Trials@uspto.gov
Tel: 571-272-7822

Paper No. 22
Entered: May 10, 2019

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

AMAZON WEB SERVICES, INC., AMAZON.COM, INC., and VADATA, INC.,
Petitioner,

v.

SAINT REGIS MOHAWK TRIBE,
Patent Owner.

Case IPR2019-00103
Patent 7,149,867 B2

Before KALYAN K. DESHPANDE, JUSTIN T. ARBES, and CHRISTA P. ZADO, *Administrative Patent Judges*.

ZADO, *Administrative Patent Judge*.

DECISION
Denying *Inter Partes* Review
35 U.S.C. § 314

I. INTRODUCTION

A. Overview

Amazon Web Services, Inc., Amazon.com, Inc., and VADATA, Inc. (collectively, "Petitioner")[1] filed a petition requesting *inter partes* review of claims 1, 3–9, and 11–19 (the "challenged claims") of U.S. Patent No. 7,149,867 B2 (Ex. 1001, "the '867 patent"). Paper 1 ("Pet."). Saint Regis Mohawk Tribe ("Patent Owner")[2] filed a Preliminary Response. Paper 20 ("Prelim. Resp.").

35 U.S.C. § 314 provides that an *inter partes* review must not be instituted "unless . . . the information presented in the petition . . . shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition." 35 U.S.C. § 314(a). Upon considering the evidence and arguments presented, we determine the Petition does not demonstrate a reasonable likelihood that Petitioner would prevail with respect to at least one of the challenged claims.

Accordingly, we do not institute an *inter partes* review.

B. Related Proceeding

The parties advise that the '867 patent has been subject to, or relates to, the following district court proceeding: *SRC Labs and Saint Regis Mohawk Tribe v. Amazon Web Services, Inc., Amazon.com, Inc., and VADATA, Inc.*, No. 2:18-cv-00317 (W.D. Wash.). Pet. 2; Paper 17, 1.

---

[1] Petitioner identifies only itself as real parties-in-interest to the Petition. Pet. 1–2.

[2] Patent Owner identifies only itself as a real party-in-interest to this proceeding. Paper 17, 1.

C. The ’867 Patent

The ’867 patent, titled “System and Method of Enhancing Efficiency and Utilization of Memory Bandwidth in Reconfigurable Hardware,” generally relates to “implementing explicit memory hierarchies in reconfigurable processors that make efficient use of off-board, on-board, on-chip storage and available algorithm locality.” Ex. 1001, 1:18–21.

The ’867 patent explains that microprocessors “enjoyed annual performance gains averaging about 50% per year,” wherein most of the gains were attributable to higher clock processor speeds, more memory bandwidth, and increasing utilization of instruction level parallelism (“ILP”) at execution time. *Id.* at 1:26–30. However, as microprocessor speeds increased, challenges arose to designing memory hierarchies that could keep up. *Id.* at 1:31–33. The ’867 patent identifies two measures of the gap between microprocessor and memory hierarchy speeds—bandwidth efficiency and bandwidth utilization. *Id.* at 1:35–37. Because potential performance gains from using a faster microprocessor were reduced or negated by corresponding drops in bandwidth efficiency and bandwidth utilization, significant effort had been spent, according to the ’867 patent, on development of memory hierarchies that could maintain high bandwidth efficiency and utilization. *Id.* at 1:45–50.

The ’867 explains that one approach to bridging the gap was the utilization of cache memories. *Id.* at 1:51–53. In designing cache memories, a number of considerations had to be taken into account. *Id.* at 59–60. For example, for programs that exhibit a high degree of spatial locality (i.e., it is likely that other data within the same cache line will be needed), wide cache lines are efficient. *Id.* at 1:64–2:4. However, for programs that have low levels of spatial locality, narrow cache lines are

more efficient. *Id.* at 2:4–7. The ’867 patent provides additional examples of considerations in cache design. *Id.* at 2:14–3:40. The ’867 patent states that the various considerations and tradeoffs made cache design challenging for a multipurpose computer that executes a wide variety of programs. *Id.* at 3:30–32. Cache designers tried to derive the program behavior of the “average” program, and optimize the cache for the “average” program. *Id.* at 3:32–36. As a result, the cache was sub-optimal for most programs, because most programs that actually run on the microprocessor are not “average.” *Id.* at 3:36–39.

