

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MERCEDES-BENZ USA, LLC,
Petitioner,

v.

DAEDALUS PRIME LLC,
Patent Owner.

IPR2023-01344
Patent 9,575,895 B2

Before JAMES P. CALVE, GREGG I. ANDERSON, and
ARTHUR M. PESLAK, *Administrative Patent Judges*.

ANDERSON, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision
Determining All Challenged Claims Unpatentable
35 U.S.C. § 318(a)
Denying as Moot Patent Owner's Motion to Exclude
37 C.F.R. § 42.64

I. INTRODUCTION

A. Background

Mercedes-Benz USA, LLC (“Petitioner”) filed a Petition requesting *inter partes* review of claims 1–17 (the “challenged claims”) of U.S. Patent No. 9,575,895 B2 (Ex. 1001, the “’895 patent”). Paper 2 (“Pet.”). Daedalus Prime LLC (“Patent Owner”) filed a Preliminary Response. Paper 6 (“Prelim. Resp.”). We instituted *inter partes* review. Paper 7 (“Inst. Dec.”). Patent Owner filed a Response (Paper 10, “PO Resp.”), Petitioner filed a Reply (Paper 13, “Reply”), and Patent Owner filed a Sur-Reply (Paper 20, “Sur-Reply”).

Patent Owner filed Motion to Exclude Evidence (Paper 24, “Motion,” “Mot.”), Petitioner filed an Opposition to the Motion (Paper 25, “Opposition,” “Opp.”), and Patent Owner filed a Reply to the Opposition (Paper 26, “Reply to Opposition,” “Opp. Reply”). We deny the Motion as moot.

A hearing was held December 9, 2024, and a transcript has been made of record. Paper 30 (“Tr.”).

We have jurisdiction under 35 U.S.C. § 6. This decision is a Final Written Decision issued pursuant to 35 U.S.C. § 318(a). For the reasons we discuss below, we determine that Petitioner has proven by a preponderance of the evidence that all of the challenged claims of the ’895 patent are unpatentable.

B. Related Proceedings

The parties identify the following district court and ITC proceedings involving the '895 patent: (1) *Daedalus Prime LLC v. Arrow Electronics, Inc.*, 1:22-cv-01107 (D. Del.); (2) *Daedalus Prime LLC v. Mazda Motor Corporation*, 1:22-cv-01108 (D. Del.); (3) *Daedalus Prime LLC v. Mazda Motor Corporation*, 1:22-cv-01109 (D. Del.); and (4) *In re Certain Semiconductors and Devices and Products Containing the Same, Including Printed Circuit Boards, Automotive Parts, and Automobiles*, Inv. No. 337-TA-1332 (ITC August 23, 2022). Pet. 3; Paper 3 (“Patent Owner’s Mandatory Notices”), 2.

C. Real Parties-in-Interest

Petitioner identifies the following real parties-in-interest: Mercedes-Benz USA, LLC; Mercedes-Benz Intellectual Property GmbH & Co. KG; Mercedes-Benz Group AG; and Mercedes-Benz AG. Pet. 2. Patent Owner identifies Daedalus Prime LLC as the real party-in-interest. Paper 3, 2

D. The '895 Patent

The '895 patent was filed January 30, 2015. Ex. 1001 at code [22]. The '895 patent is a continuation of Application Ser. No. 13/324,053, filed December 13, 2011, now U.S. Pat. No. 8,984,228. *Id.* at code [63].

1. Technical Background

The '895 patent relates to multiprocessor (MP)-socket computer systems, which may include different “semiconductor components realized as integrated circuits (ICs).” Ex. 1001, 1:11–16. “The ICs include processors, memories, chipsets, input/output hubs (IOHs). . . .” *Id.*

Problems arise once an input/output (IO) component is integrated on the same chip with a multiprocessor because the IO has “a separate caching agent.” Ex. 1001, 1:25–27. This means the dedicated logic associated with the IO component has to “snoop” the central processing unit (CPU) to maintain cache coherency. *Id.* at 1:20–22, 1:28–32. In MP systems, this becomes a major scaling problem, resulting in degraded performance. *Id.* at 1:32–36.

2. *The Multicore Processor of the '895 Patent*

The '895 patent describes a solution for the problem of integrating an IO component on the same chip with a multiprocessor in order to avoid scaling issues. Ex. 1001, 1:59–61. The solution relates to configuring a CPU caching agent to support both CPU traffic and IO traffic. *Id.* at 1:62–64. Instead of treating the IO component as a separate caching agent, the '895 patent employs “a CPU caching agent [] configured to support both CPU traffic and IO traffic as well, thereby resolving the scaling issue.” *Id.*

“In modern multiprocessor (MP)-socket computer systems, various topologies are possible.” Ex. 1001, 1:11–12. Referring to Figure 1, a socket can include “4, 8, or another such number of cores” where “one or more levels of cache memories can be present within the cores.” *Id.* at 2:15–21. “Each of cores 120 may be coupled to a shared cache memory 130₀, which may be a last level cache (LLC).” *Id.* at 2:21–23. “LLC 130 can include a caching agent 135₀.” *Id.* at 2:23–24.

Figure 2 of the '895 patent is reproduced below.




FIG. 2

Figure 2 is a block diagram of a caching agent.

Ex. 1001, 2:41–43. Although a “single structure” is shown in Figure 2, “a caching agent can be distributed such that each of the different portions of the caching agent can be associated with a corresponding core and LLC bank or slice.” *Id.* at 3:6–10.

With this approach, the IIO module [(integrated IO module)] proxies through the CPU caching agent to access memory or other IO devices, therefore reducing the overhead of allocating dedicated resources for an integrated IO caching agent. This also reduces the amount of snoop traffic needed since a reduced number of caching agents per system can be realized. Thus in various embodiments, a system can include a single caching agent per multicore processor socket, where each socket includes multiple cores and an IIO module.

Id. at 3:11–19. “[C]aching agent 200 may be logic interposed between one or more cores of a processor and a LLC 230.” *Id.* at 2:43–45.

Figure 6 of the '895 patent is reproduced below.

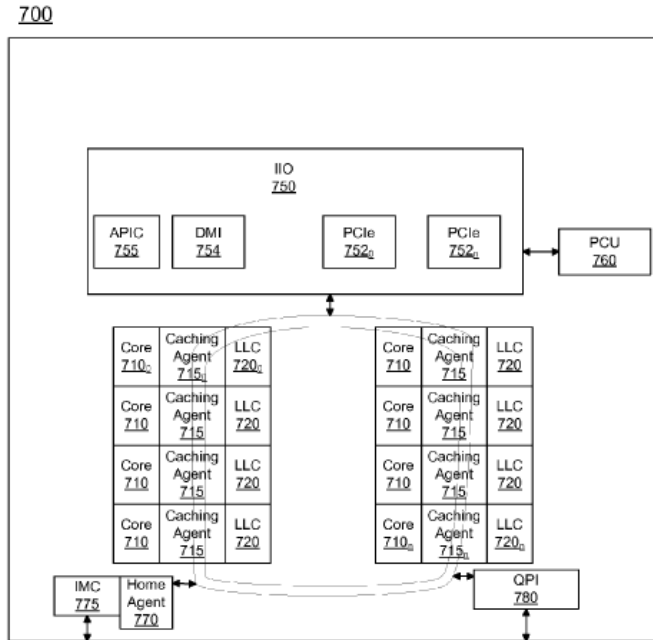


FIG. 6

Figure 6 is a block diagram of one embodiment of a multicore processor.

Ex. 1001, 6:30–33. As shown in Figure 6, the “processor 700 includes a distributed configuration having partitions or slices each including a core 710 and a partition of a caching agent 715 and a LLC 720.” *Id.* at 6:34–36. The ’895 patent explains “that while distributed caching agents are shown, understand that these distributed portions form a single caching agent, and which is configured to handle cache coherency operations both for the cores as well as an I/O module 750.” *Id.* at 6:35–39.

E. Illustrative Claim

The Petition challenges claims 1–17, of which claims 1, 11, and 15 are independent claims to, respectively, a processor, a non-transitory

computer readable medium, and a system. Claim 1 is illustrative of the claimed subject matter and reproduced below:¹

1. [1p] A processor comprising:
 - [1a] a plurality of cores,
 - [1b] a shared cache memory,
 - [1c] a memory controller to interface with a memory coupled to the processor,
 - [1d] an integrated input/output (IIO) module to interface between the processor and an IO device coupled to the processor and
 - [1e] a caching agent to perform cache coherency operations for the plurality of cores and the IIO module,
 - [1f] wherein the processor is to receive an allocation transaction from the IO device and directly store data of the allocation transaction into the shared cache memory,
 - [1g] wherein the caching agent is a single caching agent for the processor and includes a plurality of distributed portions each associated with a corresponding one of the plurality of cores.

Ex. 1001, 7:48–61.

