

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MERCEDES-BENZ USA, LLC,

Petitioner,

v.

DAEDALUS PRIME LLC,

Patent Owner.

Case No.: IPR2023-01343

U.S. Patent No. 8,898,494

PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

I.	INTRODUCTION AND SUMMARY OF ARGUMENTS..	1
II.	BACKGROUND, THE '494 PATENT, AND CLAIM CONSTRUCTION	4
A.	Background	4
B.	The '494 Patent	5
C.	The Challenged Claims	11
D.	Claim Construction and the Level of Ordinary Skill in the Art	13
	1. Workload	15
E.	The IPR Petition.....	16
III.	THE PETITION FAILS TO DEMONSTRATE THAT ANY CHALLENGED CLAIM IS UNPATENTABLE	17
A.	Bose Fails to Teach or Suggest the Invention Recited in Challenged Independent Claim 1	19
	1. Overview of Bose	19
	2. Bose does not anticipate claim 1	23
	3. Claim 1 is not obvious in view of Bose.....	28
B.	The Remaining Challenges Based on Bose Also Fail.....	33
C.	The Combination of White and Simeral Fails to Teach or Suggest the Invention Recited in Challenged Independent Claim 1	34
	1. Overview of White	35
	2. Overview of Simeral.....	38

3.	The combination of White and Simeral proposed by Petitioner fails to teach or suggest a core workload monitor configured to determine a core workload	42
4.	Claims 3, 14, and 15 are also patentable over White and Simeral.....	45
5.	The remaining challenges based on White and Simeral also fail	45
D.	White Fails to Teach or Suggest the Invention Recited in Challenged Independent Claim 4	46
1.	Claim 7 is also patentable over White	52
2.	The remaining challenges based on White also fail	52
IV.	CONCLUSION	53

TABLE OF AUTHORITIES

CASES

<i>ACTV, Inc. v. Walt Disney Co.</i> , 346 F.3d 1082 (Fed. Cir. 2003).....	14
<i>Alza Corp. v. Mylan Labs., Inc.</i> , 464 F.3d 1286 (Fed. Cir. 2006).....	29
<i>CFMT, Inc. v. Yieldup Int’l. Corp.</i> , 349 F.3d 1333 (Fed. Cir. 2003).....	33, 34, 45, 52
<i>Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.</i> , 800 F.3d 1375 (Fed. Cir. 2015).....	53
<i>Harmonic Inc. v. Avid Tech., Inc.</i> , 815 F.3d 1356 (Fed. Cir. 2016).....	53
<i>Hartness Int’l. Inc. v. Simplimatic Engineering Co.</i> , 819 F.2d 1100 (Fed. Cir. 1987).....	passim
<i>In re Fine</i> , 837 F.2d 1071 (Fed. Cir. 1988).....	33
<i>In re Gordon</i> , 733 F.2d 900 (Fed. Cir. 1984).....	18
<i>In re Magnum Oil Tools Int’l, Ltd.</i> , 829 F.3d 1364 (Fed. Cir. 2016).....	53
<i>In re Peterson</i> , 315 F.3d 1325 (Fed. Cir. 2003).....	18
<i>In re Ratti</i> , 270 F.2d 810 (CCPA 1959)	18
<i>In re Warner</i> , 379 F.2d 1011 (CCPA 1967)	17

<i>Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.</i> , 821 F.3d 1359 (Fed. Cir. 2016).....	18
<i>KSR Int’l v. Teleflex Inc.</i> , 550 U.S. 398 (2007).....	17, 19
<i>Net MoneyIN, Inc. v. VeriSign, Inc.</i> , 545 F.3d 1359 (Fed. Cir. 2008).....	26
<i>Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co. Ltd.</i> , 868 F.3d 1013 (Fed. Cir. 2017).....	14
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	14
<i>PPG Industries v. Guardian Industries Corp.</i> , 75 F.3d 1558 (Fed. Cir. 1996).....	28
<i>SAS Inst. Inc. v. Iancu</i> , 138 S. Ct. 1348 (2018).....	17
<i>Trintec Industries, Inc. v. TOP-USA Corp.</i> , 295 F.3d 1292 (Fed. Cir. 2002).....	26
<i>Vitronics Corp. v. Conceptronic, Inc.</i> , 90 F.3d 1576 (Fed. Cir. 1996).....	14
<i>Zenith Elecs. Corp. v. PDI Commc’ns Sys., Inc.</i> , 522 F.3d 1348 (Fed. Cir. 2008).....	28

OTHER AUTHORITIES

<i>Changes to the Claim Construction Standard for Interpreting Claims in Trial Proceedings Before the Patent Trial and Appeal Board</i> , 83 Fed. Reg. 51,340 (Oct. 11, 2018)	13
Patent Trial and Appeal Board Consolidated Trial Practice Guide, Nov. 2019	26

EXHIBIT LIST

EXHIBIT NO.	DESCRIPTION
2001	Declaration of Michael C. Brogioli, Ph.D.
2002	Transcript of Deposition of David Wyatt

I. INTRODUCTION AND SUMMARY OF ARGUMENTS

Mercedes-Benz USA, LLC's ("MB USA" or "Petitioner") challenges to the claims of U.S. Patent No. 8,898,494 (the "'494 patent") fail to demonstrate that any challenged claim is unpatentable.

As to MB USA's challenges based on Bose as a primary reference, those challenges fail because Bose does not teach or suggest a "balancing control adapted to receive the bus workload from the bus workload monitor and to dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit and a comparison between the bus workload and a bus workload threshold," as recited in claim 1. The "phase prediction technique" described by Bose involves various power vs. performance tradeoffs as reflected in specified power modes that are assigned by a global power management control unit (PMCU). Contrary to the allegations in the petition, the power vs. performance tradeoffs, which are characterized in a power-performance tradeoff table, do ***not*** represent a workload *threshold* for a given component working in an associated power mode. Indeed, inasmuch as the power/performance values change over time, a person of ordinary skill in the art would recognize these values are algorithmic in nature.

Moreover, Bose fails to teach or suggest any comparison with a

workload threshold as alleged in the petition. In Bose there is no need for a comparison of bus workload to a bus workload threshold because direct information concerning actual power level decreases or increases for different power modes is available from the power-performance tradeoff table. Hence, the global PMCU can directly determine whether a power level associated with a prospective power mode will “fit” within a total power budget.

As to the challenges based on White and Simeral, the petition relies on White for allegedly teaching power management logic that monitors a core workload using a core workload monitor. This is not so. Instead, White’s power management logic makes decisions regarding future operating characteristics (operating voltage, operating frequency) of a core based on its present operating characteristics. This is not the same as monitoring workload using a core workload monitor, as recited in claim 1. In fact, the ’494 patent itself explains that performance metrics such as frequency can be altered independently of workload, and so a person of ordinary skill in the art would appreciate that merely monitoring operating voltage and operating frequency does not necessarily entail determining activity as monitoring workload does.

Finally, as to the challenges based on White, these too fail because

White does not teach or suggest receiving a bus workload and a processing element workload, as recited in independent claim 4. White is clear that what is “received” by the power management logic is not workload – i.e., is not activity levels. Instead, what is received is current operating voltages and operating frequencies, and the petition admits as much. As the petition fails to allege any modifications to White that would result in a system that *receives* a bus workload and a processing element workload, the petition fails to make out a prima facie case of obviousness. Moreover, when attempting to allege that White discloses receiving a bus workload and a processing element workload based on an integrated circuit reaching a limit, the petition confuses the actions of White’s power management unit and temperature control unit and, in doing so, fails to make a case for obviousness of claim 4.