Because of the above-discussed issues, there was a growing need, according to the ’867 patent, to develop improved memory hierarchies that limited overhead of a memory hierarchy without also reducing bandwidth efficiency and utilization. *Id.* at 3:57–60. To address the need, 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. “Unlike conventional static hardware platforms,” the memory hierarchy is reconfigurable so that computational demands and memory bandwidth can be matched. *Id.* at 7:17–22. 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 an example of configuring the data prefetch unit depending on the needs of the computational logic. For

example, Figures 9A and 9B show an external memory organized into a 128 byte (16 word) block structure that is optimized for stride 1 access of a cache. *Id.* at 7:56–59. However, the data prefetch unit can be configured to extract only 8 bytes of data in the memory block, discarding the remaining 120 bytes if only the 8 bytes are needed. *Id.* at 8:3–11. In another example relating to a computational intensive matrix multiplication problem, the '867 patent explains that

> On a conventional microprocessor with static execution resources, these loops [representing matrix multiplication] would be arranged to give stride-one data access where possible and also block or tile these uses to facilitate data cache hits on the B and A matrices, which are read many times. With the configurable memory hierarchy of the present invention, matrix B may be stored in on-board BRAM memory 307 and rows of matrix A in registers.

Id. at 10:33–40.

D. Asserted Grounds of Unpatentability

Petitioner challenges claims 1, 3–9, and 11–19 of the '867 patent on the following grounds. Pet. 3.

Reference	Ground	Claims
Lange[3]	§ 103(a)	1, 3–9, 11–19
Zhong[4]	§ 103(a)	1, 4, 6, 7, 9

[3] Holger Lange & Andreas Koch, "Memory Access Schemes for Configurable Processors," *Field-Programmable Logic and Applications: The Roadmap to Reconfigurable Computing*, 10th International Conference, FPL 2000, Villach, Austria, 615–25 (Aug. 27–30, 2000) (Ex. 1003) ("Lange").

[4] Peixin Zhong & Margaret Martonosi, "Using Reconfigurable Hardware to Customize Memory Hierarchies," *High-Speed Computing, Digital Signal Processing, and Filtering Using Reconfigurable Logic*, SPIE—The International Society for Optical Engineering, Boston, MA, 237–248 (Nov.

Petitioner relies on the declaration of Brad L. Hutchings, Ph.D., to support the Petition. Ex. 1002 ("Hutchings Declaration").

E. Challenged Claims

Of the challenged claims, claims 1, 9, and 13 are independent. Claim 1, reproduced below, is illustrative.

> 1. 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.

II. DISCUSSION

A. Level of Ordinary Skill

Petitioner asserts that a person of ordinary skill in the art in the field of the '867 patent in the relevant time frame would have had a bachelor's degree in electrical engineering, computer engineering, or a related field, with two to three years of experience working with reconfigurable systems. Pet. 3 (citing Ex. 1002 ¶ 24). Petitioner asserts that "[w]ith more education,

---

20–21, 1996) (Ex. 1004) ("Zhong").

such as additional graduate degrees or study, less experience is needed to attain the ordinary level of skill." *Id.*

The Preliminary Response provides no assessment of the level of ordinary skill in the art.

For purposes of this decision and based on the record before us, we adopt Petitioner's assessment of the level of ordinary skill in the art.

B. Claim Construction

In an *inter partes* review involving a petition filed before November 13, 2018, claim terms in an unexpired patent are given their broadest reasonable construction in light of the specification of the patent. 37 C.F.R. § 42.100(b) (2016). Consistent with this standard, we assign claim terms their ordinary and customary meaning, as would be understood by one of ordinary skill in the art at the time of the invention, in the context of the entire patent disclosure. *See In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). Only those terms that are in controversy need be construed, and only to the extent necessary to resolve the controversy. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (citing *Vivid Techs., Inc. v. Am. Sci. & Eng'g, Inc*., 200 F.3d 795, 803 (Fed. Cir. 1999)). We determine that the term "data prefetch unit" requires construction.

Each of the challenged independent claims recites a "data prefetch unit." Claim 1 recites

> 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

Ex. 1001, 12:43–54.

