

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.: IPR2023-01333

U.S. Patent No. 10,049,080

PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

I.	INTRODUCTION AND SUMMARY OF ARGUMENTS..	1
II.	BACKGROUND, THE '080 PATENT, AND CLAIM CONSTRUCTION.....	4
A.	Background	4
B.	The '080 Patent	6
C.	The Challenged Claims	9
D.	Claim Construction and the Level of Ordinary Skill in the Art	10
1.	Operating System	12
E.	The IPR Petition.....	13
III.	THE PETITION FAILS TO DEMONSTRATE THAT ANY CHALLENGED CLAIM IS UNPATENTABLE	14
A.	Sutardja Fails to Teach or Suggest the Invention Recited in the Challenged Independent Claims	17
1.	Sutardja '748.....	17
2.	Sutardja '785.....	19
3.	Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as claimed	20
B.	Claims Depending from Claims 1, 9, and 17 are Not Obvious in View of Sutardja.....	28
C.	The Remaining Challenges that Rely on Sutardja as a Primary Reference Likewise Fail	29
D.	Mathieson-Sutardja Fails to Teach or Suggest the Invention Recited in the Challenged Independent Claims	31

1.	Overview of Mathieson	31
2.	Mathieson-Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as claimed.....	34
3.	Mathieson-Sutardja does not teach or suggest a first plurality of cores and a second plurality of cores that support a same instruction set, wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores, as claimed.....	38
4.	The Challenged Dependent Claims are Also Not unpatentable in View of Mathieson-Sutardja	43
E.	The Remaining Challenges that Rely on Mathieson-Sutardja Likewise Fail.....	43
IV.	CONCLUSION	44

TABLE OF AUTHORITIES

CASES

<i>ACTV, Inc. v. Walt Disney Co.</i> , 346 F.3d 1082 (Fed. Cir. 2003).....	11
<i>Arendi S.A.R.L. v. Apple Inc.</i> , 832 F.3d 1355 (Fed. Cir. 2016).....	21
<i>CFMT, Inc. v. Yieldup Intern. Corp.</i> , 349 F.3d 1333 (Fed. Cir. 2003).....	passim
<i>Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.</i> , 800 F.3d 1375 (Fed. Cir. 2015).....	45
<i>Harmonic Inc. v. Avid Tech., Inc.</i> , 815 F.3d 1356 (Fed. Cir. 2016).....	44
<i>Hartness Int’l. Inc. v. Simplimatic Engineering Co.</i> , 819 F.2d 1100 (Fed. Cir. 1987).....	29, 30, 43, 44
<i>In re Gordon</i> , 733 F.2d 900 (Fed. Cir. 1984).....	16
<i>In re Magnum Oil Tools Int’l, Ltd.</i> , 829 F.3d 1364 (Fed. Cir. 2016).....	45
<i>In re Peterson</i> , 315 F.3d 1325 (Fed. Cir. 2003).....	16
<i>In re Ratti</i> , 270 F.2d 810 (CCPA 1959)	15
<i>In re Warner</i> , 379 F.2d 1011 (CCPA 1967)	15
<i>Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.</i> , 821 F.3d 1359 (Fed. Cir. 2016).....	15

<i>KSR Int’l v. Teleflex Inc.</i> , 550 U.S. 398 (2007).....	15, 16, 20
<i>Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co. Ltd.</i> , 868 F.3d 1013 (Fed. Cir. 2017).....	11
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	10, 11
<i>SAS Inst. Inc. v. Iancu</i> , 138 S. Ct. 1348 (2018).....	14, 15
<i>Takeda Chem. Indus., Ltd. v. Alphapharm Pty, Ltd.</i> , 492 F.3d 1350 (Fed. Cir. 2007).....	20
<i>Vitronics Corp. v. Conceptronic, Inc.</i> , 90 F.3d 1576 (Fed. Cir. 1996).....	11

REGULATIONS

37 C.F.R. § 42.104(b)(4)	28, 38
--------------------------------	--------

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)	10
--	----

EXHIBIT LIST

EXHIBIT NO.	DESCRIPTION
2001	US PGPUB 2008/0263324 A1 Sutardja et al.
2002	Transcript of Deposition of Dr. Horst
2003	US Patent 5,889,679 Henry et al.
2004	US Patent 8,892,931 Kruglick
2005	US Patent 9,086,883 Thomson et al.
2006	Declaration of Michael C. Brogioli, Ph.D.
2007	Excerpts from Microsoft Computer Dictionary, Fifth Edition (2002), pp. 378 (“operating system”), 518 (“thread”)

I. INTRODUCTION AND SUMMARY OF ARGUMENTS

Mercedes-Benz USA, LLC's ("MB USA" or "Petitioner") challenges to the claims of U.S. Patent No. 10,049,080 (the "'080 patent") fail to demonstrate that any challenged claim is unpatentable.

Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as required by all of the challenged claims, and the petition fails to provide any reasoned analysis that explains why a person of ordinary skill in the art would have modified Sutardja to provide such an operating system. Indeed, MB USA is inconsistent in defining where exactly the claimed "power management hardware" can be found in Sutardja. Initially, the petition alleges, "The hardware controlled by the core switching module corresponds to the claimed power management hardware." Pet. at 32-33. But later it claims, "A POSITA would have understood that the hardware on which the disclosed hypervisor is installed corresponds to the claimed power management hardware." *Id.* at 33-34. Still later we are told, "the operating system (kernel module and PMS module) controls the power management hardware (core switching module) based on the demand." *Id.* at 35. Regardless, MB USA's argument for unpatentability fails no matter whether one considers the core

switching module, the hardware controlled by the core switching module, or the hardware on which the hypervisor is installed as the claimed core switching module.

First, Sutardja's core switching module is not "power management hardware." Indeed, it is not "hardware" at all as Sutardja advises the core switching module, whether it is integrated in the hypervisor module or not, is loaded into high-speed memory from RAM or other storage after bootup. Hardware cannot be "loaded into high speed memory from RAM;" thus, the core switching module of Sutardja is software, not hardware.

Second, the petition itself relies on Sutardja's hypervisor "operat[ing] below the kernel." Hence, if one considers the recited "power management hardware" as Sutardja's the hardware on which the hypervisor is installed, then it cannot be that Sutardja's operating system controls this hardware. Instead, because the hypervisor operates "below the kernel" the hypervisor, not the operating system, controls this hardware. This is, in fact, the whole point of a so-called "bare metal hypervisor."

Third, if the recited "power management hardware" is the hardware controlled by Sutardja's core switching module, Sutardja is clear that, the kernel and the PMS module (the alleged "operating system") do not "control" this hardware and instead only make system calls to the core

switching module which actually effects the control. Accordingly, all of MB USA's challenges that depend on Sutardja as a primary reference fail.

The same is true of MB USA's challenges based on the combined teachings of Mathieson and Sutardja. For these challenges, MB USA relies on Mathieson for teachings concerning the power management hardware and operating system limitations of independent claims 1, 9, and 17. In particular, the petition alleges that Mathieson describes a controller within a CPU that monitors demand and that a POSITA would recognize such a controller performs the functions of an operating system and therefore meets the requirements of the claims.

Not so. If anything, the POSITA would understand that because Mathieson directly states that transitions between modes, which involve the controller transferring processor state from one set of cores to the other set of cores, are done *transparently* to the operating system. That is, it is **not** the operating system that is effecting the transition. Instead, for such transparency to exist, the transition between modes must be provided by a mechanism at a layer ***below the operating system***. Accordingly, and contrary to the allegation in the petition, the controller does **not** perform the functions of an operating system and instead operates below the operating system when controlling any power management hardware.

Moreover, while Mathieson might discuss power consumption of cores on the basis of operating frequency, nothing is said concerning such power consumption on the basis of an applied supply voltage. Hence, the petition fails to demonstrate that Mathieson-Sutardja, teaches or suggests, “wherein the second plurality of cores consume less power, for a same applied operating frequency *and supply voltage*, than the first plurality of cores.”

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 '080 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 semiconductor die. Ex. 1001 at 1:22-25. 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 die means a multi-core processor can execute multiple such instructions simultaneously. Ex. 2006 at ¶ 28.

The amount of power consumed by a processor is strongly correlated with its performance capabilities. Ex. 1001 at 2:26-29. Because decreasing the processing performance of the system corresponds to power savings, processor power management typically scales up the performance of the system as the system's workload increases, and scales down the performance as workload decreases. *Id.* at 2:22-27. A typical way to scale processing performance and power consumption with workload is to enable or disable entire cores and raise or lower their supply voltages and operating frequencies (and thereby their power consumption) in response to system workload. *Id.* at 2:30-34.

But solutions such as these have limitations. For example, merely adjusting the operating frequency (also called “clock”) and/or voltage at which a core operates to adjust its power consumption/processing capability balance presents core management difficulties. Ex. 2006 at ¶ 30. So too do other approaches to power management, such as the use of low-power states, have drawbacks, including latencies due to the required time to raise the cores out of the low power state when needed. *See* Ex. 1033 at 1:26-33 (“One approach to reducing power in a computing platform when there is relatively little activity, is to place the processor in a low-power state. However, placing a processor in a low-power state or returning a processor

from a low-power state may require a non-trivial amount of time. Therefore, it may or may not be worth the time required to place a processor in a low-power state or to return the processor from a low-power state.”).