For these and further reasons as discussed in this Patent Owner Response, the Board should not find any of the challenged claims unpatentable.

II. BACKGROUND, THE '494 PATENT, AND CLAIM CONSTRUCTION

A. Background

Modern computing devices can make use of multi-core processors that combine a plurality of processor cores on a single integrated circuit. *See* Ex. 1001 at 1:20-26. The different processor cores are the functional units that read and execute instructions contained in a computer program. The presence of multiple cores on a single integrated circuit means a multi-core processor is capable of maintaining several independent architectural states. Ex. 2001 at ¶ 28.

The consolidation of multiple processing cores and other components on a single integrated circuit did not, however, mean that power consumption and performance of the different cores and components were always considered or addressed at a system level. Ex. 1001 at 1:47-52. Instead, individual devices were often associated with their own, separate power saving policies. *Id.* And, even where system-level policies were defined, they often did not consider the need to balance power consumption and performance of the different system elements to provide an optimal system-level solution. *Id.* at 1:52-57.

B. The '494 Patent

The '494 patent proposes improved power management techniques in integrated circuits, such as computer processors, by balancing power allocation between processing cores and other elements, including other processing elements and communication buses. *See, e.g., id.* at 3:15-30; 7:40-61; 10:6-51; 14:59-64; 15:4-21; and 16:10-41, *et seq.* Whereas previously, power management policies were implemented separately amongst system-on-chip components (e.g., each component received a power budget and implemented its own power savings techniques), or merely in a uniform blanket fashion across the integrated circuit (e.g., components were uniformly capped), the inventors of the '494 Patent devised of an integrated circuit that improved system performance by balancing allocation between processing components such as a communication bus and a processor core, based on a power limit and a workload threshold comparison. *See id.* at 16:10-41, *et seq.* Using the systems and methods described in the '494 Patent, processors are able to increase overall performance without increasing power consumption.

For example, according to the methods described in the '494 patent, if the workload of one component in a processor system is too high, then that component's frequency/power may be increased. *Id.* at 17:8–11. In this

regard, a component's activity may be taken as being indicative of its workload:

Determining a workload includes any known method for determining activity of hardware, a workload of software, a number of instructions executed or to execute, or any other metric for determining use of a processing element. As one example, a workload for core 505 includes an amount time core 505 is active during a period of time. Activity may be determined as when core 505 is in an active performance state versus a sleep or low power state or a number of cycles executing or performing operations.

Id. at 9:47-57. In the case of a communication bus, activity may be measured by the number of cycles or number of work slots occupied on the bus. *Id.* at 16:31-34. Both processing elements and communications busses may be instrumented (e.g., using workload monitors) to determine the various measures of component activity. *Id.* at 16:10-41.

The inventors of the '494 patent determined that—contrary to prior power management approaches, which would cap or cut power supplied to high-activity processor elements (because those are the elements consuming high power)—increasing power supplied to high-activity elements improves overall system functionality by reducing bottlenecking. However, because

the amount of power available within an integrated circuit is limited, the increase in power supplied (e.g., by increasing frequency) to high-activity components must come from a limited power budget. The inventors of the '494 patent therefore determined that power allocation must be tuned between various components:

In one embodiment, power control module 160, which may include hardware, firmware, software, or a combination thereof, is able to balance and allocate power/frequency between devices of IC 100 to achieve maximum performance; even when IC 100 is under a power limit. Additionally, such allocation, in one scenario is based on actual workload of devices, instead of assumptions of workload or power consumption. For example, a power limit (either a maximum power limit or a reduced limit during power saving modes) is provided for IC 100. Previously, a maximum current (and as a corollary frequency) was reduced for all processing elements of IC 100 to meet the power limit. Instead, in this scenario, a device that is determined to be a bottleneck for performance, such as throughput performance as viewed from the perspective of an application, is allocated more current/frequency, while the other competing device is capped/limited to ensure the power limit is still met. In other words, a dynamic balance is continuously [struck] between devices to

ensure the device that needs more resources for better overall performance is allocated such extra resources, while other devices are ‘balanced’ (reduced in performance) to meet a given power limit.

Id. at 7:40-61.

In the case of balancing power (e.g., current, voltage, frequency, or other metric affecting power consumption) between processing cores and a communication bus, *see e.g., id.* at 15:11–21; 16:10–14, the ’494 patent indicates that when the communication bus has high workload, increasing the communication bus frequency may need to be “counterbalanced in the CPU.” *Id.* at 17:11–15. “[A]ny number of balancing effects may be utilized,” such as by decreasing the frequency for the CPU cores, but the ’494 patent expresses a preference for an “intelligent balancing methodology” in which “caps may be dynamically and separately altered for each component or unit.” *Id.* at 17:16–22. “Therefore, a power increase for communication bus 950 ... is dynamically balanced with a cap of CPU cores 905 ... to alleviate the bottlenecks in the system to increase performance, while not exceeding the current limit of [the integrated circuit] 900.” *Id.* at 17:35–41.

Figure 9 illustrates one embodiment of a system configured in accordance with the described approach for power management:




FIG. 9

Id. at Figure 9. Integrated circuit 900 includes multiple cores (e.g., CPU cores 905, GPU cores 910), cache or memory 960, communication bus 950, and power management service 920. *See e.g., id.* at Fig. 9; 2:21-24; 15:48-52; 16:10-41, 57-59; 17:4-15. Communication bus 950 connects the cores 905, 910, to shared resources such as cache or memory 960, and each component has an associated workload monitor (911, 935, 941). *Id.* at Fig. 9; 9:47-52; 15:38-67; 16:14-56. Power management service 920 balances power among the components based on their workloads. *Id.* at Fig. 9; 16:57-66. For example, the communication bus's frequency/power may be

increased if its workload is too high, alleviating a bottleneck between the cores and cache, and that increase may be counterbalanced by decreasing the frequency/power of the cores. *Id.* at 17:4-41; and see *id.* at Fig. 10




FIG. 10

reproduced here, showing an example of a process for balancing power between cores and a communication bus within a power limit); 18:35 – 19:46. In other examples, the core’s power/frequency may be increased and be counterbalanced by a decrease in the power/frequency of another element

to maintain operation of the integrated circuit below its thermal limit. *Id.* at 19:3-63.

C. The Challenged Claims

In keeping with the above-described teachings, independent claim 1 of the '494 patent recites:

1. A processor comprising:
 - an integrated circuit including:
 - a first core;
 - a cache memory;
 - a communication bus coupled to the first core, the communication bus to connect the first core and the cache memory;
 - a core workload monitor configured to determine a core workload for the first core;
 - a bus workload monitor configured to determine a bus workload for the communication bus; and
 - balancing control adapted to receive the bus workload from the bus workload monitor and to dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit and a comparison between the bus workload and a bus workload threshold, the power limit corresponding to a maximum thermal dissipation capacity for the

integrated circuit, wherein a power consumption of one of the first core and the communication bus is to be reduced, the power consumption reduction to be limited to maintain operation of the one of the first core and the communication bus above a low limit.

Id. at 21:17-39.

Independent claim 4 recites:

4. A non-transitory storage medium including code, when executed, to cause a machine to perform the energy efficiency operations of:

receiving a bus workload for a communication bus and a first processing element workload for a first processing element based on an integrated circuit reaching a limit, the integrated circuit including the communication bus and the first processing element, wherein the communication bus is to connect the first processing element with at least one other device of the integrated circuit;

determining if the first processing element workload indicates the first processing element is a performance bottleneck;

determining if the communication bus is able to accommodate the bus workload;

capping a bus frequency of at least the communication bus and a frequency of a second processing element in response to determining the

first processing element workload indicates the first processing element is the performance bottleneck and determining the communication bus is able to accommodate the bus workload; and

increasing a previous cap on the bus frequency of at least the communication bus in response to determining the first processing element workload indicates the first processing element is the performance bottleneck and determining the communication bus is not able to accommodate the bus workload.

