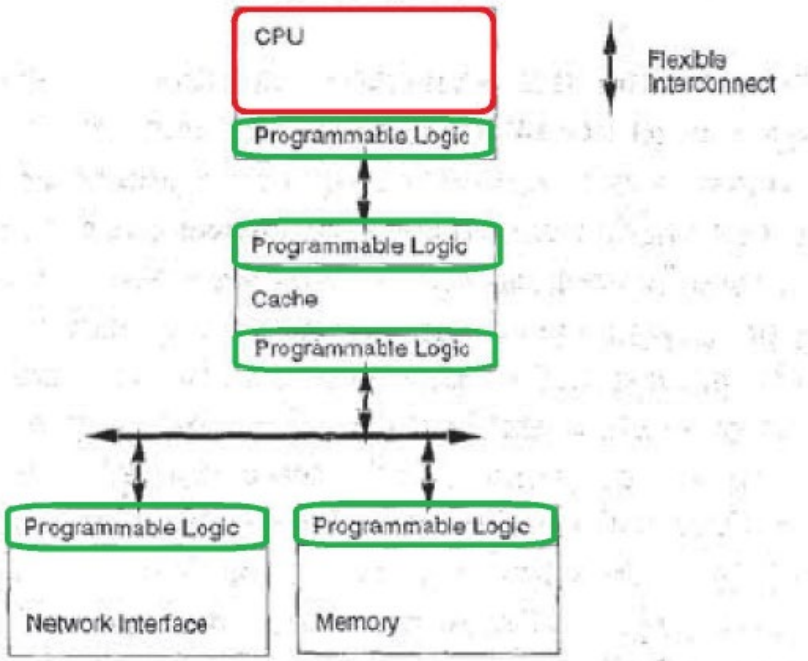

Figure 2. An Architecture for Adaptation

Figure 2 of Zhang, above, depicts an architecture for adaptation, with boxes depicting a CPU, Network Interface, Memory, and Cache, each with corresponding "Programmable Logic." Patent Owner annotates Figure 2, drawing a red line around the CPU, and green lines around the "Programmable Logic." Patent Owner contends that this "clearly shows the use of a CPU (indicated in red) that executes all algorithms that comprise the application that remain in [sic] software, and only small pockets of reprogrammable logic blocks (indicated in green)." PO Resp. 35.

Patent Owner argues that Zhang's "programmable logic is only implemented in small blocks for specific intermediate purposes, such as the interface between CPU and cache, the network interface, or the memory interface." *Id.* at 36 (citing Ex. 1003, Fig. 2). For this reason, Patent Owner contends that Zhang's architecture is "incapable of 'instantiating an

algorithm as hardware.'" *Id.* (citing Ex. 2028 ¶ 42). Patent Owner further argues that "the CPU hardware is static and cannot be adapted to a specific application . . . [and] [t]his is why *Zhang* teaches that its entire application must remain in software." *Id.* Patent Owner relies on Dr. Mangione-Smith's testimony in support of its arguments that Zhang does not "instantiate an algorithm as hardware." *E.g.*, *id.* at 35 (citing Ex. 2028 ¶ 70).

Patent Owner also argues that Zhang "fails to disclose any capability to reconfigure its small reprogrammable logic blocks on a per application basis." *Id.* at 37 (citing Ex. 2028 ¶ 68). Patent Owner argues that, at most, Zhang's architecture could "create application specific hardware instantiations for the cache, network, and memory interfaces," but "the CPU hardware is static and cannot be adapted to a specific application." *Id.* at 36. Patent Owner also argues that Zhang "discloses only relatively minor hardware architectural adaptations . . . [which] is quite different from the physical adaptation of the entire reconfigurable processor to the specific needs of a particular application taught in the '867 patent." *Id.* at 36–37 (citing Ex. 1003, 15–16; Ex. 2028 ¶¶ 66–67). Patent Owner contends that "these relatively minor hardware adaptations are not application specific, but application class specific." *Id.* at 37 (citing Ex. 2028 ¶ 42).

In Reply, Petitioner argues that Patent Owner introduces requirements that are not commensurate with the scope of the claim or the agreed construction of a "reconfigurable processor." Reply 10. Petitioner further argues that Patent Owner's arguments conflict with the disclosure in the '867 patent and the prosecution history. *Id.* at 11.

Petitioner also argues that Patent Owner mischaracterizes Zhang as "disclosing a 'static' 'conventional' CPU incapable of instantiating an algorithm as hardware." *Id.* at 12. According to Petitioner, Patent Owner's

annotated Figure 2 is incorrect, and Patent Owner is "wrong to focus its analysis only on Zhang's CPU-labeled block alone rather than the reconfigurable processor that Petitioner and Dr. Shanfield identified." *Id.* Petitioner's annotated Figure 2 of Zhang is reproduced below.

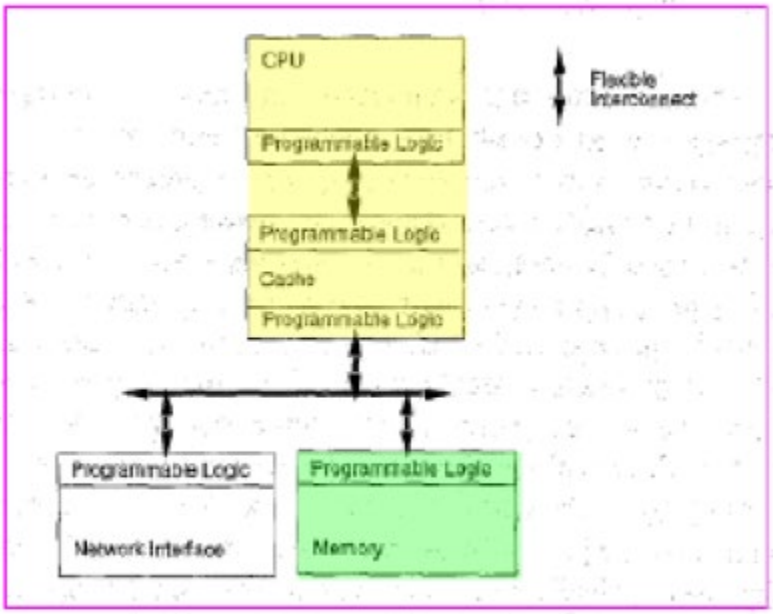

Figure 2. An Architecture for Adaptation

Figure 2 of Zhang, above, depicts an architecture for adaptation, with boxes depicting a CPU, Network Interface, Memory, and Cache, each with corresponding "Programmable Logic." Petitioner annotates Figure 2 with a pink box around the entirety of the CPU, Network Interface, Memory, and Cache boxes. Reply 12. Petitioner further argues that "Zhang expressly discloses integrating programmable logic *into the processing elements*, which Zhang identifies *separately* from the memory hierarchy, scalable interconnect, and network interface." *Id.* at 13 (citing Ex. 1003, 13; Fig. 2). According to Petitioner, "Zhang's processing elements–whether labeled 'processor' or 'CPU'–are indeed reconfigurable." *Id.*

Additionally, Petitioner argues that Zhang's reconfigurable processor can instantiate an algorithm as hardware, because it teaches "optimizing matrix multiplication computations using the customization provided by its programmable logic to improve computational processing." *Id.* at 14 (citing Ex. 1003, 12). Petitioner contends that "[n]othing about [Zhang's] statement [that the 'application remains in software'] precludes compiling one (or more) of an application's algorithms and instantiating it as hardware." *Id.* According to Petitioner, this is consistent with the teachings in the '867 patent. *Id.* Petitioner further contends that Zhang is "merely distinguish[ing] its approach from conventional co-processing architectures that must repartition hardware and software functionality and reimplement the co-processing hardware every time a new application is run." *Id.* at 15 (citing Ex. 1003, 12).

In the Sur-reply, Patent Owner contends that "[t]here is a stark difference between including some reconfigurable components in a conventional CPU as compared to using reconfigurable hardware to actually 'instantiate an algorithm as hardware' as claimed." Sur-reply 2–3. Patent Owner argues that in Zhang, "the processor running the main application is a conventional CPU, not a reconfigurable processor." *Id.* at 3. In support, Patent Owner again refers to Figure 2 of Zhang, and argues that "the reconfigurable logic is only used at the periphery of the conventional CPU." *Id.* According to Patent Owner, Figure 2 "leaves no doubts that the programmable logic is *next to* – not *contained in* – the CPU." *Id.* at 4. Patent Owner contends that the "programmable logic is only implemented in small blocks for specific intermediate purposes between components, such as the interface between CPU and cache, the network interface, or the memory interface." *Id.* Patent Owner further contends, with supporting

testimony from Dr. Mangione-Smith That Zhang's architecture is incapable of 'instantiating an algorithm as hardware.'" *Id.* (citing Ex. 2028 ¶ 42).

(2) Analysis

As discussed above, the parties agree, and we determine, that a "reconfigurable processor" is construed as "a computing device that contains reconfigurable components such as FPGAs and can, through reconfiguration, instantiate an algorithm as hardware." *See* Ex. 1001, 5:26–29. For the reasons advanced by Petitioner and as discussed below, we agree with Petitioner and find that Zhang teaches the claimed "reconfigurable processor." *See* Pet. 28–30; Reply 11–15. Patent Owner's arguments are premised on Zhang teaching a conventional, static processor and not a reconfigurable processor. In presenting these arguments, Patent Owner focuses on Zhang's CPU in Figure 2. PO Resp. 35–36; Sur-Reply 3–4. But Petitioner relies not only on the CPU to teach the "reconfigurable processor," but on the entire architecture depicted in Figure 2, which "includes a CPU, cache memory, and processor main memory, each integrated with blocks of programmable logic over a flexible interconnect." Pet. 29. This is consistent with the construction of the "reconfigurable processor." Patent Owner's arguments do not squarely address Petitioner's contentions, or explain why Zhang does not teach the "reconfigurable processor," *as construed by the parties* and which we have adopted.

Patent Owner's focus on the CPU of Figure 2 – and attempt to separate it from the rest of the components – carries through its arguments. For example, Patent Owner admits that Zhang's architecture can "create application specific hardware instantiations for the cache, network, and memory interfaces," but attempts to distinguish Zhang's CPU as "static and cannot be adapted to a specific application . . . [which] is why *Zhang* teaches

that its entire application must remain in software." PO Resp. 36; *see also* Ex. 2028 ¶ 68. But Zhang's disclosure does not support such a distinction.