Claim 9 recites one or more reconfigurable processors,

> 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 wherein the data prefetch unit operates independent of and in parallel with logic blocks using the computational data, and wherein the data prefetch unit is configured to conform to needs of the algorithm and match format and location of data in the common memory.

Id. at 13:17–26.

Claim 13 recites "transferring data between a memory and a data prefetch unit in a reconfigurable processor," *id.* at 14:2–3, and further recites that the data prefetch unit is

> 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 wherein the prefetch unit operates independent of and in parallel with the computational unit.

Id. at 14:6–11.

Petitioner proposes to construe the term "data prefetch unit" as "a functional unit that retrieves computational data needed to complete the algorithm instantiated on the reconfigurable processor during processing." Pet. 6–8. Petitioner states that this construction was proposed by Patent Owner in the related district court proceeding, and asserts that for purposes of the Petition, Patent Owner's construction should be used. *Id.* Petitioner does not explain why this construction is correct, or provide any arguments

or evidence to support this claim construction. *See generally id.*; Prelim. Resp. 17.

Patent Owner responds that Petitioner's proposed construction is incorrect. Prelim. Resp. 16–17. Patent Owner argues that the '867 patent expressly defines the term "data prefetch unit," and therefore the term should be construed in accordance with the express definition provided in the patent. *Id.* Patent Owner asserts that the '867 patent provides a heading labeled "Definitions" in the Detailed Description, and under this heading, defines "data prefetch unit." *Id.* at 16. Patent Owner argues, therefore, that "the patentee has clearly set forth a definition of the disputed term with reasonable clarity, deliberateness, and precision." *Id.* at 17.

When the specification of a patent provides a special definition for a claim term, even if it differs from the term's ordinary meaning, then the inventor's lexicography governs. *In re Morris*, 127 F.3d 1048, 1054 (Fed. Cir. 1997) (applying "the broadest reasonable meaning of the words in their ordinary usage as they would be understood by one of ordinary skill in the art, taking into account whatever enlightenment by way of definitions or otherwise that may be afforded by the written description contained in the applicant's specification"); *see also Thorner v. Sony Computer Entm't Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012) ("To act as its own lexicographer, a patentee must 'clearly set forth a definition of the disputed claim term,'" and "'clearly express an intent' to redefine the term.").

We are persuaded that the '867 patent clearly sets forth a definition for the term "data prefetch unit." The '867 patent provides an express definition for "data prefetch unit" under the heading "Definitions," thereby indicating the patentee intended to accord a special definition to the term. Ex. 1001, 5:18, 5:40–43. The definition provides "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," wherein a memory hierarchy "is a collection of memories." *Id.* at 5:39–43.

Petitioner's argument that its construction is the same as that proposed by Patent Owner in district court is not persuasive. Pet. 6–8. Petitioner does not cite any intrinsic or extrinsic evidence to support its proposed construction, much less explain why its construction is correct. *See generally id.* Petitioner has not explained, nor do we discern, a reason to deviate from the express definition for "data prefetch unit" provided in the '867 patent.

Therefore, on the record before us, we construe "data prefetch unit" in accordance with the definition set forth in the '867 patent, namely as "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 unite stride memory," wherein a "memory hierarchy" is "a collection of memories." Ex. 1001, 5:39–43.

C. Lange (Ex. 1003)

Lange, titled "Memory Access Schemes for Configurable Processors," generally describes a scalable, device-independent memory interface that supports both irregular access (via configurable caches) and regular access (via pre-fetching stream buffers). Ex. 1003, 615. Lange states that "[b]y hiding specifics behind a consistent abstract interface, it is suitable as a target environment for automatic hardware compilation." *Id.* Lange explains that reconfigurable compute elements can achieve considerable performance gains over standard central processing units ("CPUs"). *Id.* According to Lange, these reconfigurable elements often are combined with

a conventional processor, which provides control and I/O services that are more efficiently implemented with fixed logic. *Id.* In combined systems, design tools address hardware and software issues separately. *Id.* According to Lange, whereas the level of support for software is suitable, the same level of support is not provided for hardware. *Id.* Lange states that it therefore presents a "hardware target" for hardware compilers that is analogous to a software target for conventional computers. *Id.* The hardware target is a Memory Architecture for Reconfigurable Computers ("MARC"). *Id.* Figure 4 of Lange, reproduced below, shows an overview of the MARC architecture.