Id. at 21:60 – 22:20.

D. Claim Construction and the Level of Ordinary Skill in the Art

In *inter partes* reviews such as this one, the Board applies the same claim construction standard used to construe a claim in a civil action under 35 U.S.C. § 282(b). *Changes to the Claim Construction Standard for Interpreting Claims in Trial Proceedings Before the Patent Trial and Appeal Board*, 83 Fed. Reg. 51,340 (Oct. 11, 2018). Accordingly, the claims should be construed consistent with their plain and ordinary meaning as understood by a person of ordinary skill in the art as considered in the context in which the terms of the claims are used. *Phillips v. AWH Corp.*, 415 F.3d 1303,

1312-17 (Fed. Cir. 2005); *Vitronics Corp. v. Conceptronics, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996); *see also ACTV, Inc. v. Walt Disney Co.*, 346 F.3d 1082, 1088 (Fed. Cir. 2003) (“the context of the surrounding words of the claim also must be considered in determining the ordinary and customary meaning of those terms”). “The ordinary and customary meaning of a claim term is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention.” *Phillips*, 415 F.3d at 1313. Furthermore, constructions are needed only to the extent necessary to resolve any asserted ground in the petition. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co. Ltd.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (only “terms ‘that are in controversy, and only to the extent necessary to resolve the controversy’” need be construed) (quoting *Vivid Techs., Inc. c. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999)).

In its Institution Decision, the Board assessed MB USA’s proposed level of ordinary skill in the art as being “consistent with the prior art.” Paper 7 at 13. That proposed level of ordinary skill is: a person with “an undergraduate degree in electrical engineering, computer engineering, computer science, or similar disciplines” who would have had “two years of professional experience working with the research, design, and/or development of semiconductors, processors, power consumption and

thermal power management, and related firmware and software or an equivalent level of skill, knowledge, and experience.” *Id.* at 12-13. Patent Owner uses this definition in the analysis below. *See* Ex. 2001 at ¶ 25.

Although the Board did not make any claim constructions as part of its Institution decision, *see* Paper 7 at 13, certain terms of the claims will be important for understanding distinctions from the prior art cited by Petitioner:

1. Workload

All of the challenged independent claims recite the term “workload,” whether as a “workload monitor,” a “core workload,” a “processing element workload,” and/or a “bus workload.” *See* Ex. 1001 at 21:17-39 (claim 1); 21:60 – 22:20 (claim 4). A person of ordinary skill in the art would understand that in the ’494 patent, the term “workload” means “an amount of activity over a quantum or period of time.” Ex. 2001 at ¶¶ 36-38; *and see* Ex. 1001 at 8:20-22; 9:49-57 (recognizing workload as “activity of hardware,” “a number of instructions executed or to execute,” “use of a processing element,” and “an amount [of] time core 505 is active during a period of time”); 9:61-63 (“workload is measured as a number of active cycles over a period of time (i.e., a total number of cycles for

measurement)"); Ex. 2002 at 34:14 – 36:9 (Petitioner’s declarant, Mr.

Wyatt, acknowledging that “workload” would be understood as activity over time).

E. The IPR Petition

In its petition, MB USA presents multiple grounds for the alleged unpatentability of the claims of the ’494 patent. Pet. at 6-7. In particular,

Claims Challenged	35 U.S.C. §	References
1	102/103	Bose ¹
3, 14-15	103	Bose, White, ² Simeral ³
2	103	Bose, Anderson ⁴
13	103	Bose, Anderson, Naffziger ⁵
4, 7	103	White
5, 6	103	White, Anderson
17, 18	103	White, Anderson, Naffziger
1, 3, 14, 15	103	White, Simeral
2	103	White, Simeral, Anderson

¹ U.S. Patent 7,421,601

² US Patent 7,263,457.

³ US PGPUB 2009/0150689.

⁴ US Patent 8,601,300.

⁵ US Patent 8,510,582.

13	103	White, Simeral, Anderson, Naffziger
----	-----	--

Id.

III. THE PETITION FAILS TO DEMONSTRATE THAT ANY CHALLENGED CLAIM IS UNPATENTABLE

A petition must be judged as it is written. The Board may not overlook flaws in the petition or reformulate its challenges to enhance their suitability. *See SAS Inst. Inc. v. Iancu*, 138 S. Ct. 1348, 1355 (2018) (“[I]t’s the petitioner, not the Director, who gets to define the contours of the proceeding.”); *see also id.* (“[T]he petitioner is master of its complaint . . .”); *see also id.* at 1353 (forbidding the Board from “curat[ing]” the petition). Furthermore, because all of the challenges in the petition rely on conclusions of obviousness, the analyses set forth in the petition must include “some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *KSR Int’l v. Teleflex Inc.*, 550 U.S. 398, 418 (2007) (*quoting In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)). The challenger must supply the factual basis for the challenge and “may not resort to speculation, unfounded assumptions or hindsight reconstruction to supply deficiencies in its factual basis.” *In re Warner*, 379 F.2d 1011, 1017 (CCPA 1967). Additionally, in cases such as the present, where the

obviousness challenges involve a modification or combination of teachings, the Petitioner must also establish that one skilled in the art would have had a reasonable expectation of success to make the modification or combination. *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367-68 (Fed. Cir. 2016). The modification or combination must not “require a substantial reconstruction and redesign of the elements” in a reference or “a change in the basic principles under which [a reference] was designed to operate.” *In re Ratti*, 270 F.2d 810, 813 (CCPA 1959). Nor can the modification or combination render a reference unsatisfactory or “inoperable for its intended purpose.” *In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984).

Also, the prior art must be considered for all that it teaches, and a *prima facie* case of obviousness may be rebutted “by showing that prior art teaches away from claimed invention in any material respect.” *In re Peterson*, 315 F.3d 1325, 1331 (Fed. Cir. 2003).

Last but not least, obviousness cannot be established by hindsight combination to produce the claimed invention, as the Supreme Court has emphasized:

A factfinder should be aware, of course, of the distortion caused by hindsight bias and must be cautious of arguments reliant upon *ex post* reasoning. See *Graham [v. John Deere Co.]*, 383

U.S. [1,] 36, 86 S. Ct. 684, 15 L. Ed. 2d 545 [(1966)] (warning against a “temptation to read into the prior art the teachings of the invention in issue” and instructing courts to “guard against slipping into use of hindsight” (quoting *Monroe Auto Equip. Co. v. Heckethorn Mfg. & Supply Co.*, 332 F.2d 406, 412 (6th Cir. 1964)).

KSR, 550 U.S. at 421.

As explained below, the petition’s challenge fails to demonstrate that any challenged claim is unpatentable.

A. Bose Fails to Teach or Suggest the Invention Recited in Challenged Independent Claim 1

MB USA challenges the patentability of claim 1 as being anticipated by or obvious in view of Bose. Pet. at 6. However, given the number of distinctions between the systems and methods described by Bose and those recited in claim 1 of the ’494 patent, this challenge lacks merit.