F. References and Other Evidence

The Petition relies on the following references:

Name²	Reference	Publication Date³	Exhibits
Sinharoy	<i>IBM POWER7 multicore server processor</i> , IBM Journal	April 29, 2011	Ex. 1003

¹ Paragraph labeling in brackets is based on those provided by Petitioner.

² First named author or inventor only.

³ Date of publication for Exhibits 1003, 1004, and 1005 based on three Declarations of Gordon MacPherson, Exhibits 1009, 1011, and 1013, respectively. *See* Ex. 1009, 1; Ex. 1011, 1; Ex. 1013, 1. Patent Owner does not contest the publication dates or the prior art status of the references. *See generally* PO Resp.

Name ²	Reference	Publication Date ³	Exhibits
	of Research and Development, Vol. 55, Issue 3, Pages 1:1–29		
Tang	<i>DMA Cache: Using On-Chip Storage to Architecturally Separate I/O Data from CPU Data for Improving I/O Performance</i> , The Sixteenth International Symposium on High-Performance Computer Architecture	January 9–14, 2010	Ex. 1004
Xu	<i>Explorations of Optimal Core and Cache Placements for Chip Multiprocessor</i> , 2011 NORCHIP conference	November 14–15, 2011	Ex. 1005
Harikumar	US 2009/0164739 A1	June 25, 2009	Ex. 1006

In addition, Petitioner submits the Declaration of Dr. Robert Horst. (Ex. 1007, “Horst Declaration”) and Supplementary Declaration of Dr. Robert W. Horst in Support of Petitioner’s Reply (Ex. 1017, “Horst Supplementary Declaration”). Patent Owner submits the Declaration of Michael C. Brogioli, Ph.D. in Support of Patent Owner’s Response. Ex. 2002 (“Brogioli Declaration”). The parties have also submitted deposition transcripts for Dr. Horst and Dr. Brogioli. Exs. 2001 (“Horst Deposition”), 2003 (“Horst Reply Deposition”), 1018 (“Brogioli Deposition”).

G. Asserted Grounds of Unpatentability

Petitioner asserts the challenged claims are unpatentable on the following grounds:

Claims Challenged	35 U.S.C. §⁴	References
1, 3–7, 11–13, 15–17	102/103	Sinharoy
2, 10	103	Sinharoy, Xu
8, 9, 14	103	Sinharoy, Tang
1, 5–9, 11–15, 17	102/103	Tang
2–4, 10, 16	103	Tang, Harikumar

See Pet. 4.

H. Overview of the Prior Art

1. Sinharoy (Ex. 1003)

Sinharoy discloses a POWER7 server processor, which has eight cores, two memory controllers, an L2 cache, and an L3 cache. Ex. 1003, 2, 18–19. Figure 1 of Sinharoy is reproduced below.

⁴ Because the earliest application from which the '895 patent claims priority was filed before March 16, 2013, the pre-AIA (“America Invents Act”) version of § 103 applies. Leahy-Smith America Invents Act (“AIA”), Pub. L. No. 112-29, 125 Stat. 284, 285–88 (2011).

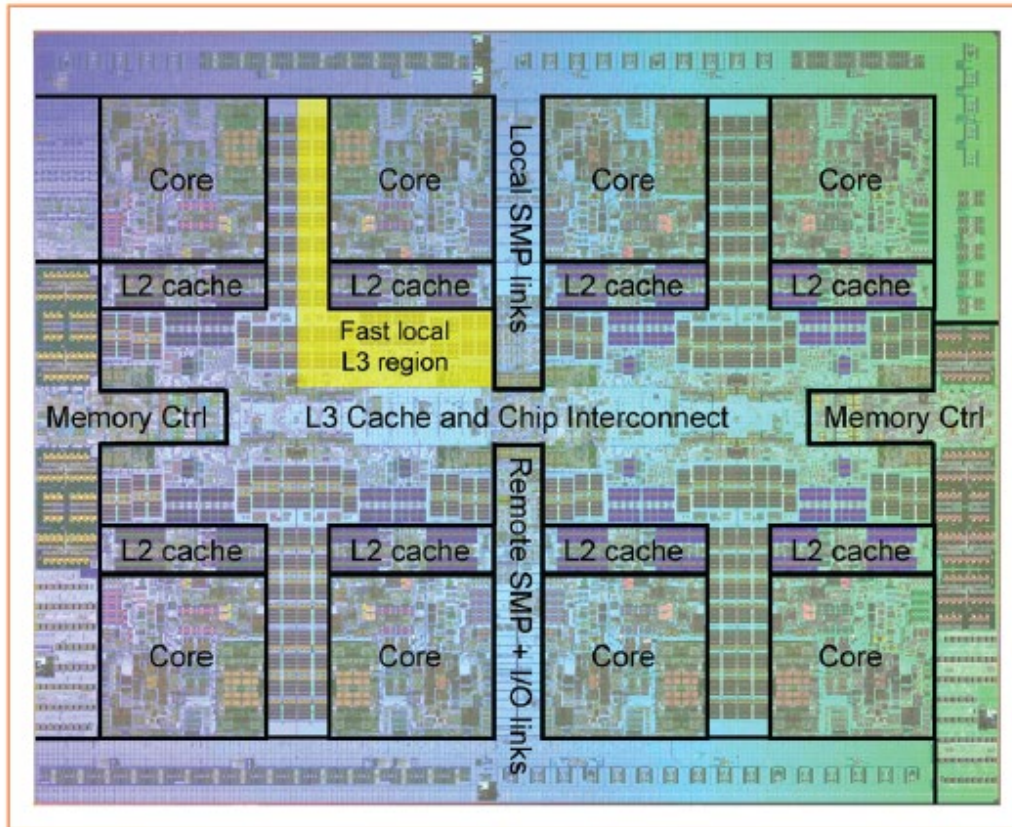


Figure 1

Die photo of the IBM POWER7 chip.

Figure 1 shows the POWER7 chip, which has eight processor cores, each with 12 execution units capable of simultaneously running four threads.

Ex. 1003, 2. Referring to Figure 1, Sinharoy discloses that the POWER7 processor has a “256-KB POWER7 L2 cache” and a “shared 32-MB POWER7 L3 cache,” “composed of eight 4-MB L3 regions.” *Id.* at 18–19, Fig. 1. The L2 cache and L3 cache each has a single centralized controller to manage the read/write request operations. *Id.* at 18–19.

The POWER7 processor supports “two integrated I/O controllers.” Ex. 1003, 22. Each of the I/O controllers supports a “4-byte off-chip read interface and 4-byte off-chip write interface, thereby connecting the POWER7 chip to up to two I/O hub chips.” *Id.* The I/O hub chips, PCI bridge, and PCI adapter are examples of input/output (I/O) devices that can be connected to the POWER7 processor via the integrated I/O controllers. *Id.* at 26, Fig. 14.

2. Tang (Ex. 1004)

Tang discloses systems and methods to directly inject or store I/O data into a cache memory because “moving I/O data in/out memory” exhausts a portion of memory bandwidth as well as accounts for a considerable part of I/O operation latency. Ex. 1004 § 1. Tang describes two designs of a DMA⁵ cache, the decoupled DMA cache (DDC) and the Partition-based DMA Cache (PBDC). *Id.* §§ 3.1.1 (DDC), 3.2.1 (PBDC). The DDC design has a separate DMA cache while the PBDC allows DMA from an I/O device to a portion of the last level cache (LLC). *Id.* In the PBDC design, Tang discloses a multiprocessor system using “several ways”⁶ of the processor’s last level cache (LLC) as the “dedicated DMA Cache.” *Id.* § 3.2.1.

⁵ The Horst Declaration explains that a person of ordinary skill in the art would understand that DMA stands for Direct Memory Access, and in this context refers to the ability of an I/O device to directly write data into the cache (also known as cache injection). Ex. 1007 ¶ 115.

⁶ A “way” or “bank” is a portion of a cache memory. Ex. 1007 ¶ 115 n.16. Memory lines are stored in a cache and can only be stored in one place in each way or bank. *Id.*

Figure 10 of Tang is reproduced below.