Zhang states "We propose an architecture that ***integrates small blocks of programmable logic into*** *key elements of a baseline architecture,* ***including processing elements***, components of the memory hierarchy, and the scalable interconnect, to provide *architectural adaptation* – the *customization of architectural mechanism and policies to match an application*." Ex. 1003, 13 (first and third emphases added). That is, Zhang states that the small blocks of programmable logic are ***integrated into*** key elements of a baseline architecture, including processing elements and components of the memory hierarchy, and this is depicted in Figure 2. Ex. 1003, 13; Fig. 2 (showing programmable logic integrated with a CPU, cache, memory, and network interface). Given this disclosure in Zhang, along with the inclusion of "programmable logic" in the "CPU," we do not agree that Zhang's CPU is "static and cannot be adapted to a specific application." *See* PO Resp. 36. Nor, for the same reasons, do we agree that Zhang's programmable logic is not "*contained in*" the CPU (or the memory), but is "*next to*" it, as Patent Owner argues. *See* Sur-reply 4. As Petitioner persuasively argues, Patent Owner's arguments regarding Zhang are not commensurate with the parties' agreed construction of the term "reconfigurable processor" or the claim language. *See* Reply 10–11.

Patent Owner contends that Zhang's statement that "the entire application remains in software" evidences that Zhang's architecture does not instantiate an algorithm as hardware. *See* PO Resp. 34–37; Ex. 2028 ¶ 67. Zhang states: "In addition, because the entire application remains in software while the underlying hardware is adapted for system performance,

our approach improves over co-processing architectures by preserving machine usability through software." *See* Ex. 1003, 14.

Dr. Mangione-Smith testifies that a person of ordinary skill in the art "would understand [Zhang's] disclosure to indicate that system performance is adapted *transparently* to the application and its algorithms, *i.e.*, that these adaptations do not teach that any part of the application actually run[s] in *Zhang*, that is any of the algorithms that actually make up the application that remains in software, are instead instantiated in hardware." Ex. 2028 ¶ 70 (emphasis in original). He states that "this fact is confirmed by the last part of the same cited sentence, [which states] our approach improves over co-processing architectures *by preserving machine usability through software*." *Id.* According to Dr. Mangione-Smith, "[m]achine useability through software aims to retain the flexibility of implementing software on general purpose hardware such as the conventional CPU used by *Zhang*." *Id.*

Patent Owner's argument as to the instantiation relies on a finding that Zhang's CPU is a static, conventional processor. *See* PO Resp. 36 ("This architecture is incapable of 'instantiating an algorithm as hardware'" (citing Ex. 2028 ¶ 42); "[T]he CPU hardware is static and cannot be adapted to a specific application . . . [t]his is why *Zhang* teaches that its entire application must remain in software."). As stated above, we reject that argument. Moreover, Patent Owner admits that Zhang's architecture could "*create application specific hardware instantiations* for the cache, network, and memory interfaces." PO Resp. 36 (emphasis added).

We do not find Dr. Mangione-Smith's testimony availing, given the disclosure in Zhang. Specifically, we do not agree that Zhang's statement that "our approach improves over co-processing architectures by preserving

machine usability through software" supports that the architecture does not instantiate an algorithm as hardware. Rather, we find persuasive Petitioner's contention that in that statement Zhang is simply distinguishing itself from conventional co-processing architectures. *See* Reply 15. We do not agree that this disclosure in Zhang precludes "instantiating an algorithm as hardware," given that Zhang extensively refers to "customizing architectural mechanisms and policies to match an application," "application-specific hardware assists," "application-specific optimizations," "architectural customization," "architectural adaptation," and "the underlying hardware is adapted." *E.g.*, Ex. 1003, 12–14.

Patent Owner draws a distinction between hardware applications that are application specific and hardware applications that are application class specific, arguing that Zhang teaches the latter and, therefore, does not "disclose any capability to reconfigure its small reprogrammable logic blocks on a per application basis." PO Resp. 37. Patent Owner's argument is not supported by Zhang's disclosure. *See, e.g.*, Ex. 1003, 12 ("the goal of customizing architectural mechanisms and policies *to match an application*" and "*application-specific hardware assists* and policies can provide substantial improvements in performance *on a per application basis*"); 13 ("the customization of architectural mechanisms and policies *to match an application*"), 14 ("[a]rchitectural adaptation provides the mechanisms for *application-specific hardware assists*") (emphasis added).

In summary, we agree with Petitioner, and find credible Dr. Shanfield's testimony that a person of ordinary skill in the art "would understand Zhang's processor architecture that integrates 'programmable logic' customized to 'match an application' to disclose a computing device that, through reconfiguration, instantiates an algorithm as hardware. *See* Ex.

1006 ¶ 131. Moreover, as Dr. Shanfield points out, the architecture's name is "MORPH" - (**M**ultiprocess**O**r) with **R**econfigurable **P**arallel **H**ardware. *Id.*; *see* Ex. 1003, 14.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang teaches the preamble.

b) Limitation 1(a): "a first memory having a first characteristic memory bandwidth and/or memory utilization"

Petitioner asserts that Zhang discloses this limitation. Pet. 30–31. Specifically, Petitioner relies on Zhang's L1 cache, contending that the L1 cache is a first memory with a first characteristic memory bandwidth because "L1 cache has a transfer rate (memory bandwidth) of 16B/5 cycles." *Id.* (citing Ex. 1003, 15 Table 1, Fig. 4). Petitioner also relies on Zhang's Figure 5, contending the "Cache" shown in Figure 5 also teaches a first memory. *Id.* at 52 (citing Ex. 1003, Fig. 5).

Patent Owner does not specifically respond to these arguments. *See generally* PO Resp.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang teaches limitation 1(a).

c) Limitation 1(b): "a data prefetch unit coupled to the first memory"

Petitioner contends that Zhang in combination with Gupta discloses a data prefetch unit as recited in limitation 1(b). Pet. 32–36. Petitioner relies on Zhang's data prefetcher that is coupled to LI cache and L2 cache. *Id.* at 32–33 (citing Ex. 1003, 15, Fig. 4; Ex. 1006 ¶ 138). With supporting testimony from Dr. Shanfield, Petitioner asserts that "Zhang teaches a data

prefetcher that (i) is implemented in reconfigurable (programmable) logic; (ii) is a set of logic that performs a specific operation; (iii) is coupled to the first memory (L1 cache); and (iv) moves data between a first memory (L1 cache) and a second memory (L2 cache or main memory), which together are a collection of memories constituting a 'memory hierarchy.'" *Id.* at 34 (citing Ex. 1006 ¶ 138).

Petitioner contends Gupta discloses a prototype implementation of the data prefetch unit disclosed in Zhang, also coupled to a first memory (L1 cache). Pet. 34–36 (citing Ex. 1004, 9, 11, Fig. 1). Petitioner contends that "Gupta uses application-specific prefetching to move computational data between members of the memory hierarchy." *Id.* at 34 (citing Ex. 1004, 9). Petitioner asserts that Figure 1 of Gupta "teaches an AMRM prototype board that includes a 'prefetch unit' that can move data 'coming from or going to the L1 cache,' such as data from the L2 cache or main memory." *Id.* With supporting testimony from Dr. Shanfield, Petitioner asserts that "Gupta's prefetch unit (i) is implemented in reconfigurable (programmable) logic; (ii) is a set of logic that performs a specific operation; (iii) is coupled to the first memory (L1 cache); and (iv) moves data between a first memory (L1 cache) and a second memory (L2 cache or main memory), which together are a collection of memories constituting a 'memory hierarchy.'" *Id.* at 35 (citing Ex. 1004, 10; Ex. 1006 ¶ 139).

Petitioner contends an ordinarily skilled artisan would have been motivated to combine the teachings of Zhang and Gupta because (i) Gupta discloses a prototype implementation of the architecture that Zhang disclosed, (ii) Zhang teaches the general architecture for data prefetching in a reconfigurable processor, and Gupta teaches a specific prototype implementation of that architecture, including a prefetch unit positioned

between the L1 cache and the rest of the system accessible in parallel with the L2 cache in a way that allows it to accept and supply data coming from or going to the L1 cache, and (iii) "[i]t would therefore have been obvious for a [person of ordinary skill in the art] to look to Gupta to better understand one way to implement the data prefetch architecture taught by Zhang." Pet. 35–36 (citing Ex. 1003, 16; Ex. 1004, 9, 11; Ex. 1006 ¶ 140).

Patent Owner argues that the combination of Zhang and Gupta does not teach "a data prefetch unit" because "[t]he basic idea of prefetching data is to obtain data before it is needed by the application . . . [and] [s]ince the *Zhang* prefetch unit is triggered only by a cache miss, it does not have the ability to initiate a data transfer in advance of the requirement for data by computational logic, and thus does not *pre*-fetch any data at all." PO Resp. 38 (citing Ex. 1003, Fig. 4, 15–16; Ex. 2028 ¶¶ 47, 81, 87). Patent Owner does not contend that Petitioner has not provided sufficient rationale to combine Zhang and Gupta. *See* PO Resp. 38–39.

Dr. Mangione-Smith testifies that a person of ordinary skill in the art "would understand that [in the agreed construction of a 'data prefetch unit'] 'movement of data between members of a memory hierarchy' in a data *prefetch* unit implicitly refers to moving data *before* it is needed, *i.e.*, it is not '*fetched when needed*' but '*pre-fetched before it is needed*.'" Ex. 2028 ¶ 87. According to Patent Owner "[t]he fact that *Zhang* refers to its disclosure as a 'prefetch unit' does not change its actual post-fetch (*i.e.*, based only on a cache miss) operation" and "[t]he mere fact that some of the data may be used later does not teach a [person of ordinary skill in the art] that *Zhang's* prefetch unit intentionally moves data before it is needed." PO Resp. 39.

In Reply, Petitioner contends that neither the agreed construction of a "data prefetch unit" nor the claim language precludes prefetching after a

read miss. Reply 17. For example, Petitioner contends that the '867 patent contemplates a prefetcher that operates after computational logic. *Id.* (citing Ex. 1001, 7:36–41; Ex. 1044, 64:12–67:3). Petitioner further argues that both Zhang and Gupta explicitly refer to "prefetching," and "prefetching" has existed since the mid-1960s, as Patent Owner and Dr. Mangione-Smith admit. *Id.* (citing Pet. 32–36; Ex. 1003, 15–17; Ex. 1004, 9–11; Ex. 1016, 16; Ex. 1044, 63:23–64:11). According to Petitioner, a person of ordinary skill in the art "would not plausibly read these express references to 'prefetch' as 'actually not prefetching data at all.'" *Id.* Moreover, Petitioner argues that Patent Owner focuses on Zhang's first case study, but Zhang's second cases study "teaches *prefetching* by packing/gathering only the used fields of the matrix element structure and forwarding them directly to cache." *Id.* at 17–18 (citing Pet. 36–40; 49–53; Ex. 1003, 16).