1. Overview of Bose

Bose describes “systems and methods for a power/performance monitoring and control unit.” Ex. 1005 at 1:16-22. The system includes a global power-performance monitor and control unit (PMCU) 110, that

“chooses appropriate power optimizations” for various core elements 112-1, 112-2, 112-R, an L2/L3 cache hierarchy 106, and an interconnect 108. Ex-1005 at 7:8-11. The global PMCU receives reports of power usage and temperature usage from local PMCUs 114, 116, 118-1, 118-2, 118-R associated with these elements. *Id.* at 6:67 – 7:3. This arrangement is illustrated in Bose’s FIG. 1:




FIG. 1

Bose’s approach to power assignment among the different system elements is to “package” various power optimizations into different power modes according to power vs. performance tradeoffs. *Id.* at 7:11-13.

“Higher” modes are associated with greater power consumption and higher

performance, while “lower” modes are associated with less power consumption and lower performance. *Id.* at 7:13-17. The power modes thus detail assignment of an overall chip power budget among the chip components and a particular power mode is chosen based on a system objective. *Id.* at 7:18-31. For example, “When the main objective of the system is performance, the global PMCU 110 will choose the power modes and assign power for each core element 112-1, 112-2, 112-R, 106, and 108 such that all elements can run at the highest possible performance within the budget.” *Id.* at 7:24-28.

The global PMCU “operates in a looping manner,” as described in Bose’s FIG. 3, reproduced here. *Id.* at 8:62 – 9:6. Following initialization,




FIG. 3

the global PMCU checks whether the actual total power being used by the system is at or below the total system power budget (303). If so, all cores are run in the highest power mode and the power budget is divided equally among the cores (312). *Id.* at 9:3-10.

Otherwise, if the actual power being used exceeds the system power


budget, a power mode is selected given consideration to the total system power budget, real time statistics reported by the local PMCUs, and power vs. performance tradeoffs of the various power modes (311). *Id.* at 9:14-24. In particular, a lower power mode is assigned to a core if its assigned portion of the power budget is insufficient, and a higher power mode is assigned to a core if that core is exceeding its assigned portion of the power budget such that “the sum of power usage for all components fits into the overall power budget.” *Id.* at 9:25-30. The local PMCUs operate with the newly assigned power modes/budgets and continue to report information back to the global PMCU. *Id.* at 9:26-43.

The various information relied upon by the global PMCU when making power budget allocations

is stored in a collection of registers and tables, as illustrated in Bose’s FIG. 2, reproduced here.

Connections from sensors associated with the local PMCUs 240 supply performance information for a phase history

table 213, while information concerning the overall power budget is stored



in total power budget register 212. *Id.* at 7:41 – 8:10. The global PMCU uses this power budget and performance information along with phase prediction information (stored in phase prediction table 204), previous decision information (stored in decision history table 205), and power/performance tradeoff information (stored in power-performance tradeoff table 206) to make decisions concerning allocation of the power budget. *Id.* at 8:19-51.

In particular, the power-performance tradeoff table 206 “includes the predicted power and performance for each power mode,” *id.* at 8:35-36, as “a structure (with update semantics and control) maintained to record the information needed to perform accurate prediction of the changes and duration of distinct workload phases, characterized by power and performance data as recorded within the [phase history table].” *Id.* at 12:54-58. The phase history table maintains a record of information related to power, temperature, voltage, and performance metrics that is periodically supplied by the resources monitored by the local PMCUs. *Id.* at 12:49-54.

2. Bose does not anticipate claim 1

Among other things, challenged claim 1 recites, “balancing control adapted to receive the bus workload from the bus workload monitor and to dynamically tune power allocation between the first core and the

communication bus based on a power limit for the integrated circuit and a comparison between the bus workload and a bus workload threshold.” Ex. 1001 at 21:28-33. The petition alleges that Bose’s global PMCU functions as the recited “balancing control,” Pet. at 23, and that “Bose discloses dynamic tuning of power allocation based on a comparison between bus workload and bus workload threshold [because a] POSITA would have understood that in Bose, selecting or adjusting power mode for a component entails comparing a workload threshold, such as predicted power and performance associated with each power mode with real time statistics, such as the monitored workload of a component.” *Id.* at 25-26. This is not true and Bose does not teach at least this element of claim 1. *See* Ex. 2001 at ¶¶ 44-55.

The “phase prediction technique” described by Bose, Ex. 1005 at 7:4-7, makes use of power modes that have associated power vs. performance tradeoff characteristics. *Id.* at 7:11-13. Power modes are assigned by the global PMCU according to a total power budget for the chip, current power consumption statistics, and the power vs. performance tradeoffs associated with the different modes. *Id.* at 9:14-32. The power vs. performance tradeoffs are represented in a power-performance tradeoff table, which stores “the predicted power and performance for each power mode.” *Id.* at 8:35-36.

The petition alleges, that the “predicted power and performance represents a workload threshold for a given component working in the associated ‘power mode.’” Pet. at 25-26 (citing Wyatt Decl. at ¶ 53). This is false. Dr. Brogioli notes that the information in Bose’s power-performance tradeoff table is dynamically updated over time. Ex. 2001 at ¶ 45. According to Bose, “The global PMCU controller 216 updates the Power/Performance Trade-off Table 206 through connection 209, after considering information from the Phase Prediction Table 204 and the Decision History Table 205.” Ex. 1005 at 8:39-43. Dynamically changing data of this kind, which is apparently derived algorithmically from historical information and current data samples, does not represent a threshold, a singular value against which a measured quantity is evaluated, and a person of ordinary skill in the art would appreciate as much. Ex. 2001 at ¶¶ 44-48. Indeed, Bose indicates, “The [phase prediction table] includes a structure (with update semantics *and control*),” Ex. 1005 at 12:54-55, and a POSITA would appreciate that if a structure includes semantics and control, that is likely something pertaining to the algorithm itself to be used rather than data (like a threshold) to use in variables. Ex. 2001 at ¶ 49.

The petition goes on to allege that, “A POSITA would have understood global PMCU compares a workload threshold as exemplified by

the predicted power and performance of each power mode and the monitored workload to select and adjust the power mode of a component.” Pet. at 26 (citing Wyatt Decl., ¶¶ 54-55). Notably, the petition fails to allege that Bose actually *teaches* such a comparison. Anticipation requires that a reference “disclose[] within the **four corners** of the document not only all of the limitations claimed but also all of the limitations arranged or combined in the same way as recited in the claim.” *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1371 (Fed. Cir. 2008); *see also Trintec Industries, Inc. v. TOP-USA Corp.*, 295 F.3d 1292, 1294, 1296 (Fed. Cir. 2002) (recognizing that anticipation is a highly technical condition that requires a single prior art reference to “expressly or inherently describe each and every limitation set forth in the patent claim” and that such “cases involving novelty, with its strict identity requirement are quite rare” such that a claimed method requiring a color *photocopier* was not anticipated by a prior art catalog that disclosed a color *printer*.). By relying on Mr. Wyatt’s declaration rather than the teachings of Bose itself, the petition tacitly admits that Bose fails to anticipate claim 1. *See* Patent Trial and Appeal Board Consolidated Trial Practice Guide, Nov. 2019 (“Consul. TPG”) at 36 (“Expert testimony [] cannot take the place of a disclosure in a prior art reference, when that disclosure is required as part of the unpatentability analysis.”).