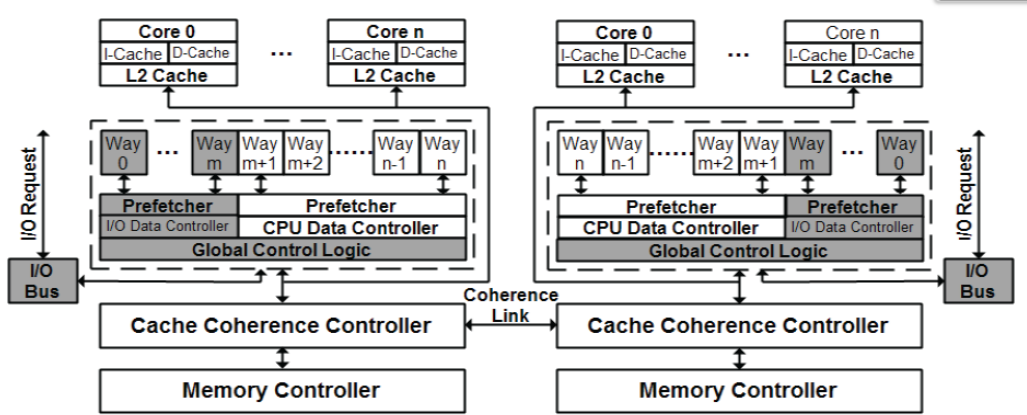


Figure 10 Organization of Partition-Based DMA Cache (PBDC)

Figure 10 that illustrates the organization of PBDC in a multiprocessor system, PBD can use several ways of the processor’s LLC as the dedicated DMA Cache.

Ex. 1004 § 3.2.1. Figure 10 illustrates a multicore processor system, one processor on each side of the figure. Ex. 1007 ¶ 114 (citing annotated Figure 10). Each processor includes multiple cores (Core 0 to Core n), an I/O bus, and a memory controller. Ex. 1004 § 3.2.1. As seen in Figure 10, each processor includes a “Cache Coherence Controller.” *Id.* at Fig. 10. Each processor also includes “Global Control Logic (GCL), I/O Data Controller (IOD-Ctrler), CPU Data Controller (CPUD-Ctrler), two Prefetchers, and Configuration Module (CM).” *Id.* § 3.2.2. The “I/O Data Controller (IOD-Ctrler) and CPU Data Controller (CPUD-Ctrler) are responsible for write policy, maintaining cache coherence, [and] replacement policy” and the LLC controller. *Id.*

3. Xu (Ex. 1005)

Xu is a study and analysis of

optimal core and cache placements for modern Chip Multiprocessors (CMPs). As the number of cores increases, traditional on-chip interconnects such as bus and crossbar suffer from poor scalability and low efficiency. Ring based design has been proposed and implemented to mitigate these problems.

Ex. 1005, 1. Xu discloses that “continuation growth of number of cores will render the ring interconnect infeasible.” *Id.*

4. Harikumar (Ex. 1006)

Harikumar relates to running multiple operating systems on the same hardware and “facilitates handling and isolating events to a partition.” Ex. 1006 ¶ 12.

Figure 4 of Harikumar is reproduced below.

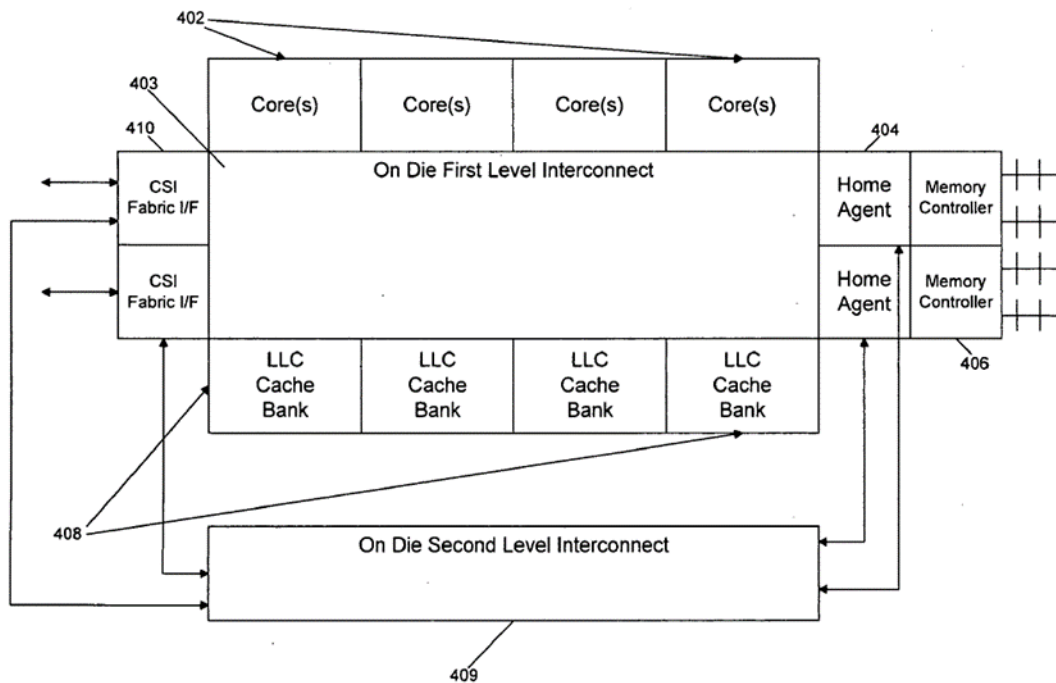


FIG. 4

Figure 4 is a socket architecture of the invention.

Ex. 1006 ¶ 9. As shown in Figure 4, Harikumar teaches a multicore processor with a distributed LLC (Last Level Cache banks). *Id.* ¶ 31. Harikumar discloses the cores and the distributed LLC 408 are connected to each other within the socket by a first level interconnect 403. *Id.* Harikumar discloses “the first level interconnect 403 is an on-die ring interconnect.” *Id.*

II. ANALYSIS OF THE CHALLENGED CLAIMS

A. Level of Ordinary Skill in the Art

Petitioner contends a person of ordinary skill in the art of the '895 patent “would have had a good working knowledge of computer science or electrical engineering and multicore processor systems.” Pet. 8 (citing Ex. 1007 ¶ 30). Such a person “would have gained this through an undergraduate Bachelor’s degree in computer science or electrical engineering, or a comparable field, as well as some experience in the field (*e.g.*, in multicore processor systems).” *Id.* Further,

[t]he more education one has, the less experience needed to attain an ordinary level of skill. Similarly, more experience in the field may serve as a substitute for formal education.

Id.

In the Institution Decision we determined Petitioner’s description was consistent with the prior art before us. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001) (prior art itself may reflect an appropriate level of skill). For the purpose of that decision, we adopted Petitioner’s proposal in the Institution Decision. Inst. Dec. 10–11. Now Patent Owner argues for a “more precise definition” offered in the Horst Declaration and adopted in the Brogioli Declaration. PO Resp. 6–7 (citing Ex. 2002 ¶ 25).

Patent Owner does not argue how this different level of ordinary skill effects any patentability issue. The only difference between the two is the addition of “two years of experience working with multicore processor systems, central processing unit (CPU) , cache, and memory systems” as the experience level of a person of ordinary skill in the art. *Compare* Inst. Dec. 10–11 (citing Pet. 8) *and* Ex. 1007 ¶ 30 *and* Ex. 2001 ¶ 25.

No mention of the difference between the two levels of ordinary skill is made in any of the papers filed. Regardless, inasmuch as both experts agree to the level of ordinary skill, we find the experience associated with the level of ordinary skill is “two years of experience working with multicore processor systems, central processing unit (CPU) , cache, and memory systems.” The level of education is unchanged from the Institution Decision.

B. Claim Construction

We construe claim terms only as relevant to the parties’ contentions and only to the extent necessary to resolve the issues in dispute. *See Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999); *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017).

Petitioner contends that “[u]nless indicated otherwise, all claim terms herein are given such ordinary and customary meaning.” Pet. 8 (citing 37 C.F.R. §42.100(b)). Except as detailed below, Patent Owner also alleges all claim terms “should be construed consistent with their plain and ordinary meaning as understood by a person of ordinary skill in the art when read in the context of the specification and prosecution history.” PO Resp. 5.

1. “*single caching agent*” (claims 1, 11, and 15)

In the Institution Decision, we preliminarily construed a “‘single caching agent’ [to mean] one caching agent for the processor, which may be formed from distributed portions.” Inst. Dec. 12. Patent Owner alleges the term should be interpreted “consistent with its ‘plain and ordinary meaning’ of indicating only *one*, and not, for example, one of many.” PO Resp. 7. (citing Ex. 2002 ¶ 30).

Patent Owner relies on various parts of the description of the ’895 patent to support its proposed plain and ordinary meaning. Patent Owner’s quotes from the ’895 patent are listed below:

“[A] CPU caching agent can be configured to support both CPU traffic and IO traffic as well” Ex. 1001, 1:62–64;

“In various embodiments, *this* caching agent may be *a combined* caching agent both for the CPU as well as for an integrated IO agent.” *Id.* at 2:24–26;

“[V]ia caching agent 200, cores of the processor and an IIO module can maintain coherency *without the need for additional caching agents.*” *Id.* at 1:62–64);

“[A] *single caching agent* per multicore processor socket, where each socket includes multiple cores and an IIO module.” *Id.* at 3:16–19;

“Note that while distributed caching agents are shown, understand that these distributed portions *form a single caching agent*, and which is configured to handle cache coherency operations both for the cores as well as an IIO module 750.” *Id.* at 6:30–39;

“[A] *single caching agent* shared by the IIO module and the cores of the processor.” *Id.* at 7:13–17.