In Sur-reply, Patent Owner contends that "[t]he fact that the data prefetch unit of the '867 Patent can operate 'prior to, in parallel with, or after' the computational logic does not support that the computational logic must wait for data to be fetched by the *pre-fetch* unit." Sur-reply 5. According to Patent Owner, "[t]he statement relates only to the operation of the prefetch unit, not to the data that is fetched at a given time, which . . . must always be fetched before it is needed." *Id.* (citing Ex. 1001, 7:66–8:2). Patent Owner further argues that Zhang does not teach that its prefetch unit "intentionally moves data before it is need." *Id.* at 6. Patent Owner also argues that Zhang's second case study "discloses only speculative fetching triggered by a read-miss." *Id.*

For the reasons advanced by Petitioner and as set forth below, we agree with Petitioner that Zhang and Gupta teach this limitation. As discussed above, the parties agree that a "data prefetch unit" is "a functional

unit that moves data between members of a memory hierarchy. The movement may be as simple as a copy, or as complex as an indirect indexed strided copy into a unit stride memory." *See* Section III.C.1. Patent Owner's arguments are not commensurate with the parties' agreed construction. For example, Patent Owner contends that the data must be "intentionally" moved and data is "move[d] data before it is needed." PO Resp. 39; Sur-reply 6; Ex. 2028 ¶ 87. Under the agreed construction, which we have adopted, Petitioner need only show that the Zhang-Gupta combination discloses "a functional unit that moves data between members of a memory hierarchy." As Petitioner has shown, Zhang teaches a "prefetcher" that moves data between L1 Cache and L2 Cache, and Gupta teaches a prototype of that "prefetcher" implementation. *See, e.g.*, Pet. 32–36; Ex. 1003, Fig. 4; Ex. 1004, 9 ("The basic machine architecture supports application-specific cache organization and policies . . . prefetching and dynamic cache structures . . . that optimize the movement and placement of application data through the memory hierarchy"). Petitioner has also provided sufficient articulated reasoning, supported by rationale underpinning, to combine the teachings of Zhang and Gupta, which we adopt. Pet. 35–36.

Moreover, both Zhang and Gupta explicitly refer to "prefetching," a "prefetcher," and/or a "prefetch unit." *See, e.g.*, Ex. 1003, 15–16; Ex. 1004, 9–11. In light of this disclosure, we do not find credible Patent Owner's contention that Zhang does not teach "prefetching." In addition, the fact that Zhang's prefetching may be triggered by read misses does not support that Zhang still does not obtain data before it is needed. For example, Zhang explicitly discloses that it "targets records spanning multiple cache lines and . . . prefetches all fields of a matrix element structure whenever some field of

the element is accessed," and "targets pointer fields that are likely to be traversed when their parent structures are accessed." Ex. 1003, 15. Neither the claim language, nor the definition of the "data prefetch unit", precludes prefetching after a cache miss or read miss.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang in combination with Gupta teaches limitation 1(b), and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

> *d) Limitation 1(c): "wherein the data prefetch unit retrieves only computational data required by the algorithm from a second memory of second characteristic memory bandwidth and/or memory utilization and places the retrieved computational data in the first memory"*

Petitioner asserts that Zhang in combination with Gupta discloses this limitation. Pet. 36–41. Petitioner contends Zhang's L2 cache or main memory teach the claimed "second memory of second characteristic memory bandwidth and/or memory utilization." *Id.* at 36–38 (citing Ex. 1003, pp. 15–16, Figs. 4, 5; Ex. 1006 ¶¶ 141–142). Petitioner further contends Zhang's prefetcher retrieves data from the L2 cache or main memory (second memory) and places the retrieved data in cache memory closer to the processor, L1 cache or L2 cache (first memory), by writing it there. *Id.* (citing Ex. 1003, 15, Fig. 4; Ex. 1006 ¶¶ 141–142).

Petitioner contends that Zhang's prefetcher "retrieves *only computational data required by the algorithm* from a second memory . . . and places the retrieved computational data in the first memory." Pet. 36–

40. Specifically, Petitioner asserts "[i]n its sparse matrix-matrix multiplication example, Zhang teaches that the data prefetcher is configured to retrieve only the matrix data elements required for a given computation using dynamic scatter-gather memory access operations to improve memory bandwidth." *Id.* at 39 (citing Ex. 1003, 16); Pet. 39–40 (citing Ex. 1003, 15, Fig. 4; Ex. 1006 ¶ 144). Petitioner contends that "whether placed directly in L1 cache or placed first in L2 cache, Zhang's prefetcher retrieves only the computational data required by the instantiated algorithm and places that data into the L1 cache memory." *Id.* at 40.

Petitioner contends Gupta's prefetcher, which is a prototype implementation of the data prefetch unit disclosed in Zhang, is configured to prefetch computational data from the L2 cache (a second memory) or from main memory (a second memory) and move (write) the computational data into the L1 cache (a first memory). Pet. 40 (citing Ex. 1004, 11, Fig. 1; Ex. 1006 ¶ 145).

Petitioner further contends an ordinarily skilled artisan would have been motivated to combine the teachings of Zhang and Gupta because (i) Gupta discloses a prototype implementation of the architecture that Zhang disclosed, (ii) Zhang teaches a general architecture for data prefetching in a reconfigurable processor and the use of a scatter-gather memory access technique, and Gupta teaches a specific prototype implementation of that architecture and technique, including using "hardware assist *gather* in the memory controller" to enable efficient prefetching of data scattered throughout physical memory, and (iii) "[i]t would therefore have been obvious for a [person of ordinary skill in the art] to look to Gupta to better understand one way to implement the data prefetch architecture and scatter-

gather technique taught by Zhang." Pet. 40–41, 56–57 (citing Ex. 1003, 16; Ex. 1004, 9; Ex. 1006 ¶ 146).

Patent Owner argues that Zhang does not teach retrieval of "only computational data required by the algorithm." PO Resp. 40–44. Patent Owner does not address Gupta, or contend that Petitioner has not provided sufficient rationale to combine Zhang and Gupta. *See generally id*. According to Patent Owner, Zhang describes retrieval of computational data for two case studies that is "already optimized" because it contains "only non-zero matrix elements" while "all non-essential matrix elements (*i.e.*, matrix elements with the value of zero)" have already been eliminated by an optimized data storage scheme for a sparse matrix library. *Id.* at 40 (citing Ex. 1003, 14–15, Figs. 2 and 4; Ex. 2028 ¶ 74). Patent Owner argues "[t]ransfer of only required elements by virtue of all elements being required (*Zhang*) is materially different from transfer of only required elements by virtue of the prefetch unit being adapted to purposefully load 'only data required by the algorithm' (as in the '867 patent)" because "[in Zhang] it is not the hardware adaptation . . . which ensures that only data needed for computation is retrieved[, i]nstead, the hardware adaptation [of Zhang] simply transmits all data that is stored in the sparse matrix because all data is needed." *Id.* at 40–41.

Patent Owner also contends that Zhang's first and second case studies do not teach retrieval of "o*nly computational data required by the algorithm*." PO Resp. at 41–44. Dr. Mangione-Smith testifies that a person of ordinary skill in the art would understand that, due to Zhang's disclosure of moving fixed cache line sizes, "the entire last cache line is loaded, regardless of whether it is full or not," and, therefore, a person of ordinary skill in the art "would therefore not glean the teaching from *Zhang* that only

data required by the algorithm is loaded." *Id.* at 41 (citing Ex. 2028 ¶ 74). Patent Owner contends that Dr. Shanfield agrees that Zhang prefetches the entire row or column. *Id.* at 42 (citing Ex. 2029, 218:20–219:7). Patent Owner also argues that, in Zhang's second case study, a whole row (or column) of used fields is prefetched and packed into cache memory, but these additionally retrieved matrix element fields are for "potential (speculative) later use" and might or might not be used later, and, therefore, are not needed for the computation that triggered the prefetch operation. *Id.* at 43 (bold emphasis omitted).

In Reply, Petitioner contends that even if Zhang's data is "already optimized," this "does not preclude that prefetcher from retrieving or reading only computational data required by the algorithm." Reply 18. Petitioner also argues that Patent Owner's argument that Zhang's second case study retrieves elements for "potential (speculative) use" is contradicted by Zhang. *Id.* at 19 (citing Ex. 1003, 16). Petitioner argues that the combination's prefetcher "does *not* transfer a complete cache line" because its prefetcher packs/gathers "*only used fields* of matrix elements . . . and plac[es] only those fields into a split cache expressly '*for the prefetched matrix elements only*.'" *Id.* at 19–20. Petitioner relies on Figure 5 of Zhang, which Petitioner contends shows only matrix element values (Val1, Val2, Val3) retrieved from memory and packed into cache. *Id.* at 20 (citing ex. 1003, Fig. 5; Ex. 1006 ¶ 164; Pet. 54–55). Petitioner further asserts that Gupta discloses that "its prefetch hardware is combined with gather logic in the memory controller 'that works well with data structures *that do not quite fit into a single cache line*.'" *Id.* at 20 (citing Ex. 1004, 9). Petitioner also argues that Patent Owner mischaracterizes Dr. Shanfield's testimony, because Dr. Shanfield confirms that Zhang's second case study uses

adjustable cache line sizes, not fixed cache line sizes. *Id.* (citing Ex. 2029, 210:20–212:8). Finally, Petitioner argues that Dr. Mangione-Smith "admits his opinion is both (i) in direct contradiction to Zhang's explicit statements to the contrary," and "(ii) unsupported by any express statement in Zhang." *Id.* at 21 (citing Ex. 1044, 120:3–17; 120:25-121:24, 120:8–17, 124:22–125:6).

In Sur-reply, Patent Owner reiterates its arguments, and in particular, Dr. Mangione-Smith's testimony that "[i]t is highly unlikely that even in the efficient data storage scheme disclosed in *Zhang*, the last data elements of the matrix would completely fill up the last cache line perfectly. A [person of ordinary skill in the art] would understand from *Zhang's* disclosure that the entire last cache line is loaded, regardless of whether it is full or not." Sur-reply 7 (citing Ex. 2028 ¶ 74). Patent Owner also contends that Zhang only discloses cache line sizes of 32 or 64, while the '867 patent teaches "that the cache size is specifically adapted to always match the data requirements." *Id.* at 8. Finally, Patent Owner argues that the references do not support Petitioner's argument that "the combinations' prefetcher does *not* always transfer a complete cache line." *Id.* (citing Reply 19).

Patent Owner's arguments are unavailing. Petitioner relies on Zhang's second case study to teach this limitation. *See* Pet. 39–40. In the second case study, Zhang describes a sparse matrix-matrix multiplication by "architectural customization [that] *aims to send only used fields of matrix elements during a given computation* to reduce bandwidth requirement using dynamic scatter and gather" by

> prefetching of whole rows or columns using pointer chasing in the memory module and *packing/gathering of only the used fields of the matrix element structure*. When the root pointer of a column or row is accessed, the gather logic in the main memory

> module chases the row or column pointer *to retrieve different matrix elements and forwards them directly to the cache*.

Id. at 39; *see* Ex. 1003, 16, column 1, lines 36–42 (emphases added). For the reasons advanced by Petitioner, we find that Zhang's "packing/gathering of only the used fields of the matrix element structure" teaches "retrieves only computational data required by the algorithm." Pet. 36–41; Reply 18–21. We do not agree that the "optimized" data in Zhang has any effect on the foregoing disclosure in Zhang. Moreover, Petitioner has also provided sufficient articulated reasoning, supported by rationale underpinning, to combine the teachings of Zhang and Gupta, which we adopt. Pet. 53–54.