Further, a person of ordinary skill in the art would **not** come to the understanding alleged by Petitioner and Mr. Wyatt. Ex. 2001 at ¶ 50. In Bose, power modes are assigned with, among other things, consideration of power vs. performance tradeoffs associated with each mode. *Id.*; Ex. 1005 at 9:14-24. The “information needed to perform accurate prediction of the changes and duration of distinct workload phases, characterized by power and performance data,” is maintained in a phase prediction table, Ex. 1005 at 12:54-58, and a power mode transmitted to a monitored resource will be “acted upon to place that resource into a specified configuration, such that its power level is decreased or increased (with a corresponding decrease or increase in net performance).” *Id.* at 12:37-4. In Bose it is not necessary, when considering power vs. performance tradeoffs, for there to be any *comparison* of bus workload to a bus workload threshold because direct information concerning actual power level decreases or increases for different power modes is available from the power-performance tradeoff table. Ex. 2001 at ¶ 50 (citing Ex. 1005 at 2:62-65 (“The amount of such change in power or performance is available from, e.g., an (updatable) power-performance tradeoff table (PPTT) maintained at the appropriate level of the control hierarchy.”), and 12:49-54). Thus, the global PMCU can directly determine whether a power level associated with a prospective

power mode will “fit” within a total power budget. *Id.* If anything then, a person of ordinary skill in the art would appreciate there is no need for the comparison hypothesized by petitioner and Mr. Wyatt.

Accordingly, because Bose does not teach all of the requirements of claim 1, the claim is not anticipated by Bose. *PPG Industries v. Guardian Industries Corp.*, 75 F.3d 1558, 1566 (Fed. Cir. 1996) (“To anticipate a claim, a reference must disclose every element of the challenged claim and enable one skilled in the art to make the anticipating subject matter.”); *see also Zenith Elecs. Corp. v. PDI Commc’ns Sys., Inc.*, 522 F.3d 1348, 1363 (Fed. Cir. 2008).

3. Claim 1 is not obvious in view of Bose

The petition makes no distinctions between its allegations of anticipation or obviousness of claim 1 with respect to Bose. *See* Pet. at 15-27. Nevertheless, claim 1 is not unpatentable for either reason. In particular, claim 1’s requirement of balancing control adapted to receive the bus workload from the bus workload monitor and to dynamically tune power allocation between the first core and the communication bus based on . . . the bus workload and a bus workload threshold is not suggested by Bose. *See* Ex. 2001 at ¶¶ 51-55.

An obviousness analysis measures the difference between the claimed invention and the prior art to determine whether “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. *Alza Corp. v. Mylan Labs., Inc.*, 464 F.3d 1286, 1289 (Fed. Cir. 2006) (citations omitted). As explained by Dr. Brogioli, there are a number of differences between the invention recited in claim 1 and the systems and methods discussed in Bose. For example, the claimed invention does not resemble the hierarchical and distributed control system described by Bose. Ex. 2001 at ¶¶ 51-52.

Bose is architected in a hierarchical manner, with each layer of the hierarchy having its own control. *Id.* at ¶ 51 (citing Bose, Ex. 1005 at 2:25); *and see* Ex. 2002 at 17:11 – 18:6; 24:19 – 25:15. A top level monitor “provides modes of operation and budgets to lower levels” while “the lowest levels provide feedback to the higher levels to control and manage power usage in the microprocessor system globally and locally.” Ex. 1005 at 2:11-15. In addition, the architecture of the various monitors within Bose vary, and overall are heterogeneous in nature. Some monitors can communicate upstream and downstream, while others can only communicate downstream. Ex. 2001 at ¶ 51 (citing Bose, Ex. 1005 at 2:27-35). Given this construct, a person of ordinary skill in the art would understand that control of power

allocation in Bose is distributed and heterogeneous, with local monitors able to make local decisions, to some degree, and also perform functions that upper layers cannot (and consume different sets of input/output tables than higher layers). *Id.* (citing Bose at 2:27-35, 10:44 – 11:53); *and see* Ex. 1005 at 11:16-28 (“To obtain the best performance out of a given power budget and power mode, the local PMCU controller 512 uses self-monitoring of power and performance The local PMCU controller 512 uses these statistics to decide which microarchitectural techniques to be employed and also their corresponding parameters to achieve the desired power vs. performance tradeoffs.”).

With regard the use of and comparison with a communication bus workload threshold, as recited in claim 1 of the '494 patent, Dr. Brogioli notes that here too Bose differs in important aspects. Ex. 2001 at ¶¶ 53-55. For example, “Bose states that with regard to the power/performance budget for the chip, ‘monitored power and performance levels received from the lower-level resources *may be translated* by the intermediate layer(s) of the control hierarchy *into a set* of power modes that are transmitted downwards’” *Id.* at ¶ 53 (citing Ex. 1005, Bose at 2:52-58). “A power mode transmitted to a monitored resource will be ‘acted upon to place that resource into a specified configuration, such that its power level is decreased

or increased (with a corresponding decrease or increase in net performance).”” *Id.* (citing Bose, Ex. 1005 at 12:37-4). And, the distributed and seemingly heterogeneous power budget allocation algorithm of Bose “may be further guided in its decision by inputs derived from a Phase History Table (PHT), a Phase Prediction Table (PPT) and a Decision History Table (DHT).” *Id.* at ¶ 54.

The phase prediction table stores “information needed to perform accurate prediction of the changes and duration of distinct workload phases, characterized by power and performance data,” Ex. 1005 at 12:54-58. Dr. Brogioli explains, “such a table itself does not constitute a threshold, and a POSITA would appreciate that the values in this table (as well as the others described by Bose) would be consumed via Bose’s distributed and heterogeneous algorithm to make decisions on the claimed power management of Bose.” Ex. 2001 at ¶ 54 (citing Bose, Ex. 1005 at 2:66 – 3:2). Indeed, as the information in these tables includes update semantics and control, a POSITA would appreciate they are not thresholds used for comparisons, but rather algorithmic information. *Id.*; *and see* Ex. 1005 at 12:46-61; Ex. 2002 at 27:1-20 (Mr. Wyatt indicating that “semantics” would relate to modes of operation).

Finally, Bose makes it unnecessary to perform any *comparisons* of

bus workload to a bus workload threshold, as alleged by Petitioner. “By providing a power-performance tradeoff table with information concerning actual power level decreases or increases for different power modes, no such comparison is needed.” Ex. 2001 at ¶ 55; *and see* Ex. 1005 at 2:62-65 (“The amount of such change in power or performance is available from, e.g., an (updatable) power-performance tradeoff table (PPTT) maintained at the appropriate level of the control hierarchy.”); Ex. 2002 at 26:7-12. “Having direct information concerning the amount of change in power or performance for specified power modes in the power-performance tradeoff table means the global PMCU can directly determine whether a power level associated with a prospective power mode will ‘fit’ within a total power budget.” Ex. 2001 at ¶ 55. Not only does the petition fail to articulate why a person of ordinary skill in the art would modify the distributed and hierarchical power allocation system described by Bose for one that relied upon *comparisons* of bus workload to a bus workload threshold, Dr Brogioli explains why such a person of ordinary skill would not do so especially since it would be unnecessary given that the disclosed power-performance tradeoff table already provides direct information concerning the amount of change in power or performance measures for different power modes. *Id.*

Accordingly, claim 1 is not obvious in view of Bose. *See In re Fine*,

837 F.2d 1071, 1074 (Fed. Cir. 1988) (any conclusion that claimed subject matter is prima facie obvious must be supported by evidence, as shown by some objective teaching in the prior art or by knowledge generally available to one of ordinary skill in the art that would have led that individual to combine the relevant teachings of the references to arrive at the claimed invention); *CFMT, Inc. v. Yieldup Int'l. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003) (“obviousness requires a suggestion of all limitations in a claim”) (citing *In re Royka*, 490 F.2d 981, 985 (CCPA 1974)).