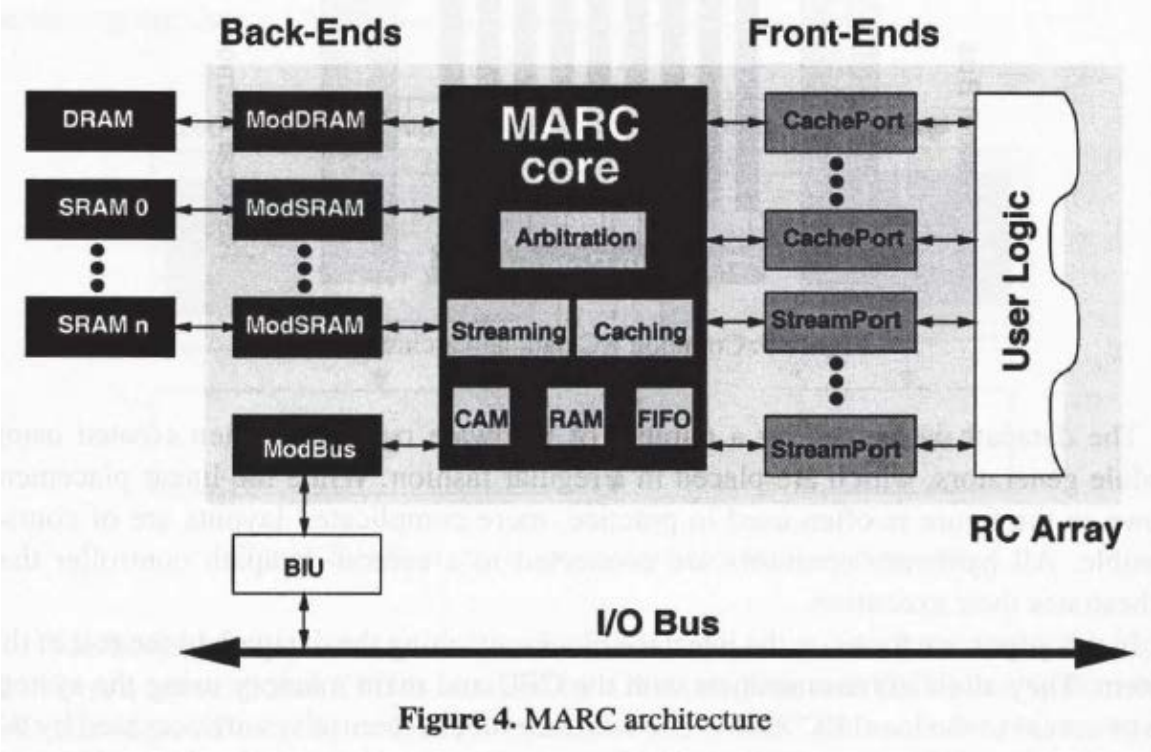

Figure 4. MARC architecture

Id. at 618. Figure 4 shows a MARC core with a caching port interfaced with front-end ports (CachePorts and StreamPorts) that interface with User Logic, and a streaming port interfaced with back-end ports interfaced with dynamic random access memory (DRAM) and n static random access memories (SRAMs). *Id.* The MARC core also interfaces, through a Back-End port,

with a bus interface unit ("BIU") to an I/O bus. *Id.* The MARC core includes a First-In First-Out ("FIFO") buffer and Random Access Memory ("RAM").

D. Zhong (Ex. 1004)

Zhong, titled "User Reconfigurable Hardware to Customize Memory Hierarchies," generally describes implementing mechanisms like victim caches and prefetch buffers in reconfigurable hardware to improve application memory behavior. Ex. 1004, 237. Zhong states that microprocessor speeds have increased much more quickly than memory speeds. *Id.* As a result, there is a processor-memory performance gap such that many significant applications suffer from substantial memory bottlenecks, according to Zhong. *Id.* Zhong explains that typically cache memories are used to bridge the performance gap, but that cache memory still fails to provide high performance for certain applications. *Id.* To address issues with cache performance, Zhong states that prefetching techniques and use of victim caches (e.g., memory for storing data recently evicted from cache) may hide some latencies, but that these techniques result in waste of transistor space on CPU chips. *Id.*; *see also id.* at 239 (describing victim caches). Zhong proposes to address these issues by using programmable logic, such as field-programmable gate arrays (FPGAs), that can be reconfigured and customized for different functions during different sessions. *Id.* at 237. Part of Figure 1 of Zhong is reproduced below.