Dr. Mangione-Smith testifies that "[b]ecause of the efficient data storage scheme, the data is already optimized and effectively all data is retrieved" and that a person of ordinary skill in the art "would understand from *Zhang's* disclosure that the entire last cache line is loaded, regardless of whether it is full or not." Ex. 2028 ¶ 74. But this is not what Zhang states, and at his deposition, Dr. Mangione-Smith admitted as much. When asked if "in terms of what [Zhang] state[s] at that line in column 1, lines 36 to 39, you say that despite their express statement, it's not true; is that fair?", Dr. Mangione-Smith answered "Yeah. I tried to say that a couple of times now." Ex. 1044, 123:14–20; *see* 119:7–10 (identifying referring to page 16 of Zhang). In other words, despite what is expressly disclosed in Zhang – "*packing/gathering of only the used fields of the matrix element structure*" – Dr. Mangione-Smith takes the position that something different is happening. We do not find such testimony credible.

Further, Zhang's disclosure does not support Patent Owner's argument that "the additionally retrieved matrix element fields are prefetched for potential (speculative) later use," rather than those needed for

computation, as Zhang states it "*aims to send only used fields of matrix elements during a given computation*." PO Resp. 45 (emphasis omitted); Ex. 1003, 16. Thus, Zhang teaches that *only used fields of matrix elements during a given computation* are *sent* from the main memory (second memory) to the cache (first memory, *see* Zhang's Fig. 5), by prefetching using pointer chasing in the memory module (second memory) and packing/gathering only the used fields of the matrix element structure. *See* Ex. 1003, 16; *see* Pet. 39–40, 52.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang in combination with Gupta teaches limitation 1(c) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

e) Limitation 1(d): "wherein the data prefetch unit operates independent of and in parallel with logic blocks using the computional [sic] data"

Petitioner asserts that Zhang in combination with Gupta discloses this limitation. Pet. 41–46. Specifically, Petitioner asserts Zhang's prefetcher, which can be enabled or disabled at any point of a program, operates *independent* of the programmable logic blocks that are associated with a processor and use computational data to perform matrix multiplication computations. *Id.* at 41–43, 45–46 (citing Ex. 1003, 13, 15, Fig. 4). Petitioner further asserts Gupta's prefetch unit in the AMRM chip is (i) an implementation of the prefetcher disclosed in Zhang, and (ii) operates *independent* of and in *parallel* with logic blocks that use the computational data in the command processor. *Id.* at 43–46 (citing Ex. 1004, 9, 11, Fig. 1).

Petitioner contends an ordinarily skilled artisan would have been motivated to combine the teachings of Zhang and Gupta because (i) Gupta discloses a prototype implementation of the architecture that Zhang disclosed, (ii) Zhang teaches a "general architecture for data prefetching in a way that operates independent of and prior to the logic blocks that use the computational data," and "Gupta teaches a specific prototype implementation of that architecture and technique, including prefetching computational data in parallel with the logic blocks that use the computational data . . . as well as prefetching such data prior to and after those blocks," and (iii) "[i]t would therefore have been obvious for a [person of ordinary skill in the art] to look to Gupta to better understand one way to implement the data prefetch architecture taught by Zhang." Pet. 45–46 (citing Ex. 1003, 15; Ex. 1004, 9, 11; Ex. 1006 ¶¶ 153–154).

Patent Owner does not specifically respond to these arguments. *See generally* PO Resp.

After considering the evidence and arguments of the complete record, we determine that Petitioner has demonstrated by a preponderance of the evidence that Zhang in combination with Gupta teaches limitation 1(d) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

f) Limitation 1(e): "wherein at least the first memory and data prefetch unit are configured to conform to needs of the algorithm"

Petitioner asserts that Zhang in combination with Gupta discloses this limitation. Pet. 46–49. Petitioner asserts "Zhang teaches that its L1 cache and its data prefetcher are implemented in reprogrammable logic with the

goal of customizing these components to match an application." Pet. 46 (citing Ex. 1003, 12, 15, Fig. 4); *see also* Ex. 1006 ¶¶ 155–156. Petitioner also asserts "Gupta teaches a reconfigurable processor with the cache memory and data prefetcher [unit] implemented in FPGA reprogrammable logic to support application-specific cache organization and prefetching." *Id.* at 47–48 (citing Ex. 1004, 9, 11, Fig. 1; Ex. 1006 ¶¶ 157–158).

Petitioner contends an ordinarily skilled artisan would have been motivated to combine the teachings of Zhang and Gupta because (i) Gupta discloses a prototype implementation of the architecture that Zhang disclosed, (ii) Zhang teaches the general architecture for a data prefetch unit and first memory that are configured to conform to needs of the algorithm, and Gupta teaches a specific prototype implementation of that architecture and technique, including reconfigurable logic blocks in FPGAs for application-specific cache organization policies, hardware assisted blocking, prefetching, and dynamic cache structures, and (iii) "[i]t would therefore have been obvious for a [person of ordinary skill in the art] to look to Gupta to better understand one way to implement the reconfigurable data prefetch architecture taught by Zhang." Pet. 48–49 (citing Ex. 1003, 13, 14, Fig. 2; Ex. 1004, 9; Ex. 1006 ¶ 159).

Patent Owner argues that Zhang or Gupta does not disclose both a "first memory" and a "data prefetch unit" that is "configured to conform to needs of the algorithm." PO Resp. 44–49. Specifically, Patent Owner asserts that "the memory disclosed in *Zhang* and *Gupta* is not configurable at all." *Id.* at 45. Patent Owner further argues that Zhang and Gupta's data prefetch unit is not "configured to conform to [the] needs of the algorithm" because it does not "retrieve[] only computational data required by the algorithm from a second memory" and no other computational data. *Id.*

Patent Owner does not contend that Petitioner has not provided sufficient rationale to combine Zhang and Gupta. *See generally* PO Resp. 44–49.

With respect to Zhang, Patent Owner also argues that Zhang "uses reconfigurable logic only as controllers for the SRAM, DRAM, and L1 cache," but "uses conventional memories with a fixed line size of either 32 or 64 bytes." PO Resp. 46–47 (citing Ex. 1003, Fig. 2, Table 1). Patent Owner contends that Zhang "does not teach reconfiguring the cache line sizes (32 or 64 bytes) to match the needs of the particular algorithm which operates on 40-byte matrix elements." *Id.* at 45–46. Patent Owner also argues that Zhang's "prefetcher is implemented using reprogrammable prefetcher logic . . . which is then integrated with the L1 (and L2) caches . . . which comprise conventional static hardware." *Id.* at 47 (referring to Ex. 1003, Fig. 4). Patent Owner further argues that in contrast to the "single centralized L1 cache in *Zhang's* example," the first memory of the '867 patent would be configured "as multiple FIFO streams of the required width and depth in close proximity to multiple computational units performing matrix multiplication calculations in parallel, one whole row/column at a time." *Id.* at 48–49 (citing Ex. 1001, 4:8–10, 6:11–21, 7:17–19, 12:51–52; Ex. 2028 ¶¶ 88, 121).

With respect to Gupta, Patent Owner similarly argues that "[n]o indication is given that the memories themselves are implemented in programmable logic" because Gupta states that "the FPGAs on the board contain *controllers* for the SRAM, DRAM, and L1 cache." *Id.* at 48 (citing Ex. 1004, 10; Ex. 2028 ¶¶ 68, 97).

In Reply, Petitioner contends that Patent Owner's arguments that Zhang and Gupta's memory is "fixed," "conventional," and "not reconfigurable" are contradicted by the references. Reply 22–23. Petitioner

also asserts that Zhang teaches a reconfigurable cache in prefetch operations, and Gupta teaches a prefetch with a reconfigurable cache. *Id.* at 23 (citing Ex. 1003, 16; Ex. 1004, 9). Finally, Petitioner asserts that Zhang's cache line size of 32 or 64 bytes reflects what was chosen for line sizes in the simulation, and does not limit Zhang from adapting its cache size. *Id.* (citing Ex. 1003, 12, 15).

In Sur-reply, Patent Owner reiterates its argument that Zhang's memory is separate from the reconfigurable components. Sur-reply 9. Specifically, Patent Owner argues that Zhang and Gupta's reconfigurable logic "resides only in the peripheral components 'integrated with' as opposed to 'contained within' the memory and cache." *Id.* (citing Ex. 1003, 15). Patent Owner argues that "*Gupta's* figure 1 confirms that only the cache controllers are in the programmable logic." *Id.* at 10 (Ex. 1004, Fig. 1). Patent Owner contends that "[n]o indication is given that the memories themselves are implemented in programmable logic." *Id.* (citing Ex. 2028 ¶¶ 68, 97).

For the reasons advanced by Petitioner and as discussed below, we agree with Petitioner that Zhang and Gupta teach "first memory and data prefetch unit are configured to conform to needs of the algorithm." Specifically, Zhang describes a "machine architecture that integrates programmable logic into key components of the system with the goal of customizing architectural mechanisms and policies to match an application." Ex. 1003, 12; *see also id.* at 13. Zhang describes "integration of programmable logic with memory components," "adding a small amount of programmable logic to the memory units," and "integrating programmable logic into memory components." *Id.* at 14, 17, 18. Zhang "uses architectural adaptation for prefetching" to "enable[] more flexible prefetching policies,

e.g., multiple level prefetch, according to the application access pattern," and "architectural customization." *See* Ex. 1003, 15–16. Gupta describes a "machine architecture [that] supports . . . dynamic cache structures (such as stream, victim caches, stride prediction and miss history buffers)," and "an application-specific prefetching scheme that resides in dedicated hardware at arbitrary levels of the memory hierarchy, in all of them, or to bypass them completely." Ex. 1004, 9.

As discussed above in Section III.E.2.a.2, and for the same reasons as set forth for CPU in Zhang's Figure 2, we disagree with Patent Owner's argument that Zhang's memory is separate from the reconfigurable components and not configurable. That is, Zhang discloses "We propose an architecture that ***integrates small blocks of programmable logic into*** *key elements of a baseline architecture,* ***including*** processing elements, ***components of the memory hierarchy,*** and the scalable interconnect, to provide *architectural adaptation* – the *customization of architectural mechanism and policies to match an application.*" Ex. 1003, 13 (first, second, and fourth emphases added); *see also* Fig. 2.