IPR2023-01344
Patent 9,575,895 B2
PO Resp. 7–8.

Patent Owner alleges our preliminary construction, which includes “a plurality of distributed portions each associated with a corresponding one of the [first] plurality of cores,” is “not necessary to account for such a circumstance as part of the construction of ‘single caching agent’ and a person of ordinary skill in the art would appreciate as much.” PO Resp. 8–9 (citing Ex. 1001, 7:58–61 (claim 1), 8:37–39 (claim 11), 9:6–10 (claim 15); Ex. 2002 ¶ 31). Patent Owner contends “[s]uch a construction imparts unnecessary confusion into the meaning of the claim.” *Id.* at 9.

We are not persuaded to change our preliminary construction. First, Patent Owner misstates our construction of “single caching agent” by adding the language “each associated with a corresponding one of the plurality of cores.” That language is present in the claim as a separate limitation. *See* Section I.E (limitation [1g]) above. This added language is not present in our preliminary construction, which states only that a “single caching agent” means “one caching agent for the processor, which may be formed from distributed portions.” Inst. Dec. 12. Thus, that each “single caching agent” is “associated with a corresponding one of the plurality of cores” is a separate limitation.

Second, none of the above citations to the ’895 patent is contrary to our construction. Our construction states that a “single caching agent” is recited. That is what is repeated in Patent Owner’s citations.

Third, the ’895 patent’s description of Figure 6 explains that a “single caching agent” may be formed from distributed portions, as do the claim limitations in question. *See* Ex. 1001, 6:35–37. Other than tacking on the second recitation of “distributed portions” to the “single caching agent”

limitation (PO Resp. 9), Patent Owner does not address why the “distributed portions” part of our construction is incorrect.

The Brogioli Declaration repeats the above citations to the ’895 patent and concludes, without any other supporting evidence, that “each of the independent claims as reciting a solitary, that is one, caching agent for an entire processor (or multicore processor or first multicore processor), comports with the way in which the term is used in the ’895 patent.”

Ex. 2002 ¶ 30. This testimony is conclusory and adds nothing to Patent Owner’s argument to deviate from our preliminary construction.

We therefore make our preliminary construction final. Accordingly, a “single caching agent” means “one caching agent for the processor, which may be formed from distributed portions.”

C. Anticipation/Obviousness

A patent claim is unpatentable for anticipation under 35 U.S.C. § 102 when a prior art reference describes “each and every claim limitation and enable[s] one of skill in the art to practice an embodiment of the claimed invention without undue experimentation.” *ClearValue, Inc. v. Pearl River Polymers, Inc.*, 668 F.3d 1340, 1344 (Fed. Cir. 2012). A prior art reference anticipates a claim only if it discloses all the elements “in the same form and order as in the claim.” *Abbott Labs. v. Sandoz, Inc.*, 544 F.3d 1341, 1345 (Fed. Cir. 2008); *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1369 (Fed. Cir. 2008).

A claim is unpatentable as obvious under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said

subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) where in evidence, so-called “secondary considerations,” including commercial success, long-felt but unsolved needs, failure of others, and unexpected results. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

Petitioner alleges anticipation or obviousness based separately on the prior art to Sinharoy and Tang. *See* Pet. 4. As discussed below, the issue is whether the prior art discloses the “single caching agent” limitation under either legal theory.

D. Claims 1, 3–7, 11–13, and 15–17 as Anticipated or Obvious Over Sinharoy

Petitioner asserts that Sinharoy discloses or teaches or suggests each limitation of claims 1, 3–7, 11–13, and 15–17 and provides a limitation-by-limitation analysis. Pet. 11–43; Ex. 1007 ¶¶ 47–96. Patent Owner disputes that any claim is unpatentable based on Sinharoy. PO Resp. 15–23.

1. Claim 1

a. Preamble 1p: A processor comprising:

Petitioner contends Sinharoy teaches a processor. Pet. 11–12 (citing Ex. 1003, 1:1, Fig. 1). Patent Owner does not dispute this contention. We

determine that on this record Petitioner demonstrates that the preamble of claim 1 is disclosed, taught, or suggested by Sinharoy.⁷

b. Limitation 1a: a plurality of cores,

Petitioner contends that Sinharoy meets this limitation. Pet. 12. Petitioner contends that Sinharoy discloses “the POWER7 processor includes eight cores.” *Id.* (citing Ex. 1003, 1:1, 1:2, Fig. 1 (annotated at Pet. 12); Ex. 1007 ¶ 48).

Patent Owner does not dispute these contentions. We determine that on this record Petitioner demonstrates that limitation 1a of claim 1 is disclosed, taught, or suggested by Sinharoy.

c. Limitation 1b: a shared cache memory,

Petitioner relies on Sinharoy to meet this claim limitation. Pet. 12–13 (citing Ex. 1007 ¶ 49). Petitioner points to Sinharoy’s disclosure that “shared 32-MB L3 cache comprised of the 4-MB local L3 regions from the eight cores” are “shared by the eight cores.” *Id.* at 13 (citing Ex. 1003, 1:17–18; Ex. 1007 ¶ 49).

Patent Owner does not dispute these contentions. We determine that on this record Petitioner demonstrates that limitation 1b of claim 1 is disclosed, taught or suggested by Sinharoy.

d. Limitation 1c: a memory controller to interface with a memory coupled to the processor,

Petitioner asserts that Sinharoy meets this limitation. Pet. 13–14 (citing Ex. 1007 ¶¶ 50–52). Petitioner argues that Sinharoy’s POWER7

⁷ We need not express an opinion on whether the preamble is limiting as Patent Owner does not argue it is not shown.

processor has “two memory controllers—one on each side of the chip.” *Id.* at 13 (citing Ex. 1003, 1:2). Further, according to Petitioner, “[e]ach memory controller interfaces with memory coupled to the processor, as each supports ‘four channels of double-data-rate-three (DDR3) memory’ that ‘provide 100 GB/s of sustained memory bandwidth.’” *Id.* at 13–14. Petitioner concludes that “[b]y providing the memory bandwidth to the POWER7 processor chip and connecting to the processor chip through the off-chip interface and the memory buffer chip, the DDR memory chips are ‘coupled to’ the POWER7 processor.” *Id.* at 14 (citing Ex. 1003, 1:22; Ex. 1007 ¶ 51).

Patent Owner does not dispute these contentions. We determine that on this record Petitioner demonstrates that limitation 1c of claim 1 is disclosed, taught, or suggested by Sinharoy.

e. Limitation 1d: an integrated input/output (I/O) module to interface between the processor and an IO device coupled to the processor and

Petitioner asserts that Sinharoy meets this limitation. Pet. 14–15 (citing Ex. 1007 ¶¶ 53–54). Petitioner alleges “Sinharoy discloses the POWER7 includes integrated input/output modules namely ‘two integrated I/O controllers,’ that interface between the POWER7 processor and ‘IO devices,’ including I/O hub chips connected to PCI bridges and PCI adapters, that are coupled to the processor.” *Id.* at 15 (citing Ex. 1003, 1:22, Figure 14; Ex. 1007 ¶ 53).

Patent Owner does not dispute these contentions. On this record we find that Petitioner demonstrates that limitation 1d of claim 1 is disclosed, taught, or suggested by Sinharoy.

f. Limitation 1e: a caching agent to perform cache coherency operations for the plurality of cores and the IIO module,

Petitioner asserts that Sinharoy meets this limitation. Pet. 16–17 (citing Ex. 1007 ¶¶ 55–58). Petitioner argues that “Sinharoy discloses a ‘caching agent,’ as it discloses centralized controllers for the L2 and L3 caches, including a pool of read/claim (RC) machines and write machines that perform ‘cache coherency operations’ such as fetches, read-write operations on behalf [of] core stores, and core prefetch operations.” *Id.* at 16 (citing Ex. 1003, 1:18–19). In addition, Petitioner alleges Sinharoy’s “ten write machines” handle cache coherency operations. *Id.* (citing Ex. 1003, 1:19, 1:22; Ex. 1007 ¶ 57). Petitioner concludes that “the cache controllers for the L2 and L3 caches including the pool of RC and write machines are a ‘caching agent to perform cache coherency operations for the plurality of cores.’” *Id.* (citing Ex. 1007 ¶¶ 56–58).

Petitioner further argues the same “‘caching agent’ comprising the L2 and L3 cache controllers with the pool of RC and write machines also performs cache coherency operations for the integrated I/O controllers (‘IIO modules’).” Pet. 16 (citing Ex. 1007 ¶ 57). According to Petitioner, “Sinharoy discloses that the POWER7’s integrated I/O controllers ‘support cache injection, which enables DMA write traffic to be written directly into a cache, instead of to memory.’” *Id.* at 16–17 (citing Ex. 1003, 1:22).