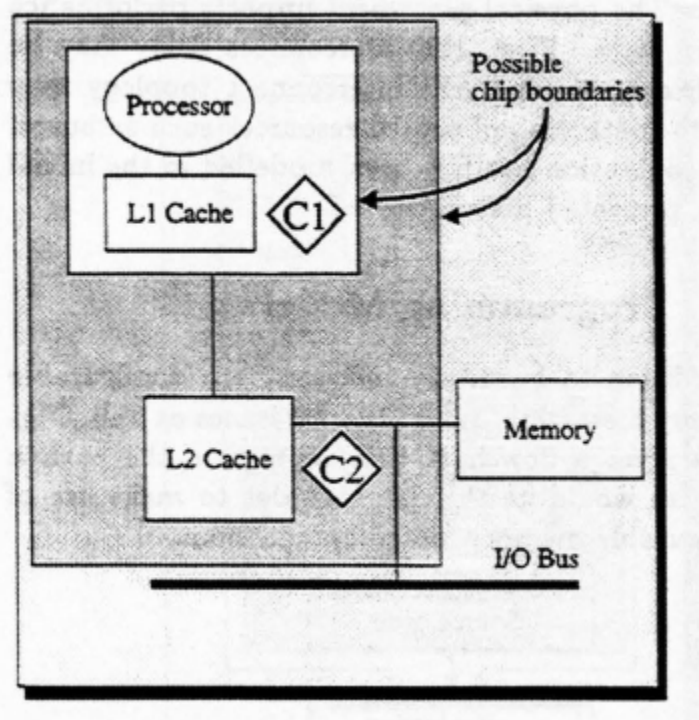

Id. at 239. The portion of Figure 1 reproduced above illustrates a computer architecture that includes configurable logic C1 on the same chip as a conventional Processor. *Id.* Applying one possible chip boundary, the chip is shown as including a Processor, C1, and an L1 cache, whereas the L2 cache and additional configurable processor C2 are off-chip. *Id.* The figure also shows an alternative chip boundary, in which the chip also includes the L2 cache and C2. *Id.* In both alternatives, the L2 cache is connected to Memory and I/O Bus. *Id.*

Zhong also discloses a prefetch buffer "to initiate main memory accesses in advance, so that the data will be closer to the processor when referenced." *Id.* at 240–241. The prefetch buffer comprises several independent slots, each of which holds several cache lines of data and works like a FIFO buffer. *Id.* at 241.

E. Principles of Law

Section 103(a) forbids issuance of a patent when "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." In *Graham v. John Deere Co.*, 383 U.S. 1 (1966), the Court set out a framework for applying the statutory language of § 103: under § 103, the scope and content of the prior art are to be determined; differences between the prior art and the claims at issue are to be ascertained; and the level of ordinary skill in the pertinent art resolved.

The Supreme Court has made clear that we apply "an expansive and flexible approach" to the question of obviousness. *KSR Int'l Co. v. Teleflex, Inc.*, 550 U.S. 398, 415 (2007). Whether a patent claiming the combination of prior art elements would have been obvious is determined by whether the improvement is more than the predictable use of prior art elements according to their established functions. *KSR Int'l Co.*, 550 U.S. at 417. Reaching this conclusion, however, requires more than a mere showing that the prior art includes separate references covering each separate limitation in a claim under examination. *Unigene Labs., Inc. v. Apotex, Inc.*, 655 F.3d 1352, 1360 (Fed. Cir. 2011). Rather, obviousness requires the additional showing that a person of ordinary skill at the time of the invention would have selected and combined those prior art elements in the normal course of research and development to yield the claimed invention. *Id.*

F. Patentability

As we discussed above, *supra* Sec. II.B, each of the challenged independent claims recites a "data prefetch unit." Petitioner's arguments, however, are based on a construction that we do not adopt. Petitioner does

not provide any arguments under the construction set forth above. For the reasons discussed below, Petitioner has not demonstrated a reasonable likelihood that Petitioner would prevail with respect to at least one of the challenged claims.