We also find that Patent Owner's other arguments are unavailing. First, we note that the claim language does not recite or require a memory that "match[es] the exact needs of the algorithm" or a memory configured "as multiple FIFO streams of the required width and depth," as argued by Patent Owner. *See* PO Resp. 45–46, 48–49. Second, we disagree with Patent Owner that Zhang discloses a fixed cache line size. Rather, we agree with Petitioner and find that Zhang's cache line size of 32 or 64 bytes reflects what was chosen for line sizes in the simulation, and does not limit Zhang from adapting its cache size. *See* Reply 23–24. Zhang states that

Table 1, which shows the L1 and L2 Cache Line Size of 32B or 64B "shows the simulation parameter used." Ex. 1003, 15.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang in combination with Gupta teaches limitation 1(e) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

g) Limitation 1(f): "the data prefetch unit is configured to match format and location of data in the second memory"

Petitioner asserts that Zhang discloses this limitation. Pet. 49–53. In particular, Petitioner asserts that "Zhang's prefetcher is configured to match the format and location of data in the second memory using 'two pieces of application-specific information: the address ranges [location of data] and memory layout [format of the data] of the target data structures [second memory' from which the data is to be prefetched." *Id.* at 50 (citing Ex. 1003, 15) (bolding omitted). Petitioner asserts that the target data structures can comprise either the L2 cache or the main memory. *Id.* (citing Ex. 1003, 15–16; Fig. 4). Petitioner contends that "Zhang's data prefetcher is architecturally adaptable to utilize dynamic scatter gather to reduce data traffic between the memory and the processor by sending only used fields of data for a given computation to reduce bandwidth requirements." *Id.* at 51 (citing Ex. 1003, 16; Ex. 1006 ¶ 162). Petitioner contends that "Zhang's use of dynamic gather prefetching techniques to access data from [memory having an irregular data structure with the matrix elements stored in disparate (non-continuous) locations] is an example of the data prefetch unit

being configured to match the format and location of the data in the second memory." *Id.* at 51–52 (citing Ex. 1003, 16). Petitioner also contends that, as shown in Figure 5, "[t]he gather logic associated with Zhang's data prefetcher is configured to use [the] dynamic scatter-gather to match this format and location in the same manner as the example shown and described with respect to Figure 13 of the '867 patent." *Id.* at 53 (citing Ex. 1001, 9:24–41, Fig. 13); *see* Ex. 1006 ¶ 164.

Patent Owner makes two arguments that are generally duplicative of arguments made for other limitations. PO Resp. 49–50. First, Patent Owner argues "*Zhang* does not disclose reconfiguring the memory unit to match the needs of the algorithm." *Id.* at 49. Patent Owner contends that, "because the data prefetch unit of *Zhang* and *Gupta* always retrieves a full cache line, it is not 'configured to match format and location of data in the second memory." *Id.* at 49–50. We are not persuaded by this argument for the reasons set forth above in for limitation 1(c). Second, Patent Owner argues that "Petitioner has also not shown that *Zhang* discloses physically reconfiguring the data prefetch unit" and "*Zhang* does not disclose actually reconfiguring the prefetch unit to match the data in the second memory." *Id.* at 50. We find these arguments unavailing for the same reasons as set forth above for the preamble, limitation 1(c), and limitation 1(e). Patent Owner does not contend that Petitioner has not provided sufficient rationale to combine Zhang and Gupta. *See generally* PO Resp. 49–50.

After considering the evidence and arguments of record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang teaches limitation 1(f) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

3. Analysis of Independent Claim 13

a) Limitation 13(a): "transferring data between a memory and a data prefetch unit in a reconfigurable processor"

Petitioner asserts that Zhang in combination with Gupta discloses this limitation. Pet. 63. Specifically, for the same reasons as set forth in independent claim 1, Petitioner asserts that Zhang and Gupta teach "a data prefetch unit in a reconfigurable processor that is adapted to transfer computational data between the data prefetch unit and an L2 cache and/or main memory, which are each components of the memory hierarchy in the reconfigurable processor." *Id.*

Patent Owner argues Zhang does not teach a "reconfigurable processor," for the same reasons it set forth for claim 1. PO Resp. 51–52. As discussed above in the analysis of the preamble of claim 1, we find these arguments unavailing.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang teaches the limitation 13(a) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

b) Limitation 13(b): "transferring the data between a computational unit and a data access unit"

Petitioner asserts that Zhang discloses this limitation. Pet. 64–67. Petitioner contends that Zhang teaches "computational units" in its reconfigurable processor that perform computations because Zhang discloses a "reconfigurable processor with programmable logic," that "performs 'sparse matrix computations' and 'matrix multiply' computations." *Id.* Petitioner further contends that Zhang teaches a "*data access unit*" because

"Zhang's processor uses a set of logic to access components of the memory hierarchy (L1 or L2 cache memory) and deliver the prefetched data (Val1/Val2/Val3) directly to the processor's computational units to perform the computation." *Id.* at 64–66. Petitioner further asserts that a person of ordinary skill in the art "also understands Zhang to disclose a data access unit that uses one or more registers to deliver data directly to the computational units." *Id.* at 66–67.

Patent Owner does not specifically respond to these arguments. *See generally* PO Resp.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang teaches limitation 13(b).

> *c) Limitation 13(c): "wherein the computational unit and the data access unit, and the data prefetch unit are configured to conform to needs of an algorithm implemented on the computational unit and transfer only data necessary for computations by the computational unit"*

Petitioner asserts that Zhang in combination with Gupta discloses this limitation. Pet. 67–69. Specifically, Petitioner asserts that "Zhang teaches a reconfigurable processor containing computational units, a data access unit, and a data prefetch unit implemented in reprogrammable logic such as FPGA." *Id.* at 67. Petitioner further contends that Zhang teaches that its "processing elements," "memory hierarchy," and the data prefetcher are "implemented in reconfigurable logic to allow them to be customized to match (configured to conform to) the algorithm in the application." *Id.* at 68. Similarly, Petitioner contends that Gupta's "computational units, data access unit and data prefetch unit are also implemented in reconfigurable

logic . . . [and] are configurable to conform to the requirements of algorithms implemented in the computational logic." *Id.* Petitioner further contends that "the combination discloses a prefetching technique that transfers only the data necessary for computations by the computational unit(s) using scatter-gather memory access techniques." *Id.* at 69.

Patent Owner argues Petitioner fails to establish that Zhang and Gupta teach a "computational unit . . . configured to conform to needs of an algorithm implemented on the computational unit" because "*Zhang* uses programmable logic (FPGA) only as a means to deliver data for use by a conventional CPU, as opposed to a reconfigurable computational unit." PO Resp. 51–52. Patent Owner also argues that Zhang and Gupta "do not utilize programmable logic to implement a 'reconfigurable processor.'" *Id.* at 52.

We find these arguments unavailing for the same reasons as discussed with respect to the preamble of claim 1 and limitation 1(e). After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang and Gupta teach limitation 13(c) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

d) Limitation 13(d): "wherein the prefetch unit operates independent of and in parallel with the computational unit"

Petitioner contends Zhang in combination with Gupta teaches this limitation. Pet. 69. Specifically, Petitioner relies in part on its analysis of limitations 13(b) and 1(d) and contends that "Zhang's computational units include logic blocks that use data to perform computations (multiply, add), .

. . and the combination's data prefetch unit operates independent of and in parallel with the logic blocks that use computational data." *Id.*

Patent Owner makes no specific arguments directed toward these limitations. *See generally* PO Resp.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang teaches limitation 13(d) and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

4. Analysis of Claim 2

Claim 2 depends from claim 1 and recites

> wherein the data prefetch unit is coupled to a memory controller that controls the transfer of the data between the memory and the data prefetch unit and transmits only portions of data desired by the data prefetch unit and discards other portions of data prior to transmission of the data to the data prefetch unit.

Ex. 1001, 12:55–60.

Petitioner contends the combination of Zhang and Gupta teaches the limitations in claim 2. Pet. 53–58. Specifically, Petitioner contends Zhang "teaches transferring data between a second memory (L2 cache or main memory) and the data prefetch unit." *Id.* at 54. Petitioner also relies on the second case study in Zhang, and in particular that "Zhang teaches that the prefetcher uses scatter-gather to retrieve only computational data required by the algorithm." *Id.* at 54–57. With respect to Gupta, Petitioner contends "Gupta discloses a prototype implementation of the data prefetch unit and scatter-gather technique disclosed in Zhang." *Id.* at 56. Petitioner contends that the "combination's memory controller therefore discloses the function of transmitting only the portions of data desired by the data prefetch unit."

Id. at 57–58. Petitioner argues that transmitting only the portions of data desired by the prefetcher may be accomplished in one of two ways: (1) extracting only the requested data from the cache line; or (2) extracting an entire cache line and discarding the portions that were not requested prior to transmission. *Id.* at 58. Petitioner contends that a person of ordinary skill in the art "would understand discarding other portions of data prior to transmission was obvious to try, choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success." *Id.*

Patent Owner makes similar arguments as for claim 1. For example, Patent Owner argues that "all remaining data in the optimized storage scheme is required and is therefore transferred and not discarded." PO Resp. 50–51. Patent Owner also argues that "[s]ince all fields of the matrix structure are prefetched, no data portions are discarded prior to data transmission." *Id.* at 51. Patent Owner does not contend that Petitioner has not provided sufficient rationale to combine Zhang and Gupta. *See id.* at 50–51.

We do not find Patent Owner's arguments persuasive for the same reasons as discussed above with respect to limitation 1(c). After considering the evidence and arguments of record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang and Gupta teach the limitations in claim 2 and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

5. Analysis of Claims 4–8 and 14–19

Petitioner contends that dependent claims 4–8 and 14–19, which depend from one of independent claims 1 and 13, would have been obvious over Zhang and Gupta, and provides explanation as to how the prior art

(either Zhang alone, or in combination with Gupta) teaches each claim limitation and, where applicable, provides rationale for the combination. Pet. 58–63, 69–76.

Patent Owner does not specifically respond to these arguments. *See generally* PO Resp.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang and/or Gupta teach the limitations in claims 4–8 and 14–19 and, if applicable, that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

6. Secondary Considerations of Nonobviousness

Because the parties refer to the claims collectively in addressing objective indicia of nonobviousness, we do so as well and discuss our findings as applied to all of Petitioner's grounds. *See* PO Resp. 57–59; Reply 24–26.

Patent Owner argues secondary considerations of nonobviousness demonstrate that the claims would not have been obvious to a person of ordinary skill in the art. PO Resp. 57–59. Specifically, Patent Owner argues (i) long-felt need in the industry; (ii) failure of others; and (iii) industry praise. *Id.* Patent Owner raised identical arguments in its Preliminary Response. *See* Paper 9 at 58–60. In the Institution Decision, we found that Patent Owner's arguments were insufficiently developed, in that Patent Owner had not established a nexus between the evidence of secondary considerations and the claimed invention. Dec. 68–71. Patent Owner has not further developed the record during trial. *See generally* PO Resp.

Petitioner contends, *inter alia*, that Patent Owner has not shown the required nexus. Reply 24–26.