Patent Owner does not dispute these contentions. On this record we find that Petitioner demonstrates that limitation 1e of claim 1 is disclosed, taught, or suggested by Sinharoy.

g. Limitation 1f: wherein the processor is to receive an allocation transaction from the IO device and directly store data of the allocation transaction into the shared cache memory,

Petitioner asserts that Sinharoy meets this limitation. Pet. 17–18 (citing Ex. 1007 ¶¶ 59–64). Petitioner asserts Sinharoy’s write requests are data directly stored into the shared cache memory. *Id.* at 17 (citing Ex. 1003 1:17, 1:22; Ex. 1007 ¶ 61). Sinharoy is also relied on for its disclosure that the “shared L2 and L3 caches support cache injection operations, where an ‘I/O device performing a direct memory access (DMA) write operation may target the operation to the cache, instead of to memory.’” *Id.* (citing Ex. 1003, 1:17, 1:22; Ex. 1007 ¶ 60). Petitioner concludes that when the shared L2 and L3 caches support cache injection operations the processor “directly store[s] data of the allocation transaction into the shared cache memory.” *Id.* at 18 (citing Ex. 1007 ¶ 64).

Patent Owner does not dispute these contentions. On this record we find Petitioner demonstrates that limitation 1f of claim 1 is disclosed, taught, or suggested by Sinharoy.

h. Limitation 1g: wherein the caching agent is a single caching agent for the processor and includes a plurality of distributed portions each associated with a corresponding one of the plurality of cores.

As noted previously the primary issue in this case is whether or not Sinharoy or Tang discloses “a single caching agent” as the term appears in limitation 1g. Petitioner asserts that Sinharoy meets limitation 1g, including the “single caching agent.” Pet. 18–20 (citing Ex. 1007 ¶¶ 65–66).

Petitioner cites to its showing for limitation [1e] and Sinharoy’s “caching agent.” *Id.* at 18. Petitioner further argues the “cache controllers for the L2 and L3 caches, including RC and write machines, form a logical single caching controller, with the single cache controller including a plurality of

distributed portions each associated with one of the cores.” *Id.* (citing Pet. 16–17 (showing for limitation 1e); Ex. 1007 ¶ 65).

Petitioner cites the Horst Declaration for its assertion that a “‘single caching agent’ is a necessary single logical component for the processor formed by physically distributed portions.” Pet. 18 n.7 (citing Ex. 1007 ¶ 65). The basis for the opinion is that the ’895 patent discloses the following:

“a caching agent can be distributed such that each of different portions of the caching agent can be associated with a corresponding core and LLC bank or slice.” (’895 Patent, 3:6–10.) The ’895 Patent further discloses, with respect to FIG. 6 thereof, that processor 700 includes a distributed configuration having “partitions or slices each including a core 710 and a partition of a caching agent 715 and a LLC 720.” (*Id.* 6:33–35.) In particular, it discloses that “while distributed caching agents are shown, understand that these distributed portions **form** a single cache agent.” (*Id.* 6:35-37 (emphasis added).).

Ex. 1007 ¶ 65 n.7.

Patent Owner disputes that Petitioner has shown that Sinharoy discloses, teaches, or suggests limitation 1g. PO Resp. 15–23. Patent Owner argues that Petitioner’s assertion that “separate centralized controllers for the L2 and L3 caches described by Sinharoy are distributed portions of a single cache controller” is wrong. *Id.* at 21. Patent Owner argues the centralized controllers for the L2 and L3 caches “are not distributed portions of one cache controller for the entire L2 and L3 cache hierarchy.” *Id.* at 16 (citing Ex. 2002 ¶¶ 36–39). Relying on the Brogioli Declaration, Patent Owner argues instead, Sinharoy teaches an individual controller is provided for each of the L2 and L3 caches. *Id.* (citing Ex. 2002 ¶ 36).

Patent Owner quotes Sinharoy's discussion that

[s]torage accesses that miss both the L2 cache and the 4-MB local L3 region are broadcast to the coherence fabric *and snooped by the memory controller, other L2 caches, possibly other L3 caches, and by the seven remote 4-MB L3 regions that comprise the remaining 28 MB of the on-chip L3 cache.*

PO Resp. 17 (citing Ex. 1004, 19). Patent Owner argues

a person of ordinary skill in the art would understand that when a cache access by a particular core misses that core's private L2 cache and its allocated portion of the shared L3 cache, that miss is broadcast and must be (and is) snooped by the cache controllers for each of the other cores' respective L2 caches and by the respective, separate cache controllers for the other allocated regions of the shared L3 cache.

Id. (citing Ex. 2022 ¶¶ 36–37). Patent Owner closes this argument by asserting that snooping of cache misses by “all of the individual cache controllers of other cores' respective L2 and L3 cache controllers indicates that no one single cache controller is or is acting as *a single*—that is, a solitary—*caching agent for the processor*, as recited by claim 1.” *Id.*

Patent Owner disputes the Horst Declaration testimony that a “‘single caching agent’ is a necessary single logical component for the processor formed by physically distributed portions” because there is “no explanation as to why such a ‘single caching agent’ arrangement is necessary.” PO Resp. 18 (citing Ex. 1007 ¶ 65). Patent Owner asserts that Sinharoy discloses individual controllers for each L2 and L3 cache each operating “independent of one another.” *Id.* (citing Ex. 2002 ¶ 38). That the caches are independent is based in part on the following testimony from the Brogioli Declaration:

In other words, each 4-Mbyte localized L3 region serves as both a victim cache^[8] for its associated L2 cache and as a victim cache for the other seven cores' respective 4-Mbyte L3 regions. [Ex. 1003, 20] To maintain these dual roles, when a particular core's L3 region evicts a cache line, that L3 region first decides whether to route the line to memory or to another cores' L3 region. . . . *This need for such decision making by the targeted L3 region demonstrates that there is no single cache controller operating or controlling the actions of the various caches.* If there were, the need for such decision making by the candidate L3 region would not exist.

Ex. 2002 ¶ 38 (emphasis changed).

Patent Owner contends that “Dr. Horst relies on a statement in Sinharoy that reads, ‘possibly other L3 caches’ may need to snoop L2/L3 misses for his conclusion that there must only be a single logical cache controller.” PO Resp. 21 (citing Ex. 2001, 18:3–9). Patent Owner argues this “ignores the remainder of that same sentence that specifies the controllers of ‘the seven remote 4-MB L3 regions that comprise the remaining 28MB of the on-chip L3 cache’ must perform such snooping.” *Id.* (citing Ex. 1003, 19⁹).

At the final hearing, both parties argued that *if* the L2 and L3 caches are independent of each other, they are not a “single caching agent.” Tr. 14:8–17 (Petitioner), 42:19–23 (Patent Owner). Petitioner argues “Sinharoy

⁸ “The symbiotic relationship between the fully functional L2 RC machines and the lightweight L3 RC machines is possible because of the tight coupling between a given 256-KB L2 cache and its corresponding 4-MB local L3 region, as well as the *partial victim cache management policy that governs their interactions.*” Ex. 1003, 19 (emphasis added).

⁹ “[O]perations that miss the 4-MB local L3 region but hit in the remaining 28 MB of the L3 cache access the cache via the snoop dispatch ports.” Ex. 1003, 19.

discloses the processor receives requests to write data into the shared cache memory (*i.e.*, ‘allocation transactions’) from IO devices (such as the I/O hub chips).” Pet. 17 (citing Ex. 1003, 17, 22; Ex. 1007 ¶¶ 61–63). Conversely, Patent Owner argues Sinharoy discloses individual controllers for each L2 and L3 cache each operating “independent of one another.” PO Resp. 18 (citing Ex. 2002 ¶ 38).

We find that the L2 cache and the L3 cache described in Sinharoy do not act independently but rather, as Sinharoy explains:

As illustrated in Figure 10, the capacity, data flow logic, and control flow logic of a 4-MB local L3 region are shared between the L2 cache to which the local L3 region is coupled, and the other seven 4-MB L3 regions on the chip.

Ex. 1003, 19.

We are not persuaded that the fact that each of the L2 and L3 regions “comprise[s] a single centralized controller” (Ex. 1003, 19) means that the cache regions are independent of each other. Indeed, the Brogioli Declaration, quoted above, discusses shared interactions between the two, including that the “localized L3 region serves as both a victim cache for its *associated L2 cache* and as a victim cache for the other seven cores’ respective 4-Mbyte L3 regions.” Ex. 2002 ¶ 38 (emphasis added).