B. The Remaining Challenges Based on Bose Also Fail

MB USA challenges the patentability of claims 2, 3, and 13-15, each of which depends directly or indirectly from claim 1, as being obvious in view of Bose when considered in combination with one or more additional references. Pet. at 6-7. In particular, claim 2 is alleged to be obvious in view of Bose and Anderson, claims 3, 14, and 15 are alleged to be obvious in view of Bose, White, and Simeral, and claim 13 is alleged to be obvious in view of Bose, Anderson and Naffziger. *Id.* However, in each of these assertions, the petition relies on its analysis of the teachings of Bose with respect to claim 1 insofar as those elements of the various dependent claims are concerned. *See Id.* at 28 and 35-36 (claims 3, 14, and 14); 39 and 45-46

(claims 2 and 13). As demonstrated above, that analysis vis-à-vis claim 1 in view of Bose is deficient. Accordingly, as to the challenges to dependent claims 2, 3, and 13-15 that depend on Bose as a primary reference, the petition fails to make out a prima facie case of obviousness. Each respective dependent claim is not unpatentable in view of the cited references for at least the same reasons as independent claim 1 is not unpatentable in view of Bose. *Hartness Int'l. Inc. v. Simplimatic Engineering Co.*, 819 F.2d 1100 (Fed. Cir. 1987)) (a claim that depends from a nonobvious independent claim is nonobvious because it contains all of the limitations of that independent claim plus a further limitation); *CFMT, Inc.*, 349 F.3d at 1342.

C. The Combination of White and Simeral Fails to Teach or Suggest the Invention Recited in Challenged Independent Claim 1

MB USA challenges the patentability of claim 1 as being obvious in view of the combined teachings of White and Simeral. Pet. at 7. Like the challenges based on Bose, however, these too fail.

1. Overview of White

White describes, “A processor or other integrated circuit [that] may include multiple logic cores, each of which may be configured to operate at frequencies and voltages independently of other logic cores.” Ex. 1006 at

2:15-18. White’s FIG. 1, reproduced here, depicts “a block diagram illustrating an integrated circuit including multiple logic cores configured to operate at independent voltages and/or




FIG. 1

frequencies.” *Id.* at 5:18–31; 6:43–

46. The common bridge logic 110, is “configured to communicate with the logic cores as well as with other elements or peripherals either internal or external to integrated circuit 100, such as bus 150, peripheral 160 and/or memory subsystem 170[, and] may be configured to operate at a voltage or frequency different from, and independent of, the voltages or frequencies at which the logic cores 120A and 120B are operating.” *Id.* at 7:3-11.

According to White, the operating voltage and/or operating frequency of the various logic cores of integrated circuit 100 may be controlled by power management logic 200, as illustrated in FIG. 2A, reproduced here. *Id.* at 8:61-




FIG. 2A

65. In this way, overall power consumption for the integrated circuit may be adjusted. *Id.* at 9:1-6. For example, the voltage and/or frequency of the cores 120A, 120B may be lowered when the system is operating on battery power, and raised in other cases. *Id.* at 9:6-12. Similarly, the voltage and/or frequencies supplied to the common bridge logic 110 may be adjusted by the power management logic 200, for example, to reduce potential bottlenecking of the common bridge logic. *Id.* at 9:30-53.

White explains a method for independent power management of the logic cores with reference to FIGs. 4 and 5, reproduced below.⁶ When a need

⁶ An additional approach is described with respect to FIG. 6 and involves reducing the operating characteristics for one or more logic cores in steps or stages. *Id.* at 20:58 *et seq.*

to change an operating characteristic, e.g., frequency (FIG. 4) or voltage (FIG. 5), of a logic core is detected (400, 500), power management logic 200 determines a new operating voltage/frequency for the core (410, 510). *Id.* at 15:42 – 16:62; 18:20 – 19:42. New frequencies may be determined based on a “standard algorithm” or “from a look-up table.” *Id.* at 16:58-63. New




FIG. 4




FIG. 5

voltages may be determined via a voltage determining algorithm . . . based upon current operating conditions for the system, such as current operating temperature, current power availability or current power consumption, or may be determined based upon a fixed table of operating voltages. *Id.* at 19:30-38.

Assuming the new frequency is compatible with the operating voltage of the core (420), or after making adjustments therefor (430, 440), the power management logic sets the new operating frequency for the core (450), *id.* at 16:63 – 17:28, and thereafter, adjusts the operating conditions for the common bridge logic to be compatible with the new operating frequency of the core. (460). *Id.* at 17:29 – 18:8.

In the case of a new operating voltage for a core, after determining that new operating voltage the power management logic sets the new voltage as the operating voltage for the core (520), *id.* at 19:43-47, determines whether the operating frequency of the core is compatible with the new voltage (530), *id.* at 19:53-57, and makes adjustments to the operating frequency as necessary (540, 550). *Id.* at 19:57-65; 20:19-23. Thereafter, the power management logic adjusts the operating conditions for the common bridge logic to be compatible with the new operating voltage and/or frequency of the core. (560). *Id.* at 20:39-57.

2. Overview of Simeral

Simeral describes “a method for integrating power management functions into a data path controller to manage power consumption of processors and peripheral devices.” Ex. 1008 at [0007]. Simeral does not

offer a precise definition of the term “data path controller,” but identifies it as “resid[ing] in a computing device” and “form[ing] data paths over which data are exchanged among components of a computing device.” *Id.* at [0019]. These “components” might be a processor and devices with which a processor communicates, e.g., peripherals, graphics processors, hard drive controllers, display drivers, and “regulatory devices” such as voltage regulators or frequency/clock generators. *Id.* The data path controller can be “reactive,” taking action to decrease power consumption as a result of low or negligible activity levels by modifying operational characteristics of the components to conserve power while providing sufficient resources to support a specific workload, or “proactive,” taking action to decrease power consumption by decreasing activities levels of a component through modification of operational characteristics of the component as a corrective

action to bring one or more noncompliant activity levels back into compliance. *Id.* at [0020].

Simeral provides FIG. 1A, reproduced here, which shows a data path controller 100 with a power management controller 106 and a power-monitoring interface 101. Power-monitoring interface 101 includes a number of general-purpose interface signal ("GPIS") ports 104a, 104b, and 104c. The GPIS ports are power monitoring ports that can be configured to monitor and quantify activity levels of a processor or a peripheral device for power management purposes. For example, a GPIS port can provide clock clamping, frequency reduction/modification, and the like. GPIS port "k" (191) can monitor activity associated with bus 130. *Id.* at [0019].

According to Simeral,

[D]ata path controller 100 can manage power for a computing device by first monitoring the activity level of the bus (e.g., the data




FIG. 1A

transfer rate of the bus), and then disabling (e.g., shutting down) components that have very low (i.e., negligible) activity levels to save power. So if the data rate is below some threshold (i.e., low amounts of data are passing over the bus), then data path controller 100 can either decrease (i.e., “throttle down”) the bus width or shut down one or more bus controllers, thereby conserving power that would otherwise be expended by maintaining a higher bandwidth.

Id. at [0021]. It is the power management controller 106 of data path controller 100 that “implement[s] the power management functions.” *Id.* at [0024]. In particular, “power management controller 106 monitors the activity levels of the component, such as the data being exchanged between a component and data path controller 100 . . . compares that monitored activity level to a threshold activity level set forth in [a] performance profile, [and] if the monitored activity level falls out of compliance, then power management controller 106 initiates corrective action . . . include[ing] modifying the behavior of the noncompliant component so that it complies with its threshold activity level [or] recalibrate[ing] activity levels of other less essential components to compensate for the increased activity level of the noncompliant component.” *Id.* A performance profile can include

operational characteristics such as voltages, frequencies, data throughput on buses, etc. *Id.* at [0026].