"For objective evidence of secondary considerations to be accorded substantial weight, its proponent must establish a nexus between the evidence and the merits of the claimed invention." *ClassCo, Inc. v. Apple, Inc.*, 838 F.3d 1214, 1220 (Fed. Cir. 2016) (citing *In re Huai-Hung Kao*, 639 F.3d 1052, 1068 (Fed. Cir. 2011)). The question of nexus is highly fact specific and it is Patent Owner's burden to establish a sufficient nexus. *Fox Factory, Inc. v. SRAM, LLC*, 944 F.3d 1366, 1373 (Fed. Cir. 2019). "To determine whether the patentee has met that burden, we consider the correspondence between the objective evidence and the claim scope." *Id.* (citation omitted). A patentee is entitled to a rebuttable presumption of nexus "when the patentee shows that the asserted objective evidence is tied to a specific product and that product embodies the claimed features, and is coextensive with them." *Id.* (citation omitted).

Patent Owner does not provide an analysis demonstrating that its products embody the features of the challenged claims of the '867 patent, and that those products are coextensive, or nearly coextensive, with those claims. Instead, Patent Owner speaks generally about "SRC's inventions, including those claimed in the '867 patent" (PO Resp. 57), SRC's "innovative technology covering numerous aspects of reconfigurable computing" (*id.*), "SRC's intellectual property" (*id.*), "SRC Computers' proven systems" (*id.* at 58), and "SRC's technology" (*id.* at 59). Patent Owner even states that "SRC was issued *over 30 U.S. patents* (with over 2,090 forward citations) for its innovative technology covering numerous aspects of reconfigurable computing." PO Resp. 57 (citing Ex. 2005 ¶¶ 52–53) (emphasis added). However, Patent Owner does not provide analysis

demonstrating sufficiently that any of its technology or products (e.g., the SRC-6) are covered by the *'867 patent* (as opposed to the other 30+ patents), or are coextensive with the challenged claims. We, therefore, find that a presumption of nexus is inappropriate on this record.

However, "[a] finding that a presumption of nexus is inappropriate does not end the inquiry into secondary considerations." *Fox Factory*, 944 F.3d at 1373. "To the contrary, the patent owner is still afforded an opportunity to prove nexus by showing that the evidence of secondary considerations is the 'direct result of the unique characteristics of the claimed invention.'" *Id.* at 1373–1374 (citation omitted).

With respect to long-felt need in the industry, Patent Owner provides testimony from Dr. Mangione-Smith that "there was considerable pressure in the industry to develop computing systems with drastically higher performance, lower operating expense, lower power usage, and lower space requirements." PO Resp. 57 (citing Ex. 2028 ¶ 44); Ex. 2011; Ex. 2012. Patent Owner argues that "SRC's inventions, including those claimed in the '867 patent, were implemented in its Memory Algorithm Processor (MAP)–a supercomputing processor–that directly addressed those needs." *Id.* (citing Ex. 2026 ¶ 73). Patent Owner states that SRC "was issued over 30 U.S. patents . . . for its innovative technology covering numerous aspects of reconfigurable computing." *Id.* at 57 (citing Ex. 2005 ¶¶ 52–53). Patent Owner also asserts that "[f]or more than a decade, SRC was the sole source provider of MAP supercomputers to Lockheed Martin on behalf of the U.S. Southern Command of high-performance reconfigurable processors." *Id.* at 57–58 (citing Ex. 2006 ¶¶ 2–23; Ex. 2007; Ex. 2008). According to Patent Owner, "[t]he inventors of the '867 patent . . . pioneered the use of Field Programmable Gate Arrays (FPGAs) as general-purpose processors to create

small energy-efficient supercomputers." *Id.* at 58 (citing Ex. 2005 ¶¶ 45–46).

With respect to failure of others, Patent Owner contends that "the '867 patent shifted the paradigm from using small pockets of reconfigurable hardware to building an entire reconfigurable processor able to instantiate a whole algorithm in hardware." *Id.* at 58–59 (citing Ex. 2028 ¶¶ 56, 64, 66, 67, 89, 95, 97, 101, 119). Patent Owner also argues that SRC "has spent over $100 million on research and development including the '867 patent." *Id.* at 59 (citing Ex. 2003 ¶ 3).

With respect to industry praise, Patent Owner contends that its "technology was received in the industry with overwhelming praise." *Id.* In support of its arguments, Patent Owner submits two articles. *Id.* (citing Ex. 2013; Ex. 2014).

Patent Owner does not provide sufficient explanation establishing a nexus to the challenged claims of the '867 patent. Specifically, Patent Owner does not provide analysis that the *claimed* systems and method solved the identified long-felt needs in the industry, failure of others, or industry praise was directed to the *claimed* systems and method. Instead, Patent Owner's evidence and arguments rely upon the SRC-6 product and/or MAP supercomputer (e.g., Ex. 2007; Ex. 2008; 2011; Ex. 2012; Ex. 2013; Ex. 2014), but, as discussed above, Patent Owner does not provide sufficient explanation or analysis demonstrating that these products implemented the systems and methods recited in any of the challenged claims.

Accordingly, we have considered Patent Owner's objective indicia of nonobviousness, and determine they carry little weight because Patent Owner has not provided persuasive evidence of a nexus between the

evidence and the challenged claims. This determination applies to all grounds advanced by Petitioner, and we weigh it accordingly.

7. Conclusion

Having considered the *Graham* factors, including the scope and content of the prior art, the differences between the prior art and the challenged claims, and the objective evidence of nonobviousness, we determine that Petitioner has shown, by a preponderance of the evidence, that claims 1, 2, 4–8, and 13–19 would have been obvious over Zhang and Gupta.

F. Ground 2: Obviousness Over Zhang, Gupta, and Chien

Petitioner contends claims 3 and 9–12 would have been obvious over the combination of Zhang, Gupta, and Chien. Pet. 5, 76–87. After reviewing the entire record developed at trial, as explained below, we determine that Petitioner has shown, by a preponderance of the evidence, that claims 3 and 9–12 are unpatentable over the combination of Zhang, Gupta, and Chien.

1. Patent Owner's General Arguments as to Chien

As with Zhang and Gupta, Patent Owner similarly argues that Chien teaches away from the claimed invention and is not an enabling reference. *See* PO Resp. 13 ("MORPH teaches away from the use of reconfigurable logic for application-specific functional units or computational logic"); *id.* at 14 ("Since the theory of *Chien* cannot be reproduced without undue experimentation, it is not an enabling prior art reference.") We find these arguments unavailing for the same reasons as set forth above in Section III.E.1.

2. Analysis of Independent Claim 9

The limitations in independent claim 9 are similar to those in independent claims 1 and 13. Petitioner asserts that the combination of Zhang, Gupta, and Chien discloses the limitations in claim 9, for many of the same reasons as set forth for claims 1 and 13. Pet. 83–86. In addition, Petitioner contends Chien teaches "a common memory." *Id.* at 83–84 (citing Ex. 1005, 11, Fig. 2). Petitioner contends that "[i]t would have been obvious to implement Zhang-Gupta's L2 cache memory as Chien's shared L2 cache memory and/or Zhang-Gupta's main memory as Chien's global shared memory." *Id.* at 84. Petitioner contends that "Zhang expressly cites to and incorporates Chien for explanation of its multiprocessor MORPH architecture . . . and Chien teaches shared L2 cache as one of only two approaches . . . which a [person of ordinary skill in the art] would be motivated to try leading to a predictable result." *Id.*

Patent Owner repeats the arguments made with respect to claims 1 and 13, which we do not find availing for the reasons discussed above. PO Resp. 53–56. Patent Owner does not contend that Petitioner has not provided sufficient rationale to combine Zhang, Gupta, and Chien. *See* PO Resp. 53–56.

After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang, Gupta, and Chien teach the limitations in claim 9 and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang, Gupta, and Chien and would have had a reasonable expectation of success in doing so.

3. Analysis of Claims 3 and 10–12

Petitioner contends the combination of Zhang, Gupta, and Chien teaches the limitations in claims 3 and 10–12. Pet. 76–81.

Patent Owner presents similar arguments as made for claims 1 and 13 for dependent claims 10 and 11, but does not separately argue claims 3 and 12. *See* PO Resp. 56–57. We find Patent Owner's arguments unavailing for the reasons discussed above. After considering the evidence and arguments of the complete record, we determine that Petitioner has shown, by a preponderance of the evidence, that Zhang, Gupta, and Chien teach the limitations in claims 3 and 10–12 and that one of ordinary skill in the art would have been motivated to combine the teachings of Zhang and Gupta and would have had a reasonable expectation of success in doing so.

4. Conclusion

Having considered the *Graham* factors, including the scope and content of the prior art, the differences between the prior art and the challenged claims, and the objective evidence of nonobviousness (*see* Section III.E.6), we determine that Petitioner has shown, by a preponderance of the evidence, that claims 3 and 9–12 would have been obvious over Zhang, Gupta, and Chien.

IV. REVISED MOTION TO AMEND

Because we conclude that Petitioner shows, by a preponderance of the evidence, that each of challenged claims 1–19 is unpatentable, we consider Patent Owner's Revised Motion to Amend. Patent Owner contingently moves to allow proposed substitute claims 20–38, should we determine that any of the original claims are unpatentable. Mot. Amend at 1–2. For the reasons below, we find that Patent Owner has not met its burden in asserting

that proposed substitute claims 20–38 have written description support in the original application that issued as the '867 patent.

A. Patent Owner's Proposed Substitute Claims

Patent Owner proposes claims 20–38 as substitutes for original claims 1–19. Mot Amend. 1. Specifically, Patent Owner proposes claims 20 and 21 as substitutes for original independent claim 1, substitute claim 28 for original independent claim 9, and substitute claim 32 for original independent claim 13. *Id.* at 18 (Appendix A). Patent Owner additionally cancels original dependent claims 6–8, and proposes substitute claim 22 for original dependent claim 2, substitute claims 23–24 for original dependent claim 3, substitute claims 25–26 for original dependent claim 4, and substitute claims 27, 29–31, and 33–38 for original dependent claims 5, 10–12, and 14–19. *Id.*

Proposed substitute claim 20 is representative, and reproduced below, using underscoring to indicate text added to original independent claim 1.

> 20. A reconfigurable processor that instantiates an algorithm as hardware comprising:
>
> a first memory having a first characteristic memory bandwidth and/or memory utilization; and
>
> a data prefetch unit coupled to the memory, wherein the data prefetch unit retrieves only computational data required by the algorithm from a second memory of second characteristic memory bandwidth and/or memory utilization and places the retrieved computational data in the first memory wherein the data prefetch unit operates independent of and in parallel with logic blocks using the computational data, and wherein at least the first memory and data prefetch unit are configured to conform to needs of the algorithm, and the data prefetch unit is configured to match format and location of data in the second memory,

> wherein the reconfigurable processor is neither integrated within nor comprises a conventional microprocessor, and
>
> wherein the reconfigurable processor operates independent of and in parallel with a conventional microprocessor.

Mot. Amend. 5–6. As shown above, Patent Owner amends claim 1 to add two new limitations.