We also find Petitioner’s argument that the L2 and L3 caches form a “logical single caching agent” persuasive based on the present record. Patent Owner’s responsive argument continues its argument that the two cache regions have separate cache controllers. Sur-Reply 11 (citing Ex. 1003, 18, 19). However, that each cache region has a separate controller does not mean that the cache regions are separate or independent. The claim term is “single caching agent,” not a single cache controller, and the “single

caching agent” “includes a plurality of distributed portions associated with a corresponding one of the plurality of cores” like each of Sinharoy’s local L3 regions that share capacity, data flow logic, and control flow logic with one another and their associated L2 cache for a particular core. Ex. 1003, 19; *see* Ex. 1001, 3:6–10 (a single caching agent can be distributed so each different portion of the caching agent is associated with a corresponding core and last level cache).

That snooping between caches occurs does not alter our determination. The L2 and L3 caches each broadcast a miss to the “coherence fabric” which is “snooped by the memory controller, other L2 caches, possibly other L3 caches, and by the seven remote 4-MB L3 regions that comprise the remaining 28 MB of the on-chip L3 cache.” Ex. 1003, 19. The caches act in concert as a group, i.e., a “single caching agent,” to solve that miss.

On this record, Petitioner demonstrates that limitation 1g of claim 1 is disclosed, taught or suggested by Sinharoy. Petitioner has demonstrated by a preponderance of the evidence that claim 1 was anticipated by or would have been obvious over Sinharoy.

2. *Claims 11 and 15*

Petitioner’s analysis of independent claim 11 incorporates much of its analysis of claim 1. Pet. 20–22. For example, the analysis of the limitations of claim 11 refers to corresponding limitations in claim 1 as follows using Petitioner’s limitation identifiers: 11a (1f); 11b (1e); 11c (1g); and 11d (1c and 1f). Similar correspondence to claim 1 is referenced in the analysis of claim 15. *Id.* at 22–25. Petitioner’s analysis for the limitations of independent claim 15 refers to corresponding limitations in claim 1 as

follows using Petitioner’s limitation identifiers: 15a (1a–e); 15b (1f); 15c (1g); 15d (1a–1e).

Limitation 15e does not have a corresponding limitation in claim 1. Pet. 24–25. Limitation 15e recites “a peripheral controller coupled to at least one of the first and second multicore processors.” *Id.* at 24. Petitioner asserts that Sinharoy meets limitation 15e. *Id.* at 24–25 (citing Ex. 1007 ¶ 79). Petitioner cites specifically to Sinharoy’s Figure 13 to contend that “[t]he first and second POWER7 processors in the system are coupled to a peripheral controller as they are connected to a cluster interconnect chip that incorporates up to three 16 × PCI (peripheral component interconnect) Express** Gen2 I/O subsystem interfaces.” *Id.* at 24.

Patent Owner’s argument for claims 11 and 15 relies on its argument for claim 1. PO Resp. 23. In summary, Patent Owner argues Sinharoy does not disclose “a single caching agent for the processor.” *Id.*

Petitioner has demonstrated by a preponderance of the evidence that claims 11 and 15 were anticipated by or would have been obvious over Sinharoy.

3. *Dependent Claims 3–7, 12, 13, 16, and 17*

Petitioner provides an analysis for each of these dependent claims in relation to Sinharoy. Pet. 25–26 (claim 3), 26–27 (claim 4), 27–28 (claim 5), 28 (claim 6), 29–30 (claim 7), 30–31 (claim 12), 31–32 (claim 13), 32 (claims 16 and 17). Patent Owner does not address these claims, relying on their dependence from the independent claims we found unpatentable. PO Resp. 23. We have reviewed Petitioner’s arguments and evidence.

Petitioner has demonstrated by a preponderance of the evidence that dependent claims 3–7, 12, 13, 16, and 17 were anticipated by or would have been obvious over Sinharoy.

E. Claims 2 and 10 as Obvious Over Sinharoy and Xu

Claims 2 and 10 both depend from claim 1 and recite respectively “wherein the plurality of distributed portions of the caching agent are coupled via a *ring interconnect*” and “a *ring interconnect* to couple the plurality of cores and the shared cache memory via a caching agent.” Ex. 1001, 7:58–61 (claim 2), 8:23–25 (claim 10) (emphasis added). Xu teaches multiprocessor cores using ring based design. *See* Section I.H.3 above.

Petitioner asserts that Sinharoy and Xu teach each limitation of claims 2 and 10 and provides a limitation-by-limitation analysis for each. Pet. 33–36; Ex. 1007 ¶¶ 97–101. Petitioner contends that a person of ordinary skill would have been motivated to combine Xu with Sinharoy because both use the POWER7 processor chip. Pet. 33 (asserting that “Xu, in the same field of art as Sinharoy and in fact describing the same POWER7 chip as Sinharoy, discloses using on-chip ring interconnect” (citing Ex. 1005, TABLE 1)), 34 (asserting that a person of ordinary skill in the art “would have understood, given the descriptions of the on-chip interconnect in Sinharoy, that Sinharoy includes a ring interconnect, and this is confirmed by Xu’s express disclosure that the on-chip interconnect in the POWER7 is a ring interconnect, that a ring interconnect couples the centralized cache controllers for each local L3 regions”).

Patent Owner does not challenge the motivation to combine these references or the showing made. Patent Owner relies on its arguments regarding claims 1, 11, and 15. PO Resp. 23–24. Patent Owner argues Xu is not prior art. *Id.* at 24–27. In Section III below we find Xu is prior art based on Exhibits 1013 and 1014.

Petitioner has demonstrated by a preponderance of the evidence that dependent claims 2 and 10 would have been obvious over the combination of Sinharoy and Xu.

F. Claims 8, 9, and 14 as Obvious Over Sinharoy and Tang

Specific to this ground, Petitioner alleges claims 8, 9, and 14 would have been obvious over Sinharoy and Tang. Pet. 36–43 (citing Ex. 1007 ¶¶ 102–111).

Petitioner argues Tang teaches what Sinharoy does not disclose: “Sinharoy does not explicitly disclose the caching agent sends a ‘drop ownership indication’ to the IIO module in receipt of the snoop request.” Pet. 38 (citing Ex. 1004 § 3.2.2). Petitioner then cites Tang for its teaching of “a multicore-processor system with shared cache with distributed portions.” *Id.* (citing Ex. 1004 § 3.2.2).

Petitioner contends a person of ordinary skill in the art “would have understood that Tang provides details on similar cache states discussed in Sinharoy, which would be common across multi-core processors implementing the cache coherence protocols described in Sinharoy and Tang.” *Id.* at 40–41 (citing Ex. 1007 ¶¶ 105–106). Giving weight to the Horst Declaration, we find the motivation for the combination in the assertion that “[a] person of ordinary skill in the art] would have been motivated to look to the cache-state transitions during cache coherence operations in Tang to understand the cache states listed in Sinharoy.” Ex. 1007 ¶ 106. Patent Owner does not challenge the motivation to combine these references or the showing made. Patent Owner relies on its arguments regarding claims 1, 11, and 15. PO Resp. 23–24.

Petitioner has demonstrated by a preponderance of the evidence that dependent claims 8, 9, and 14 would have been obvious over the combination of Sinharoy and Tang.

G. Claims 1, 5–9, 11–15, 17 as Anticipated or Obvious Over Tang

Petitioner asserts that Tang discloses or teaches or suggests each limitation of claims 1, 5–9, 11–15, 17 and provides a limitation-by-limitation analysis. Pet. 43–70; Ex. 1007 ¶¶ 116–155. Patent Owner disputes that any claim is unpatentable based on Tang. PO Resp. 27–33.

1. Claim 1

a. Limitations 1p, 1a, 1b, 1c, 1d, 1e, 1f.

We have reviewed Petitioner’s arguments and evidence regarding limitations 1p, 1a, 1b, 1c, 1d, 1e, and 1f. Pet. 44–53; Ex. 1007 ¶¶ 112–129. Patent Owner does not challenge the assertions, relying, as it did in the challenge based on Sinharoy, on its argument that Tang does not disclose limitation 1g which recites a “single caching agent for the processor.” PO Resp. 29–32. To the extent limitation 1e is incorporated into Petitioner’s showing for limitation 1g, it is discussed immediately below.

On this record we find that Petitioner demonstrates that these limitations 1p, 1a, 1b, 1c, 1d, 1e, and 1f are disclosed, taught, or suggested by Tang.