B. The '080 Patent

The '080 patent proposes a different approach to processor power management. Rather than simply relying on adjustments to operating frequency and/or voltage or employing low power states when a core is inactive, the '080 patent describes a multi-core processor that makes use of different types of processor cores that consume different amounts of power and offer different performance capabilities (for a given operating frequency and/or voltage), yet which support the same instruction set. Ex. 2006 at ¶ 32 (citing Ex. 1001 at 3:50-62). To achieve this, the different types of processor cores can employ different power consumption design features, such as different transistor technologies. *Id.* Consequently, while the different types of processor cores can each execute the same instructions, “low power” cores will “exhibit inherently lower power consumption (and processing performance) than the higher power core(s),” even when the low power core is operated at the same clock frequency as a higher power core. *Id.* (citing Ex. 1001 at 3:53 – 4:16).

Because the two sets of cores consume different amounts of power and offer different performance capabilities by virtue of their power consumption design features, the cores are referred to as “asymmetric.” *Id.* (citing Ex. 1001 at 1:17-20). The use of such asymmetric cores allows a single processor to “scale up [it]s processing performance ... as the system’s workload increases, and scale down [it]s processing performance ... as the system’s workload decreases,” just as prior multicore processors did, but because processors designed according to the teachings of the ’080 patent incorporate inherently low power cores, such processors are able to “achieve an even lower power consumption in the lower/lowest performance/power states.” Ex. 1001 at 2:19-29, 3:50-62, 4:20-46.

One power management strategy for processors that employ asymmetric cores in accordance with the teachings of the ’080 patent is described with reference to the patent’s Fig. 6, reproduced below. The process begins with a multi-core processor in which multiple high power cores and at least one low power core are operating. *Id.* at Fig. 6 (601), 4:47–59. As demand on the processor decreases, the various high power cores are disabled one at a time, until all are disabled. *Id.* at 4:54 – 5:6. If all of the high power cores are disabled and the demand on the processor continues to drop, then in a similar fashion the low power cores are disabled one by one

until only one low power core remains enabled and the processor's lowest power state is reached. *Id.* at 5:25–35. During this process, before a core (high power or low power) is disabled, its individual operating frequency and/or supply voltage, or the operating frequencies and/or supply voltages of some or all of the enabled cores, may be reduced in response to the drop in demand. And, if only one low power core is operating, its operating frequency and/or supply voltage may similarly be reduced until a lowest power consumption state is reached. *Id.* at 5:7-24, 28-49.




FIG. 6

C. The Challenged Claims

In keeping with the above-described teachings, claim 1 of the '080 patent recites:

A multi-core processor comprising:

a first plurality of cores and a second plurality of cores that support a same instruction set, wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and

power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores, wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.

Id. at 7:56 – 8:3.

Independent claim 9 recites a method that includes operating a multi-core processor such that a first plurality of cores and a second plurality of cores execute a same instruction set, wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and disabling with power

management hardware, from a state where the first plurality of cores and the second plurality of cores are enabled, all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores, wherein an operating system executing on the multi-core processor monitors a demand for the multi-core processor and controls the power management hardware based on the demand. *Id.* at 8:38-53.

Independent claim 17 recites a non-transitory machine-readable medium containing program code that, when processed by a machine, causes a method as recited in claim 9 to be performed. *Id.* at 9:24-41.

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. Conceptronic, 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 10 at 15. That proposed level of ordinary skill is: a person with “a bachelor’s degree in electrical engineering, computer science, computer engineering, material science, physics, applied physics, or a related field” who would have had “at least two years of experience in the research, design, development, or testing of electronic circuits or components or software for controlling electronic circuits or components, or the equivalent,

with additional education substituting for experience and vice versa.” Pet. 14–15. The proposed definition is flawed insofar as it encompasses people with more than ordinary skill in the art—that is, by encompassing people with “at least” two years of experience in the designated field, it includes persons with far greater levels of experience and, therefore, is not limited to just persons of *ordinary skill* in the art. Patent Owner therefore uses the proffered definition of a person of ordinary skill in the art, but limits such individuals to those with two years of experience in the designated field in the analysis below. *See* Ex. 2006 at ¶ 25.

Although the Board did not make any claim constructions as part of its Institution decision, *see* Paper 10 at 15, certain terms of the claims will be important for understanding distinctions from the prior art cited by Petitioner:

1. Operating System

All of the challenged independent claims recite “an operating system” that executes on a multi-core processor to monitor demand for the multi-core processor and control power management hardware based on the demand. Ex. 1001 at 7:66 – 8:3 (claim 1); 8:49-53 (claim 9); 9:37-41 (claim 17). The ’080 patent cites examples of “operating systems,” including Apple’s

MacOS, Microsoft’s Windows, UNIX, and others. *Id.* at 7:31-35. A person of ordinary skill in the art would thus understand that the term “operating system” is used in the ’080 patent in its conventional sense. Ex. 2006 at ¶ 40. An “operating system” is, and would be understood by a person of ordinary skill in the art to be, “The software that controls the allocation and usage of hardware resources such as memory, central processing unit (CPU) time, disk space, and peripheral devices.” *Id.* (citing Ex. 2007 at 378 (Microsoft Computer Dictionary); *and see* Ex. 2002 at 32:1 – 34:19.