1. Statutory Requirements and Burden

"Before considering the patentability of any substitute claims, . . . the Board first must determine whether the motion to amend meets the statutory and regulatory requirements set forth in 35 U.S.C. § 316(d) and 37 C.F.R. § 42.121." *See Lectrosonics, Inc. v. Zaxcom, Inc.*, IPR2018-01129, Paper 15 at 4 (PTAB Feb. 25, 2019) (precedential). Patent Owner bears the burden of persuasion to show that the Revised Motion to Amend meets those requirements. 37 C.F.R. § 42.121(d)(1).

Pursuant to 35 U.S.C. § 316(d), a motion to amend may "cancel any challenged patent claim" or, for each challenged claim, "propose a reasonable number of substitute claims." However, the motion to amend "may not enlarge the scope of the claims of the patent or introduce new matter." 35 U.S.C. § 316(d)(3). New subject matter is any addition to the claims that lacks sufficient support in the subject patent's original disclosure. *See TurboCare Div. of Demag Delaval Turbomach. v. Gen. Elec. Co.*, 264 F.3d 1111, 1118 (Fed. Cir. 2001) ("When [an] applicant adds a claim . . . , the new claim[] must . . . find support in the original specification.").

Corresponding Rule 42.121 provides, "[a] motion to amend may be denied where . . . [t]he amendment seeks to . . . introduce new subject matter." 37 C.F.R. § 42.121(a)(2)(ii). Rule 42.121(b) requires the motion to amend to "set forth: (1) The support in the original disclosure of the patent

for each claim that is added or amended; and (2) The support in an earlier-filed disclosure for each claim for which benefit of the filing date of the earlier filed disclosure is sought." 37 C.F.R. § 42.121(b); *Lectrosonics*, Paper 15 at 7–8 (explaining that the motion to amend *itself* must set forth the written description support for each proposed substitute claim as a whole, and not just the features added by the amendment).

2. Written Description Support

The determination of whether there is sufficient written description support turns on whether the original disclosure of the application relied upon reasonably conveys to a person of ordinary skill in the art that the inventor had possession of the claimed subject matter as of the filing date. *Ariad Pharms., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336, 1351 (Fed. Cir. 2010) (en banc).

Patent Owner asserts that proposed claim 20 does not introduce new matter. Mot. Amend 3–4; Reply Amend 3–5. Patent Owner contends that "[t]he claim listing in Appendix A clearly indicates the specific changes for each proposed amended claim and sets forth: (1) the support in the original disclosure of the patent for each added or amended claim; and (2) the support for each claim in an earlier-filed disclosure for which benefit of the filing of the earlier filed disclosure is sought." Mot. Amend at 4. The entries from Appendix A for the two new limitations in proposed claim 20[10] are reproduced below:

Amended Claim 20	Support in '867 Patent
wherein the reconfigurable processor is neither integrated	Abstract; 3:64-4:3; 5:19-29; 5:34-37; 5:59-6:4; 6:5-31, 6:47-58; Figs.

[10] Patent Owner provides citations indicating "Support in the '867 Patent" for each limitation in claim 20; we only address the citations provided for the two new limitations.

within nor comprises a conventional microprocessor, and	1-7 and related descriptions.
wherein the reconfigurable processor operates independent of and in parallel with a conventional microprocessor.	Abstract; 3:64-4:3; 5:19-29; 5:34-37; 5:59-6:4; 6:5-31, 6:47-58; Figs. 1 and 3 and the related embodiment; Figs. 1, 2, and 4-7 and related descriptions.

Mot. Amend 20 (Appendix A). In the Revised Motion to Amend, Patent Owner provides no discussion of the alleged support provided in Appendix A. *See id.* at 4, 18. Rather, Patent Owner states "Appendix A thereby shows that the amended claims would reasonably convey to one of ordinary skill in the art that the inventors were in possession of the claimed subject matter as of the filing date of the '867 patent." *Id.* at 4.

Petitioner contends that the Revised Motion to Amend should be denied because Patent Owner fails to adequately explain the written description support for the claims.[11] Opp. Amend 4–11. According to Petitioner, Patent Owner does not meet the requirements set forth in *Lectrosonics*, because Patent Owner (1) improperly places its purported support in Appendix A, rather than the motion itself; (2) cites to the issued patent rather than the original application's disclosure; (3) fails to set forth support in the provisional's disclosure and thus fails to show benefit to that earlier filing date; and (4) does not proffer any written description support

[11] Petitioner also argues that (1) Patent Owner fails to show it proposes a reasonable number of claims; (2) Patent Owner fails to show that its amendments do not enlarge the claim scope; (3) Patent Owner fails to show that certain amendments respond to a ground of unpatentability; and (4) the proposed claims are unpatentable over the prior art. Opp. Amend 3, 11–25. Because we find that Patent Owner has failed to provide written description support for the amended claims, we do not address these arguments.

for the fifteen substitute dependent claims. *Id.* at 4–6. In addition to its *Lectrosonics* arguments, Petitioner argues that the citations Patent Owner provides fail to provide written description support for the new limitations in the proposed substitute independent claims. *Id.* at 6–11.

As set forth below, we find that Patent Owner has not sufficiently demonstrated written description support for the two new limitations. In addition, we find that Patent Owner's string citations to various disclosures in the '867 patent are insufficient to demonstrate written description support absent some additional explanation.

a) Compliance with the Rule Requirements as set forth in Lectrosonics

First, Petitioner contends that Patent Owner improperly places its written description support in Appendix A, rather than in the motion itself. Opp. Amend 5. *Lectrosonics* states that "[t]he written description support must be set forth in the motion to amend itself, not the claim listing." *Lectrosonics*, Paper 15 at 8. *Lectrosonics* further states that "the claim listing may be filed as an appendix to the motion to amend and shall not count toward the page limit for the motion. The appendix, however, shall not contain any substantive briefing. All arguments and evidence in support of the motion to amend shall be in the motion itself." *Id.* Although we agree with Petitioner that, under *Lectrosonics*, Patent Owner should have provided the substantive briefing associated with the written description in the motion itself, rather than Appendix A, we also acknowledge Patent Owner's argument that although that section is labeled "Appendix," it still falls within the 25-page limit for motions to amend. *See* Opp. Amend 5; Reply Amend 3. Therefore, we will consider the arguments presented, and

do not rely on this deficiency as a basis to deny the Revised Motion to Amend.

Second, Petitioner contends Patent Owner improperly cites to the issued patent. Opp. Amend 5–6. There is no dispute that Patent Owner cites to Exhibit 1001, the issued patent, rather than to the original application's disclosure (Exhibit 1002), or to the provisional's disclosure (Exhibit 1017). *See* Opp. Amend 5–6; Reply Amend 3–4. Patent Owner contends that it is "irrelevant" because the cited art predates both the '867 patent, as well as the provisional application, so "citation to the actual patent disclosure is sufficient for the purposes of this IPR." PO Sur-Reply 4. Although we agree with Petitioner that the Revised Motion to Amend is deficient in citing to Exhibit 1001, rather than the original application (*see Lectrosonics*, Paper 15 at 8), neither party identifies any material differences between Exhibit 1001 and the original application in this regard, or the necessity of the earlier filing date, which would require citation to the provisional application. Therefore, we will consider the arguments presented, and do not rely on this deficiency as a basis to deny the Revised Motion to Amend.

Third, Petitioner contends that Patent Owner has not provided written description support for the proposed dependent claims. Opp. Amend 6. The chart in Appendix A only includes proposed independent claims 20, 21, 28, and 32, and Patent Owner does not provide written description support for the proposed dependent claims elsewhere in the Revised Motion to Amend. *See* Mot. Amend 19–26. Patent Owner contends that the Revised Motion to Amend provides written description support for each element of the proposed independent claims, which are incorporated by reference in the dependent claims, and are amended only by virtue of depending from that revised proposed independent claim. Reply Amend 4. However, this, does

not meet the requirements set forth in *Lectrosonics*, which states that "[i]n addition, the motion must set forth written description support for each proposed substitute claim as a whole, and not just the features added by amendment. *This applies equally to independent claims and dependent claims, even if the only amendment to a dependent claim is in the identification of the claim from which it depends*." *Lectrosonics*, Paper 15 at 8 (emphasis added). Therefore, Patent Owner does not provide written description support for proposed dependent claims 22–27, 29–31, and 33–38. We now turn to Patent Owner's asserted written description support for proposed substitute claim 20.

b) "wherein the reconfigurable processor is neither integrated within nor comprises a conventional microprocessor"

Petitioner argues that this is a negative limitation, and Patent Owner's citations do not provide written description support. Opp. Amend 7 (citing *Santarus, Inc. v. Par Pharm., Inc.*, 694 F.3d 1344, 1351 (Fed. Cir. 2012); MPEP § 2173.05(i)). For example, Petitioner argues that cited portions of the Abstract, columns 3, 4, and 5, as well as Figures 1–7 and their related disclosures do not mention a "conventional processor." *Id.* Petitioner contends that although the cited portions of column 6 refer to a "conventional computer" and a "conventional hardware platform," there is no discussion of a "conventional microprocessor" or its relation to a reconfigurable processor. *Id.* at 7–8. Petitioner also argues that Patent Owner's string citations, without more, do not meet its burden. *Id.* at 5; Sur-reply Amend 4.

Patent Owner does not directly respond to these arguments as to this limitation, but more generally states, relying on column 6, lines 20–25 of the

'867 patent, that "[t]his description clarifies that the [reconfigurable processor] is in the memory subsystem, separate and apart from the primary conventional processor, and operates independent of and in parallel." Reply Amend 5. Patent Owner also states that "[f]urther written description support for this[12] proposed claim element is found at '867 Patent at abstract; 3:64–4:3; 5:19–29; 5:34–37; 5:59–6:4; 6:5–31; 6:47–58; Figs. 1–7 and related the descriptions."[13] *Id.*

We agree with Petitioner that none of the citations provided by Patent Owner describes a "reconfigurable processor [that] is neither integrated within nor comprises a conventional microprocessor." At the outset, we agree with Petitioner that this is a negative limitation, in that it recites that the claimed reconfigurable processor is neither integrated within nor comprises a conventional microprocessor. "Negative claim limitations are adequately supported when the specification describes a reason to exclude the relevant limitation. Such written description support need not rise to the level of disclaimer. In fact, it is possible for the patentee to support both the inclusion and exclusion of the same material." *Santarus*, 694 F.3d at 1351. For example, in *Santarus*, the Court found that the claim limitation "wherein the composition contains no sucralfate" was supported by the specification because the specification described certain disadvantages to using sucralfate. *Santarus*, 694 F.3d at 1350–1351. "[P]roperly describing alternative features – without articulating advantages or disadvantages of each feature – can constitute a 'reason to exclude' under the standard articulated in *Santarus*." *Inphi Corp. v. Netlist, Inc.*, 805 F.3d 1350, 1355 (Fed. Cir.