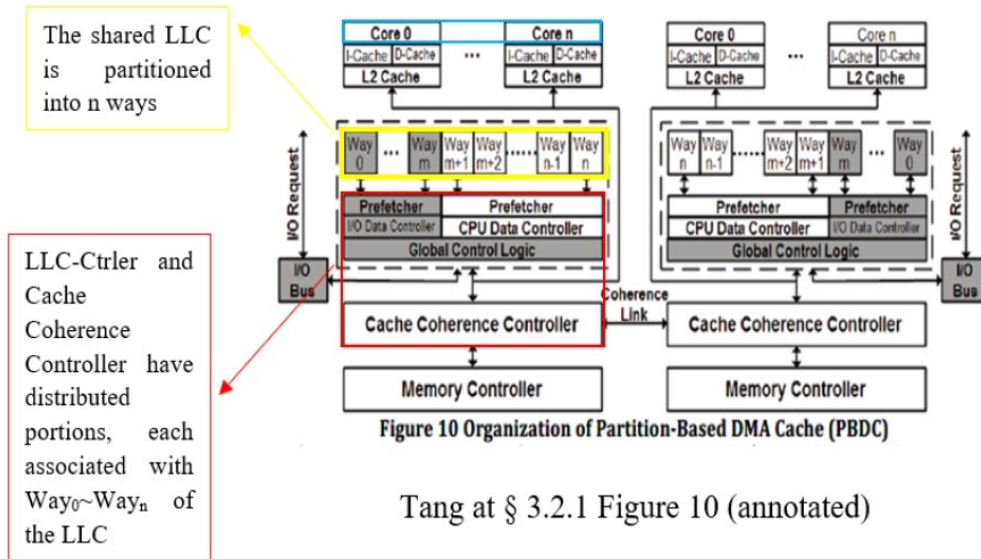
b. Limitation 1g: wherein the caching agent is a single caching agent for the processor and includes a plurality of distributed portions each associated with a corresponding one of the plurality of cores.

Petitioner asserts that Tang meets this limitation. Pet. 53–55 (citing Ex. 1007 ¶¶ 127–129). As to the “caching agent” recitation, Petitioner refers to its showing regarding limitation 1e. *Id.* at 53–54. For limitation 1e, Petitioner argues “Tang, in Figure 10, discloses a Partition-Based DMA Cache (PBDC) system includes a cache coherence controller and a last level cache controller (‘LLC-Ctrler’), which are a ‘caching agent’ and ‘perform cache coherency operations.’” *Id.* at 50 (citing Ex. 1007 ¶ 123; Ex. 1004 § 3.2.2, Fig. 10).

As to the “single caching agent for the processor” recitation in limitation 1g, Petitioner argues Tang’s “cache coherence controller and LLC-Ctrler form a logical single caching agent, and this single caching agent is further formed of a plurality of distributed portions each associated with one of the cores.” Pet. 54 (citing Ex. 1007 ¶ 127). Petitioner also argues that “in the PBDC design, Tang discloses that the LLC is partitioned into ‘several ways’ as the ‘dedicated DMA Cache.’” *Id.* (citing Ex. 1004 §§ 3.2.1, 6).

As to the “plurality of distributed portions . . . with a corresponding one of the plurality of cores” recitation in limitation 1g, Petitioner argues “[b]ecause the LLC-Controller and the cache coherence controller (i.e., the caching agent) manage cache coherence protocol for the plurality of cores (core 0 to core n) and the n-way LLC (Way₀ to Way_n), they have partitioned regions or units (i.e., distributed portions), each associated with a given core.” Pet. 54 (citing Ex. 1007 ¶ 129).

Petitioner cites Tang's Figure 10 (see Section I.H.2 above) as including two processors but shows it is relying on one of the two for its showing. Annotated Figure 10 of Tang is reproduced below.



(Tang, Fig. 10 (annotated), §3.2.1.)

Annotated Figure 10 illustrating the organization of PBDC in a multiprocessor system, PBDC can use several ways of the processor's LLC as the dedicated DMA Cache.

Pet. 54 (citing Tang § 3.2.1, annotated Fig. 10). Petitioner clarifies its position on Figure 10 in the Reply, stating “the Petition has clearly laid out that the cache controller and the LLC controller of, for example, just the processor on the left of Figure 10 form a ‘single caching agent’ for that processor.” Reply 16 (citing Pet. 51); *see also* Tr. 28:5–8 (“Now, the first processor on the left, this processor here, is performing the cache coherence for what we saw in the claim. It’s performing cache coherence operations for the cores, and it’s performing cache coherence operations for an IIO module.”).

Petitioner notes that the data ways Way_0 to Way_n shown above in annotated Figure 10 “are partitioned LLC, and the caching agent therefore has distributed portions to handle the cache coherence for the n cores and the n -way LLC of the processor.” Pet. 54 (citing Ex. 1004, Fig. 10; Ex. 1007 ¶ 129).

Patent Owner disputes that Petitioner has shown that Tang discloses, teaches, or suggests limitation 1g. PO Resp. 29–32. First, Patent Owner specifically argues Petitioner has failed to show that Tang’s “cache coherence controller and the LLC controller are collectively a ‘single caching agent for the processor’ as recited in the challenged independent claims.” PO Resp. 29 (citing Pet. 53–54). Patent Owner argues Tang does not teach such in that “Tang shows a pair of cache coherence controllers which communicate with one another via a ‘coherence link.’” *Id.* (citing Ex. 1004, Fig. 10 (reproduced in Section I.H.2 above)). Patent Owner argues that if “the two coherence controllers were part of the same ‘single caching agent for the processor’ no coherence link would be needed as the two units would be part of the same whole.” *Id.* (citing Ex. 2002 ¶ 44).

Patent Owner contends that the argument made in the Petition and the Horst Declaration that these two coherence controllers are part of the same “logical” unit is unavailing because “Tang makes no such claim.” PO Resp. 29. Patent Owner notes that the Horst Declaration asserts the controllers are logically the same because “there is a need to maintain cache coherence throughout the processor.” *Id.* at 29–30 (citing Ex. 1007 ¶ 127). Patent Owner argues that Tang’s Figure 10 shows two separate processors and that “Tang would simply have drawn the two units as a single one, or at least placed them in a common dotted outline so as to show” a single “logical” unit. *Id.* at 30 (citing Ex. 2002 ¶ 44).

Patent Owner next argues that “[e]ven [if] Tang can be read as disclosing a single caching agent for the processor, it is evident that that caching agent does not include a plurality of distributed portions *each associated with a corresponding one of the plurality of cores*, as further required by the challenged claims.” PO Resp. 30 (citing Ex. 2002 ¶ 45). Patent Owner again references Tang’s Figure 10 as depicting that “all of the ways of this last level cache are shared across all of the cores of the processor.” *Id.* at 31 (citing Ex. 2002 ¶¶ 45–46, Ex. 1004, 7). Patent Owner argues that Tang explains that on “a particular core’s L2 cache miss, the request is provided to the CPU data controller to see if the requested data is present in any of the CPU data ways.” *Id.* According to Patent Owner, “[t]his highlights the shared nature of the LLC cache ways across all of the processor’s cores.” *Id.* (citing Ex. 2002 ¶¶ 45–46).

Patent Owner further argues that the claimed “single caching agent includes a plurality of distributed portions each associated with a corresponding one of the plurality of cores.” PO Resp. 31 (emphasis omitted). Patent Owner argues that “in Tang, the different ways of the LLC cache are shared across the cores of the processor and so it is not true that the cache controller has partitioned regions or units each associated with a given core.” *Id.* at 31–32 (citing Ex. 2002 ¶ 46). Patent Owner contends “[t]he partitioned units may be associated with different ways of the last level cache, but those ways are not each associated with a given core and so neither are their associated cache controllers.” *Id.* at 32.

For the following reasons we determine Tang meets limitation 1g. Referencing the Horst Declaration, Patent Owner’s oral argument repeated its contention that the

last level cache in Tang is shared by all the cores of the processors, and so, too, are all of the components of the controller, the CPU data controller, the pre-fetcher, the global control logic, the cache coherence controller, all of those are also *shared* by all of the cores of the processor.

Tr. 54:6–14 (emphasis added); *see also* PO Resp. 31–32 (“the different ways of the LLC cache are *shared* across the cores of the processor”) (citing Ex. 2002 ¶ 46) (emphasis added). That the ways (Way₀ to Way_n) are shared as between processors supports Petitioner’s position that Tang’s two linked processors form “a single logical caching agent.” Furthermore, we agree with the position of Petitioner that one of the processors of Tang’s Figure 10, e.g., the left processor, also discloses “a single caching agent.”

That the ways are “shared” between physical processors does not preclude them being “distributed portions each associated with a corresponding one of the plurality of cores.” Figure 10 of Tang shows the ways associated with a core, and distributed logic of the cache coherence controller is associated with each core as indicated by arrows from cache coherence controller to each Core 0 to Core n. *See* Ex. 1004, 7 (the ways include I/O data ways and CPU data ways (Way_{m+1} – Way_n) for Core 0 to Core n); Reply 18–19 (annotating Figure 10 to depict distributed portions of the logic associated with a corresponding core).

We are not persuaded differently by the Brogioli Declaration, which opines that sharing between cores is the antithesis of an association between the ways and one of the cores. Ex. 2002 ¶ 46. Sharing establishes that there is a single logical caching agent established by the shared ways.

On this record we find that Petitioner demonstrates that limitation 1g is disclosed, taught, or suggested by Tang.