3. The combination of White and Simeral proposed by Petitioner fails to teach or suggest a core workload monitor configured to determine a core workload

The petition alleges that in a White-Simeral combination, it is White that discloses a core workload monitor configured to determine a core workload for the first core, as recited in claim 1. Pet. at 72 (“White discloses Claim 1[e]”). In particular, the petition identifies White’s power management logic 200 and alleges “A POSITA would have understood White monitoring the operating voltages and frequencies of the logic core 120A and common bridge logic 110 is monitoring the workload of those components.” *Id.* This is not true.

In White, power management logic 200 controls operating characteristics of cores of an integrated circuit, such as operating voltage and/or operating frequency. Ex. 1006 at 8:61-65; 9:2-6. For example, if the power management logic determines that a core is consuming more power than is allocated to it under a particular mode of system operating, the power management logic may issue a request for the core to adjust its frequency of operation. *Id.* at 9:56-62. And, the power management logic may lower the operating characteristics of one or more cores when the system is operating

on battery power, or raise the operating characteristics when the system is not on battery power. *Id.* at 9:6-12. In yet another example, the power management logic may adjust the operating characteristics of one core to compensate for the operating characteristics of another core. *Id.* at 9:30-36.

From White's description, a person of ordinary skill in the art "would understand that White's power management logic controls adjustments to a core's operating characteristics by monitoring the core's current power consumption." Ex. 2001 at ¶ 64. White explains that any adjustments to a core's operating characteristics are made via voltage regulator 130 and/or frequency regulator 140, as illustrated in White's FIG. 2, above. *Id.* at ¶ 65 (citing White, Ex. 1006 at 10:11-24). "The voltage/frequency regulators receive requests for alterations in the operating characteristics of the cores from the power management logic, and the power management logic determines current output voltage/frequency levels for a core by reading configuration registers of the voltage/frequency regulators. *Id.* (citing White, Ex. 1006 at 10:40-46); Ex. 2002 at 30:12 – 31:23. And, when adjusting the operating voltage or frequency of a core, the power management logic uses a table of compatible or desirable voltage-frequency pairs to ensure desired operating characteristics are feasible. Ex. 1006 at 10:50 – 11:23; Ex. 2002 at 32:19 – 33:8.

Dr. Brogioli explains that, the “focus on operating voltage and operating frequencies is an indication that White’s power management logic makes decisions regarding future operating characteristics of the core based on present operating characteristics.” Ex. 2001 at ¶ 65. Monitoring operating characteristics, as is done by White’s power management logic, is not the same as monitoring workload using a core workload monitor, as recited in claim 1. *Id.* at ¶ 66. This is because, as Dr. Brogioli notes, “performance metrics such as frequency can be altered independently of workload.” *Id.* (citing the ’494 patent, Ex. 1001 at 12:61-65 (“[I]f a power limit is reached, then GPU 605’s performance metric, such as frequency, may be reduced, since it doesn’t need full performance to complete its workload that was below an activity threshold.”)). In contrast, according to the ’494 patent, monitoring the workload of a core involves “determining activity of hardware, a workload of software, a number of instructions executed or to execute, or any other metric for determining use of a processing element. As one example, a workload for core 505 includes an amount time core 505 is active during a period of time. Activity may be determined as when core 505 is in an active performance state versus a sleep or low power state or a number of cycles executing or performing operations.” Ex. 1001 at 9:50-57.

A person of ordinary skill in the art would not conflate monitoring of

operating characteristics, such as voltage or frequency as in White, with monitoring of workload, as is called for in claim 1 of the '494 patent. Ex. 2001 at ¶ 66. As the '494 patent makes clear, one is not a substitute for the other as operating characteristics need not reflect the actual workload of the core. *Id.* Thus, claim 1 is not obvious in view of White and Simeral, as alleged in the petition. *CFMT, Inc.*, 349 F.3d at 1342.

4. Claims 3, 14, and 15 are also patentable over White and Simeral

Challenged claims 3, 14, and 15 depend from claim 1. Ex. 1001 at 21:50-59; 24:1-7. Each respective dependent claim is not unpatentable in view of White and Simeral for at least the same reasons as independent claim 1. *Hartness Int'l.*, 819 F.2d 1100.

5. The remaining challenges based on White and Simeral also fail

Claim 2 is challenged based on the combined teachings of White, Simeral, and Anderson. Pet. at 7. Claim 13 is challenged based on the combined teachings of White, Simeral, Anderson, and Naffziger. *Id.* However, neither Anderson nor Naffziger is cited for curing the deficiencies of White and Simeral discussed above with respect to independent claim 1.

See Pet. at 83 (citing White alone for teaching a second core and a second core workload monitor); 84 (referencing Pet. at 64-65, citing Naffziger for teaching a ‘consumer/producer relationship’ between processor elements). Accordingly, claims 2 and 13 are not unpatentable in view of the combined teachings of White, Simeral, Anderson, and Naffziger for at least the same reasons as independent claim 1. *Hartness Int’l.*, 819 F.2d 1100.

D. White Fails to Teach or Suggest the Invention Recited in Challenged Independent Claim 4

MB USA challenges the patentability of claim 4 as being obvious in view of White. Pet. at 7. However, this challenge fails because independent claim 4 recites “receiving a bus workload for a communication bus and a first processing element workload for a first processing element based on an integrated circuit reaching a limit,” and White fails to teach or suggest receiving such workloads based on an integrated circuit reaching a limit. Ex. 2001 at ¶¶ 67-72.

Similar to its argument with respect to claim 1, the petition alleges, “White discloses power management logic receives workloads for common bridge logic 110 and logic core 120A, such as their operating voltages and frequencies.” *Id.* at 53. The alleged rationale for this statement is that, “A POSITA would have understood that operating voltage/frequency reflect the

activity and use, or workload, of a component.” *Id.* But, as discussed above with respect to claim 1, this is not true.

White’s focus on operating voltage and operating frequencies is an indication that White’s power management logic makes decisions regarding future operating characteristics of a core based on present operating characteristics. Ex. 2001 at ¶ 67. Monitoring operating characteristics, as is done by White’s power management logic, is not the same as monitoring a processing element’s workload using a workload monitor, as recited in claim 4. *Id.* In fact, the ’494 patent itself explains, performance metrics such as frequency can be altered independently of workload. *See, e.g.*, Ex. 1001 at 12:61-65 (“[I]f a power limit is reached, then GPU 605’s performance metric, such as frequency, may be reduced, since it doesn’t need full performance to complete its workload that was below an activity threshold.”). And, according to the ’494 patent, monitoring the workload of a core involves “determining activity of hardware, a workload of software, a number of instructions executed or to execute, or any other metric for determining use of a processing element. As one example, a workload for core 505 includes an amount time core 505 is active during a period of time. Activity may be determined as when core 505 is in an active performance

state versus a sleep or low power state or a number of cycles executing or performing operations.” *Id.* at 9:50-57.