1. Asserted Obviousness over Lange

Petitioner asserts that, for purposes of the Petition, "the broadest reasonable interpretation of 'a data prefetch unit' is 'a functional unit that retrieves computational data needed to complete the algorithm instantiated on the reconfigurable processor during processing.'" Pet. 17. However, as we discussed above, we interpret "data prefetch unit," in accordance with the definition set forth expressly in the '867 patent, as "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." *Supra* Sec. II.B; *see also* Ex. 1001, 5:40–43. In addition, in accordance with the '867 patent's express disclosure, we interpret a "memory hierarchy" as "a collection of memories." *Supra* Sec. II.B; Ex. 1001, 5:39.

Patent Owner argues that the Petition is based on an erroneous claim construction. Prelim. Resp. 22. Patent Owner argues that "[t]he Board is not required to 'play archeologist with the record' or endeavor to discover a challenge that might have been asserted had the Petitioner identified the correct claim construction." *Id.* at 22 (citing *United Microelectronics Corp. v. Lone Star Silicon Innovations LLC*, Case IPR2017-01513, slip op. at 9 (PTAB May 22, 2018) (Paper 10)). We agree.

Applying its proposed construction of "data prefetch unit," Petitioner argues that Lange's MARC core with its front-end port interfaces incorporates data prefetch units. Pet. 17. Petitioner argues that the MARC

core, when used with the front-end ports, performs the function of prefetching the computational data needed to complete the algorithm instantiated in Lange's user logic. *Id.* at 17–18.

The Petition, however, does not specify with particularity how Lange teaches a memory hierarchy, and moving data between members of a memory hierarchy, as required under our interpretation of the term "data prefetch unit." Our rules require that a petition specify with particularity where each element of a claim is found in the prior art, and include a detailed explanation of the relevance of the prior art to the claim. 37 C.F.R. § 42.104(b)(4) ("[t]he petition must specify where each element of the claim is found in the prior art patents or printed publications relied upon"); *id.* § 42.22(a)(2) ("[e]ach petition . . . must include . . . a detailed explanation of the significance of the evidence including material facts"); *id.* § 42.104(b)(5) (the petition must "identify . . . the relevance of the evidence to the challenge raised, including identifying specific portions of the evidence that support the challenge"). As the Federal Circuit has explained, "[i]n an IPR, the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable." *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016).

In its contentions regarding independent claims 1, 9, and 13, Petitioner does not identify a memory hierarchy in Lange, much less assert that Lange teaches moving data between members of a hierarchy. With regard to claim 1, Petitioner asserts that Lange discloses a first memory (i.e., either the FIFO memory in the MARC core or BlockSelectRAM in the FPGA) and a second memory (i.e., SRAM and/or DRAM accessed by the MARC core back-end ports), as recited in the claim, but Petitioner does not specify that these memories comprise a memory hierarchy or explain why

that would be the case. Pet. 15–17 (asserting a first memory); *id.* at 21–22 (asserting second memory); *see generally id.* at 15–22 (failing to specify a memory hierarchy). With regard to claim 9, Petitioner identifies a memory, asserting that Lange discloses a "common memory" (i.e., DRAM in Figure 5), as recited in the claim, but Petitioner does not specify a memory hierarchy comprising a collection of memories. *Id.* at 33–34. With regard to claim 13, Petitioner asserts that Lange discloses a "memory" (i.e., the second memory of claim 1), as recited in claim 13, but Petitioner does not specify a memory hierarchy comprising a collection of memories. *Id.* at 37–39.

Similarly, the Petition does not address whether Lange teaches moving data between members of a memory hierarchy. *See generally id.* at 13–26, 33–34, 37–39.

By failing to address whether Lange teaches a memory hierarchy, and movement of data between members of the memory hierarchy, Petitioner has placed the burden on the Board to ascertain how the prior art allegedly reads on the challenged claims—a task that we do not undertake. The burden is on Petitioner, not the Board, to specify with particularity how Lange teaches a memory hierarchy, and moving data between members of a memory hierarchy.