[12] Patent Owner appears to be referring to both new claim limitations.
[13] These are the same string citations provided in the chart in Appendix A for this limitation.

2015). That is, "[i]f alternative elements are positively recited in the specification, they may be explicitly excluded in the claims." *Id.* at 1356. Accordingly, the issue is whether Patent Owner has shown that the specification of the '867 patent provides a "reason to exclude" reconfigurable processors that are integrated with, or comprise a conventional microprocessor.

Patent Owner addresses the citation in column 6, lines 20–25 of the '867 patent, which states:

> a number of RPs [reconfigurable processors] 100 are implemented within a memory subsystem of a conventional computer, such as on devices that are physically installed in dual inline memory module (DIMM) sockets of a computer. In this manner the RPs 100 can be accessed by memory operations and so coexist well with a more conventional hardware platform.

Ex. 1001, 6:20–25. Patent Owner contends that "[t]his description clarifies that the [reconfigurable processor] is in the memory subsystem, separate and apart from the primary conventional processor, and operates independent of and in parallel." Reply Amend 5. We find this explanation insufficient; even if this describes reconfigurable processor is implemented in the memory subsystem, we agree with Petitioner that "[t]he passage does not disclose *a reason to exclude* a reconfigurable processor from being integrated within or comprising a conventional processor, nor does it disclose a reconfigurable processor integrated within or comprising a conventional processor as an excludable alternative." Sur-reply Amend 5.

The remaining string citations provide no further insight, and Patent Owner has not explained how these citations provide written description support for this limitation. For example, Patent Owner relies on three

definitions provided in the '867 patent, but does not provide further explanation as to how these definitions describe this limitation. *See* Ex. 1001, 5:19–29 ("Direct execution logic" and "[r]econfigurable [p]rocessor"), 5:34–37 ("Functional Unit"). The remaining citations generally describe a reconfigurable processor, but do not describe how it relates to a conventional microprocessor (i.e., is integrated with or separate from), or even mention a conventional microprocessor at all. *See* Ex. 1001, code (57), 3:64–4:3, 5:59–6:4, 6:5–31, 6:47–58. At the hearing, Patent Owner admitted that "the specific citations always refer to a reconfigurable processor, and never once is there any discussion of a conventional processor." Tr. 51:23–26; *see also id.* at 52:4–7 ("So in the '867 patent, even the figures that are directed to the reconfigurable processor do not ever show a conventional processor, and there's no mention of a conventional CPU."), 52:11–18 ("So, in fact, every time the reconfigurable processor is described, it is described with other reconfigurable components, it's described with configuring the reconfigurable processor so that it can execute computations, and there is no mention of the reconfigurable processor being any way dependent on a conventional processor and in any way connected to a conventional processor."), 58:3–60:8. Patent Owner contends, without any supporting evidence, that "if the reconfigurable processor had any reliance for any function on a conventional processor, it would have been mentioned." Tr. 59:25–60:2. On this record, we are not persuaded that simply not mentioning a conventional microprocessor in the '867 specification provides a "reason to exclude" reconfigurable processors that are integrated with, or comprise a conventional microprocessor, or that not mentioning a conventional microprocessor "reasonably conveys to a person of ordinary skill in the art that the inventor had possession of a

'reconfigurable processor [that] is neither integrated within nor comprises a conventional microprocessor.'"

We also agree with Petitioner that the string citations provided in Appendix A are not sufficient to meet Patent Owner's burden. *See B.E. Tech., L.L.C. v. Google., Inc.*, No. 2015-1827, 2016 WL 6803057, at *7 (Fed. Cir. Nov. 17, 2017) (determining that B.E. did not meet its burden to show written description support for its substitute claims because it did not present argument to the Board, but instead "provided a string citation to eighteen pages of the '314 patent's original specification, without explaining how the various pages supported each of the proposed substitute limitations"); *Lippert Components, Inc. v. Days Corp.*, IPR2018-00777, Paper 28 at 52 (PTAB Sept. 24, 2019) ("It is not the Board's responsibility to search through the string citations to find sufficient written description support for each limitation, and we decline to do so."); *Intel Corp. v. Alacritech, Inc.*, IPR2017-01406, Paper 83 at 43–49 (PTAB Nov. 26, 2018) (same); *Respironics, Inc. v. Zoll Med. Corp.*, IPR2013-00322, Paper 46 at 24 (PTAB Sept. 17, 2014) ("[Patent Owner's] string citations amount to little more than an invitation to us (and to [Petitioner], and to the public) to peruse the cited evidence and piece together a coherent argument for them. This we will not do; it is the province of advocacy."), *vacated on other grounds*, 656 F. App'x 531 (Fed. Cir. 2016).

For example, it is unclear whether the string citations provided are to be understood as a combination of disclosures that, taken together, disclose the corresponding limitation, or whether Patent Owner contends that each citation on its own is sufficient to disclose the corresponding limitation. The lack of clarity is compounded considering that Patent Owner relies on the *same* string citations for *both* new limitations in claim 20. *See* Mot. Amend

20 (Appendix A). In addition, Patent Owner cites to Figures 1–7 "and related description," which covers half of the Figures and eight columns in the '867 patent. It is not the Board's responsibility to search through the string citations to find sufficient written support for this limitation when Patent Owner has failed to provide explanation.

Accordingly, for the foregoing reasons, we find that Patent Owner has not met its burden under 35 U.S.C. § 316(d)(3) and 37 C.F.R. § 42.121 to provide written description support for this limitation.

c) "wherein the reconfigurable processor operates independent of and in parallel with a conventional microprocessor"

Petitioner similarly argues that Patent Owner's citations do not provide written description support for this limitation. Opp. Amend 8. Patent Owner relies on the same string citations and arguments as for the limitation discussed above. Mot. Amend 4, 20 (Appendix A); Reply Amend 3–5. This limitation is not a negative limitation, but it similarly describes a relationship between the claimed reconfigurable processor and the conventional processor, i.e., independent and parallel with.

For the same reasons as set forth above, we find that the string citations Patent Owner provides in Appendix A are not sufficient to support the written description requirement for this limitation. We also find, for the same reasons as above, that Patent Owner has not sufficiently explained how these string citations provide written description support for this limitation. For example, we similarly are not persuaded on this record that simply not mentioning a conventional microprocessor in the '867 specification "reasonably conveys to a person of ordinary skill in the art that the inventor

had possession of a 'reconfigurable processor [that] operates independent of an in parallel with a conventional microprocessor.'"

Accordingly, for the foregoing reasons, we find that Patent Owner has not met its burden under 35 U.S.C. § 316(d)(3) and 37 C.F.R. § 42.121 to provide written description support for this limitation.

3. Proposed Independent Claims 21, 28, and 32, and Proposed Dependent Claims 22–27, 29–31, and 33–38

For each of proposed amended claims 21, 28, and 32, Patent Owner has added the same two, or comparable, limitations as discussed above for amended claim 20, and relies on the same string citations for written description support. *See* Mot. Amend 21 (claim 21), 23 (claim 28), and 24 (claim 32). Accordingly, for the same reasons discussed above, we find that Patent Owner has not met its burden under 35 U.S.C. § 316(d)(3) and 37 C.F.R. § 42.121 to provide written description support for these limitations.

Additionally, as discussed above in Section IV.A.2.a, Patent Owner has not provided written description support for proposed dependent claims 22–27, 29–31, and 33–38. Further, they lack written description support for the same reasons as the proposed claims from which they depend, i.e., proposed claims 20, 21, 28, and 32. Accordingly, for the same reasons discussed above, we find that Patent Owner has not met its burden under 35 U.S.C. § 316(d)(3) and 37 C.F.R. § 42.121 to provide written description support for the proposed dependent claims.

4. Conclusion

We conclude Patent Owner has not satisfied its burden of showing written description support for the proposed substitute claims. We deny Patent Owner's Revised Motion to Amend.

V. OBJECTIONS TO DEMONSTRATIVES

Prior to the oral argument, Patent Owner filed Objections to Patent Owner's Demonstratives. Paper 51 ("PO Obj."). Specifically, Patent Owner objects to certain slides as "misleading" (slides 3, 9); "irrelevant" and "unduly prejudicial" (slides 35, 36, 80, 84, 96); "mischaracterize[ing]" and/or "incomplete" (slides 37, 45, 48, 53, 54, 60, 61, 67, 68, 72, 76, 95, 100). PO Obj. 2–5.[14]

We have considered these objections. However, demonstratives are not evidence, and we weigh the material referenced in the demonstratives based on its probative value. *See* Paper 47 ("Order Setting Oral Argument") at 3 ("Demonstratives also are not evidence, and will not be relied upon as evidence."). We do not rely on the demonstratives in rendering this Final Written Decision, and, therefore, we dismiss Patent Owner's objections as *moot*.

VI. CONCLUSION

For the foregoing reasons, we are persuaded that Petitioner established by a preponderance of the evidence that claims 1–19 of the '867 patent are unpatentable.

In summary:

Claims	35 U.S.C. §	Reference(s)/Basis	Claims Shown Unpatentable	Claims Not shown Unpatentable
1, 2, 4–8, 13–19	103(a)	Zhang, Gupta	1, 2, 4–8, 13–19	
3, 9–12	103(a)	Zhang, Gupta, Chien	3, 9–12	

[14] Patent Owner's Objections have no page numbers. We begin page numbering with the title page as page 1.

Overall Outcome			1–19	

In addition, we *deny* Patent Owner's Motion to Amend the '867 patent, as shown in the following table:

Motion to Amend Outcome	Claim(s)
Original Claims Cancelled by Amendment	
Substitute Claims Proposed by Amendment	20–38
Substitute Claims: Motion to Amend Granted	
Substitute Claims: Motion to Amend Denied	20–38
Substitute Claims: Not Reached	

VII. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that claims 1, 2, 4–8, 13–19 of the '867 patent have been shown to be unpatentable under 35 U.S.C. § 103(a) as having been obvious over Zhang and Gupta;

FURTHER ORDERED that claims 3 and 9–12 of the '867 patent have been shown to be unpatentable under 35 U.S.C. § 103(a) as having been obvious over Zhang, Gupta, and Chien;

FURTHER ORDERED that Patent Owner's Revised Motion to Amend is *denied* as to substitute claims 20–38; 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.

For PETITIONER:

Brian Nash
Evan Finkel
Matthew Hindman
PILLSBURY WINTHROP SHAW PITTMAN LLP
brian.nash@pillsburylaw.com
evan.finkel@pillsburylaw.com
matthew.hindman@pillsburylaw.com

David Hoffman
Kenneth Darby
FISH & RICHARDSON P.C.
hoffman@fr.com
kdarby@fr.com

For PATENT OWNER:

Jay Kesan
DIMUROGINSBERG, PC DGKEYIP GROUP
jay@jaykesan.com

Ari Rafilson
SHORE CHAN DEPUMPO LLP
arafilson@shorechan.com