Further, claim 4 differs from claim 1 in an important respect. Unlike claim 1, which recites, “monitoring” a workload, claim 4 recites “receiving” a bus workload and a processing element workload. *Id.* at 21:63-65. White is clear that what is “received” by power management logic is not workload – i.e., not activity levels. Instead, what is received by the power management logic is current operating voltages and operating frequencies. Ex. 1006 at 10:43-46 (the power management logic determines current output voltage/frequency levels for a core by reading configuration registers of the voltage/frequency regulators); Ex. 2001 at ¶ 68; Ex. 2002 at 31:12-23 (Mr. Wyatt agreeing that in White, voltage and frequency are read from configuration registers of the voltage/frequency regulators). Even the petition admits, “In determining power consumption of a component, power management logic 200 ‘determine[s] the current output voltage levels ... by accessing data stored in configuration registers for voltage regulator 130.’” Pet. at 53 (emphasis omitted). Even if “operating voltage/frequency reflect the activity and use, or workload, of a component,” as the petition alleges, receiving operating voltage/frequency, as in White, is **not** receiving

workload, as recited in claim 4 and a person of ordinary skill in the art would appreciate the difference. Ex. 2001 at ¶¶ 68-89.

The petition does not allege any modifications to White that would result in a system that *receives* a bus workload and a processing element workload, as recited in claim 4. Instead, the petition seems to allege that receiving operating voltage/frequency *is* receiving workload, but this is certainly not true. *Id.* Workload is “an amount of activity (e.g. number of cycles active) over a quantum or period of time.” Ex. 1001 at 8:20-22; *and see id.* at 9:49-57 (recognizing workload as “activity of hardware,” “a number of instructions executed or to execute,” “use of a processing element,” and “an amount [of] time core 505 is active during a period of time”); 9:61-63 (“workload is measured as a number of active cycles over a period of time (i.e., a total number of cycles for measurement)”); Ex. 2001 at ¶ 69. Nothing in the petition describes how a person of ordinary skill in the art would modify receipt of an output voltage level from a configuration register of a voltage regulator to receipt of an amount of time a processing element is active during a measurement period, nor does the petition allege it would be obvious to make such a modification.

And there are further failings of the petition with respect to claim 4. The petition contends that White discloses the power management unit

receiving operating voltages/frequencies based on an integrated circuit reaching a limit, such as a battery power limit. Pet. at 54. This is false. At the referenced sections of White (15:54-64), the disclosure reads,

In another embodiment, a system including a multi-core processor may be operating from battery power and power management logic 200 may be configured to detect when the system is operating on battery power and adjust the operating frequency of one or more of the logic cores. Additionally, power management logic 200 may be configured to detect when the system transitions from battery power to wall power and adjust the operating frequency of one or more logic cores to a higher frequency to take advantage of the fact that the system is no longer on battery power.

See also, Ex. 1006 at 18:25-39; 20:6-12; 24:7-29. There is no suggestion here that any workload is received when a system reaches a “battery power limit.” Indeed, there is no mention of a battery power limit at all. Ex. 2001 at ¶ 71. If anything, White specifies that the power management logic detects when the system is operating on battery power, irrespective of any “limit” whatsoever. *Id.*

And, insofar as the petition alleges, “White discloses that power management logic 200 receives workload information based on the integrated circuit reaching a thermal design power limit, such as ‘a predetermined maximum operating temperature,’” Pet. at 55, here the petition confuses the power management logic with different “temperature control logic.” Ex. 2001 at ¶ 72. Referring to FIG. 2B, reproduced below, White explains that operating characteristics (i.e., operating voltage and frequency) for each of multiple logic cores may be controlled by temperature control logic 210. Ex.

1006 at 11:61-64;

12:38-45. The

temperature control

logic, not the power

management logic, is

configured to read

current operating temperatures and request changes in operating

characteristics of the cores based on readings from temperature sensors. *Id.*

at 12:38-58. The petition mistakenly attributes such actions to the power

management logic, Pet. at 53-54, and in doing so fails to demonstrate why

claim 4 is obvious in view of White.




FIG. 2B

Accordingly, the petition fails to make out a prima facie case of obviousness with respect to claim 4. *CFMT, Inc.*, 349 F.3d at 1342. Said differently, claim 4 is not obvious in view of White, as alleged in the petition.

1. Claim 7 is also patentable over White

Challenged claim 7 depends from claim 4. Ex. 1001 at 22:39-40. Therefore, dependent claim 7 is not unpatentable in view of White for at least the same reasons as independent claim 4. *Hartness Int'l.*, 819 F.2d 1100.

2. The remaining challenges based on White also fail

Claims 5 and 6 challenged based on the combined teachings of White and Anderson. Pet. at 7. Claims 17 and 18 are challenged based on the combined teachings of White, Anderson, and Naffziger. *Id.* However, neither Anderson nor Naffziger is cited for curing the deficiencies of White discussed above with respect to independent claim 4. *See* Pet. at 60-62 (citing Anderson for teaching a GPU core and a ring interconnect); 64-66 (citing Naffziger for teaching a ‘consumer/producer relationship’ between processor elements). Accordingly, claims 5, 6, 17, and 18 are not

unpatentable in view of the combined teachings of White and Anderson, and White, Anderson, and Naffziger for at least the same reasons as independent claim 4. *Hartness Int'l.*, 819 F.2d 1100.

IV. CONCLUSION

For at least the foregoing reasons, the Board should find that Petition has failed to prove the unpatentability of any of the challenged claims.

Harmonic Inc. v. Avid Tech., Inc., 815 F.3d 1356, 1363 (Fed. Cir. 2016) (“In an [inter partes review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.”) (citing 35 U.S.C. § 312(a)(3) (requiring inter partes review petitions to identify “with particularity . . . the evidence that supports the grounds for the challenge to each claim”)); *Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015) (the burden of persuasion never shifts to the Patent Owner); *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016) (a petitioner cannot satisfy its burden of proving obviousness by employing “mere conclusory statements”).

Dated: May 7, 2024

Respectfully submitted,

/Tarek N. Fahmi/

Tarek N. Fahmi

Reg. No. 41,402

Ascenda Law Group, PC
2150 N First St., Suite 420
San Jose, CA 95131
Tel: 866-877-4883
Email: tarek.fahmi@ascendalaw.com

WORD-COUNT CERTIFICATE

I certify that the foregoing document complies with the type-volume limitation of 37 C.F.R. § 42.24(a) and (b) and contains 10,121 words (including words in drawings created by the Patent Owner or as annotations added by the Patent Owner to images reproduced in this document) based on the word count indicated by Microsoft® Word for Mac v. 16.49, excluding those portions exempted by the rule.

/Tarek N. Fahmi/

Tarek N. Fahmi

Reg. No. 41,402

CERTIFICATE OF SERVICE

The undersigned hereby certifies that a copy of the foregoing:

PATENT OWNER'S RESPONSE

and its accompanying exhibits were served on May 7, 2024, by filing this document through the P-TACTS System as well as by delivering a copy via email directed to the attorneys of record for the Petitioner at the following addresses:

Celine J. Crowson
Joseph J. Raffetto
Scott Hughes
Helen Y. Trac
Ryan Stephenson
Nicholas Rotz
Hogan Lovells US LLP
555 13th Street N.W.
Washington, D.C. 20004

Email:

celine.crowson@hoganlovells.com
joseph.raffetto@hoganlovells.com
scott.hughes@hoganlovells.com
helen.trac@hoganlovells.com
ryan.stephenson@hoganlovells.com
nicholas.rotz@hoganlovells.com

The parties have agreed to electronic service in this proceeding.

Respectfully submitted,

Dated: May 7, 2024

/Tarek N. Fahmi/

Tarek N. Fahmi
Reg. No. 41,402

Ascenda Law Group, PC
2150 N First St., Suite 420
San Jose, CA 95131
Tel: 866-877-4883
Email: patents@ascendalaw.com