For the reasons stated above, on this record, we are not persuaded that Petitioner has demonstrated a reasonable likelihood that it will prevail in showing unpatentability of claims 1, 3–9, and 11–19 as obvious over Lange.

2. Asserted Obviousness over Zhong

As we discussed above, we interpret "data prefetch unit," in accordance with the definition set forth expressly in the '867 patent, as "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," wherein a "memory hierarchy" is "a collection of memories." *Supra* Sec. II.B; *see also* Ex. 1001, 5:39–43. The Petition, however, applies a different claim construction. Pet. 48. As we discussed above with regard to Lange, *supra* Sec. II.F, Patent Owner argues that the Petition is based on an erroneous claim construction. Prelim. Resp. 22. We agree with Patent Owner.

The Petition does not specify with particularity how Zhong teaches a memory hierarchy, and moving data between members of a memory hierarchy, as required under our interpretation of the term "data prefetch unit." As we discussed above, our rules require that a petition specify with particularity where each element of a claim is found in the prior art, and include a detailed explanation of the relevance of the prior art to the claim. 37 C.F.R. §§ 42.104(b)(4), 42.22(a)(2), 42.104(b)(5); *see also Harmonic*, 815 F.3d at 1363 (explaining that "[i]n an IPR, the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable").

Applying its proposed construction of "data prefetch unit," Petitioner relies on Zhong's disclosure of a prefetch generator depicted in Figure 4 of Zhong and Zhong's "prefetching engine" for disclosure of a "data prefetch unit." Pet. 48. However, Petitioner does not identify a memory hierarchy in Zhong, much less assert that Zhong teaches moving data between members of a hierarchy. With regard to claim 1, Petitioner asserts that Zhong discloses a first memory (i.e., prefetch buffers) and a second memory (i.e., main memory or L2 cache), as recited in the claim, but Petitioner does not specify that these memories comprise a memory hierarchy or explain why that would be the case. Pet. 47 (asserting a first memory); *id.* at 52 (asserting second memory); *see generally id.* at 46–53 (failing to specify a

memory hierarchy). With regard to claim 9, Petitioner identifies a memory, asserting that Zhong discloses a "common memory" (i.e., main memory in Zhong Figure 1), as recited in the claim, but Petitioner does not specify a memory hierarchy comprising a collection of memories. *Id.* at 60.

Similarly, the Petition does not address whether Zhong teaches moving data between members of a memory hierarchy. *See generally id.* at 44–56, 59–61.

By failing to address whether Zhong teaches a memory hierarchy, and movement of data between members of the memory hierarchy, Petitioner has placed the burden on the Board to ascertain how the prior art allegedly reads on the challenged claims—a task that we do not undertake. The burden is on Petitioner, not the Board, to specify with particularity how Zhong teaches a memory hierarchy, and moving data between members of a memory hierarchy.

For the reasons stated above, on this record, we are not persuaded that Petitioner has demonstrated a reasonable likelihood that it will prevail in showing unpatentability of claims 1, 4, 6, 7, and 9 as obvious over Zhong.

G. Additional Arguments by Patent Owner

Patent Owner argues that we should exercise our discretion to deny the Petition under 35 U.S.C. § 314(a). Prelim. Resp. 5–12. Patent Owner also argues that we should deny the Petition for failure to satisfy the requirement of 37 C.F.R. § 42.104(b)(3) that the petition set forth how the challenged claims are to be construed. *Id.* at 18–21. Because we deny the Petition on other grounds, we need not, and do not, address Patent Owner's arguments regarding § 314(a) and § 42.104(b)(3).

III. CONCLUSION

For the foregoing reasons, we determine that Petitioner has not demonstrated a reasonable likelihood it will prevail in showing unpatentability of at least one claim of the ’867 patent. Because Petitioner has not satisfied the threshold for institution as to at least one claim, we do not institute *inter partes* review.

IV. ORDER

Accordingly, it is

ORDERED that the Petition is denied and no trial is instituted.

PETITIONER:

J. David Hadden
Saina Shamilov
FENWICK & WEST LLP
dhadden-ptab@fenwick.com
sshamilov-ptab@fenwick.com

PATENT OWNER:

Alfonso Chan
Joseph DePumpo
SHORE CHAN DEPUMPO LLP
achan@shorechan.com
jdepumpo@shorechan.com