2. *Claims 5–9, 11–15, 17*

Independent claims 11 and 15 contain the same limitation 1g language addressed above. Ex. 1001, 8:38–41 (claim 11), 9:7–10 (claim 15). Petitioner relies on its showing based on Tang for limitation 1g in claim 1 for the corresponding limitations of claims 11 and 15. See Pet. 57 (limitation 11c), 59 (limitation 15c). Patent Owner’s argument is the same for claims 11 and 15 as for claim 1, addressed above. PO Resp. 29–32.

Dependent claims 5–9, 12–14, and 17 all depend from claims 1, 11 or 15. Patent Owner argues patentability of these claims based on its arguments for the independent claims. PO Resp. 33.¹⁰ No separate arguments are made. *Id.*

3. *Summary*

Petitioner has demonstrated by a preponderance of the evidence that claims 1, 5–9, 11–15, and 17 were anticipated by or would have been obvious over Tang.

H. Claims 2, 3, 4, 10, and 16 as Obvious Over Tang and Harikumar

Specific to this ground, Petitioner alleges dependent claims 2–4, 10, and 16 would have been obvious over Tang and Harikumar. Pet. 70–77. We are persuaded that a person of ordinary skill in the art would have combined Tang and Harikumar because both “disclose[] a multicore processor with a distributed LLC.” Pet. 72–73 (citing 1006 ¶ 31; Ex. 1007 ¶ 157). Dr. Horst testifies that a

[person of ordinary skill in the art] would have recognized that using the ring interconnect disclosed in Harikumar in the

¹⁰ Patent Owner does not list claim 12.

processor of Tang would be the expected use of such well-known interconnect. . . [and] would have been motivated to use the ring interconnect disclosed in Harikumar in the processor of Tang to achieve the same purpose to connect the cores, the shared cache memory, and the caching agent.

Ex. 1007 ¶ 158.

For dependent claims 2–4, 10, and 16, Patent Owner argues Petitioner relies on “Tang for disclosing a *single* caching agent that includes a plurality of distributed portions *each associated with a corresponding one of the plurality of cores.*” PO Resp. 33. Thus, Patent Owner contends these claims are patentable for the same reasons the independent claims are patentable. *Id.*

Petitioner has demonstrated by a preponderance of the evidence that dependent claims 2, 3, 4, 10, and 16 would have been obvious over the combination of Tang and Harikumar.

III. Patent Owner’s Motion to Exclude

Patent Owner moves to exclude Exhibit 1019, Copyright Office Record of “2011 NORCHIP” (11 EX5488). *See* Mot. 1. Exhibit 1019 and Exhibit 1013 (“MacPherson Declaration” re Xu) are relied on by Petitioner to show that Xu is prior art to the claims of the ’895 patent. Pet. 33 n.10.¹¹ The MacPherson Declaration is relied on as evidence that copies of Xu were made available no later than the “last date of the conference,” which was November 15, 2011. Ex. 1013. Petitioner alleges Xu is also available for

¹¹ ““Gordon MacPherson confirmed that copies of Xu were made available no later than the ‘last date of the conference [November 15, 2011].’ (EX1013.) Xu is also available for public download from the IEEE Xplore. (See <https://ieeexplore.ieee.org/document/6126728> (EX1014).)” Pet. 33 n.10.

IPR2023-01344
Patent 9,575,895 B2

public download from the IEEE Xplore. Pet. 33 n.10 (citing <https://ieeexplore.ieee.org/document/6126728> (Ex. 1014,¹² “IEE Explore”)).

Patent Owner objects to Exhibit 1019 as hearsay and as lacking authentication. Paper 14 (“Patent Owner’s Objections to Petitioner’s Evidence”). In its follow-on Motion to Exclude, Patent Owner argues that, because

there are no identifying criteria showing a website from which this page was obtained or when, and no accompanying declaration attesting to its origin or the circumstances of its production[,] it cannot be determined with any certainty when, how, or from where this exhibit was produced. The exhibit itself purports to state that a work entitled “2011 NORCHIP (11EX5488),” which is said to be a “book” having 221 pages, was published “2011-12-22.”

Mot. 2 (citing Ex. 1013). Patent Owner adds that Petitioner has not shown circumstances that guarantee its trustworthiness “because the nature of the evidence precludes any connection between the Xu reference itself and the copyright record which Ex. 1019 purports to be.” *Id.* at 3.

Petitioner argues Exhibit 1019 is admissible under the public records exception to hearsay. Opp. 1 (citing Fed. R. Evid. 803(8)). Patent Owner argues that even if Exhibit 1019 is not hearsay, it still has not been authenticated. Opp. Reply 1 (citing McCormick on Evidence § 227 at 74 (6th ed. 2006)).

Ultimately, Patent Owner argues “[o]ur contention is [Petitioner] ha[s]n’t shown Xu is prior art because they haven’t properly indicated the date of its public availability.” Tr. 56:23–25. In sum, Patent Owner does

¹² *Explorations of optimal core and cache placements for Chip Multiprocessor*, IEEE Explore, IEEE (Date of Conference November 14–15, 2011, Date Added to IEEE Explore, January 12, 2012).

IPR2023-01344
Patent 9,575,895 B2

not contest that Xu was publicly available, but rather the date Xu was publicly available.

“In an IPR, the petitioner bears the burden of establishing by a preponderance of the evidence that a particular document is a printed publication.” *Nobel Biocare Services AG v. Instradent USA, Inc.*, 903 F.3d 1365, 1375 (Fed. Cir. 2018) (citing *Medtronic, Inc. v. Barry*, 891 F.3d 1368, 1380 (Fed. Cir. 2018)). “A reference will be 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.” *Medtronic*, 891 F.3d at 1380 (citations, quotations, and alterations omitted).

The evidence of record supports our finding that Xu was publicly available at an IEEE conference occurring from November 14 through 15, 2011. Ex. 1013 ¶¶ 8–11, Ex. A (Xu). Xu has a copyright date of 2011 which, although not dispositive of the date of public accessibility, is relevant evidence of public accessibility at the conference. Ex. 1005, 1. Exhibit 1014 confirms the conference dates of November 14–15, 2011, which antedates the December 31, 2011, critical date that Patent Owner asserts for the ’895 patent (PO Resp. 26; *see* Reply 13–14) and the December 13, 2011, priority date of the ’895 patent (Ex. 1001, at code [63], 1:5–7). Ex. 1014.

Whether or not Exhibit 1019 is hearsay or authenticated, the MacPherson Declaration and IEEE Explore evidence is sufficient to show that Xu was available at the November 14–15, 2011 conference. The Motion to exclude Exhibit 1019 is *denied as moot*.

IV. CONCLUSION

For the reasons discussed above, Petitioner has shown by a preponderance of the evidence that claims 1–17 are unpatentable.¹³

V. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that claims 1–17 of the '895 patent are determined to be unpatentable;

FURTHER ORDERED that Patent Owner's Motion to Exclude (Paper 24) is dismissed as moot; and

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

In summary:

Claim(s)	35 U.S.C. §	Reference(s)/Basis	Claim(s) Shown Unpatentable	Claim(s) Not shown Unpatentable
1, 3–7, 11–13, 15–17	102	Sinharoy	1, 3–7, 11–13, 15–17	
1, 3–7, 11–13, 15–17	103	Sinharoy	1, 3–7, 11–13, 15–17	

¹³ Should Patent Owner wish to pursue amendment of the challenged claims in a reissue or reexamination proceeding subsequent to the issuance of this decision, we draw Patent Owner's attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. See 84 Fed. Reg. 16,654 (Apr. 22, 2019). If Patent Owner chooses to file a reissue application or a request for reexamination of the challenged patent, we remind Patent Owner of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. See 37 C.F.R. § 42.8(a)(3), (b)(2).

IPR2023-01344
Patent 9,575,895 B2

2, 10	103	Sinharoy, Xu	2, 10	
1, 5-9, 11-15, 17	102	Tang	1, 5-9, 11-15, 17	
1, 5-9, 11-15, 17	103	Tang	1, 5-9, 11-15, 17	
8, 9, 14	103	Sinharoy, Tang	8, 9, 14	
2, 3, 4, 10, 16	103	Tang, Harikumar	2, 3, 4, 10, 16	
Overall Outcome			1-17	

IPR2023-01344
Patent 9,575,895 B2
FOR PETITIONER:

Celine Crowson
Joe Raffetto
Scott Hughes
Nicholas Rotz
Ryan Stephenson
Helen Trac
HOGAN LOVELLS US LLP
celine.crowson@hoganlovells.com
joseph.raffetto@hoganlovells.com
scott.hughes@hoganlovells.com
nicholas.rotz@hoganlovells.com
ryan.stephenson@hoganlovells.com
helen.trac@hoganlovells.com

FOR PATENT OWNER:

Tarek Fahmi
Jonathan Tsao
ASCENDA LAW GROUP, PC
tarek.fahmi@ascendalaw.com
jonathan.tsao@ascendalaw.com