E. The IPR Petition

In its petition, MB USA presents multiple grounds for the alleged unpatentability of the claims of the ’080 patent. Pet. at 5-6. In particular,

Claims Challenged	35 U.S.C. §	References
1–4, 7–12, 15–20, 23, 24	103	Sutardja ¹

¹ With respect to the grounds based on “Sutardja,” MB USA alleges Sutardja US PGPUB 2008/0288748 A1 (referred to as Sutardja ’748) incorporates by reference the disclosure of Sutardja US PGPUB 2008/0288748 A1 (referred to as Sutardja ’785) in its entirety and refers to the combined disclosures of Sutardja ’748 and Sutardja ’785 as “combined Sutardja” or just “Sutardja,”

5, 6, 13, 14, 21, 22	103	Sutardja, Rychlik ²
7, 15, 23	103	Sutardja, Carmack ³
1–4, 7–12, 15–20, 23, 24	103	Mathieson, ⁴ Sutardja
5, 6, 13, 14, 21, 22	103	Mathieson, Sutardja, Rychlik

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

in support of its asserted “single reference obviousness ground.” Pet. at 5 n.2.

² US PG PUB 2011/0145615 A1.

³ US PG PUB 2009/0309243 A1.

⁴ US PG PUB 2011/0213950 A1. MB USA contends that Mathieson incorporates by reference the disclosure of Carmack in its entirety. Pet. 5 n.5.

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. Sutardja Fails to Teach or Suggest the Invention Recited in the Challenged Independent Claims

MB USA's first set of challenges are based on the teachings of Sutardja. Pet. at 21 *et seq.* As noted above, MB USA alleges that "Sutardja" encompasses the disclosures of both Sutardja '748 and Sutardja '785 combined. Pet. at 5 n.2. Accordingly, in this Patent Owner Response, references to and arguments concerning "Sutardja" are directed to Sutardja '748 incorporating the disclosure of Sutardja '785.

1. Sutardja '748

Sutardja '748 is a reference that was considered by the Examiner during prosecution of the application that became the '080 patent. Ex. 1004 at 84. So too was the parent Sutardja application, U.S. PGPUB 2008/0263324 ("Sutardja '324") (Ex. 2001), cited by the Examiner., *id.* at 78, and even used in connection with rejections of the claims of the then-pending '080 patent. *Id.* at 070-072. This is significant because Sutardja '324 contains the identical disclosure as Sutardja '748, now relied upon by MB USA.⁵

⁵ Under Petitioner's logic, the disclosure of Sutardja'785 was considered during prosecution of the '080 patent. Petitioner asserts that "Sutardja '748

Sutardja '748 relates to “dynamically switching cores of multi-core processing systems of mobile computing devices.” Ex. 1007 at [0002]. In particular, Sutardja '748 describes so-called “core morphing,” in which “cores are dynamically enabled (i.e., activated) or disabled (i.e., deactivated) based on the system load,” *id.* at [0218], in the context of a multi-core

incorporates by reference the disclosure of Sutardja '785 and that therefore it is proper to treat “Sutardja combined” as a single reference. Pet. at 5 n.4 Indeed, Petitioner repeatedly and throughout the Petition treats the two Sutardja references as a single reference because of this alleged incorporation. *See, e.g.*, Pet. at 5 (identifying Ground 1 as based on “Sutardja alone (Ex-1007, incorporating Ex-1008”); 21 et seq. (applying the combined Sutardja (Ex. 1007, Ex. 1008) references as one vis-à-vis the challenged claims. Since Sutardja '324 is identical to Sutardja'785 and likewise incorporates by reference Sutardja'785, see Ex. 2001 at p. 1 (63) (identifying Sutardja as a continuation-in-part of the application that became Sutardja '785), then by considering the disclosure of Sutardja '324, the Examiner also considered the disclosure of Sutardja '785 in connection with the prosecution of the '080 patent.

processing system with a low-speed, low-power (LP) core and a high-speed, high-power (HP) core. *Id.* at [0212]. When one core is active, other cores may be disabled (i.e., deactivated), for example, by being put in a standby mode or completely shut down, to save power. *Id.* at [0218]. In the HP mode, the HP core is in the active state and processes threads. The LP core may also operate during the HP mode. In other words, the LP core may be in the active state during all or part of the HP mode. *Id.* at [0223]. “When applications demand still higher performance than that provided by one HP core, multiple HP cores may be used.” *Id.* at [0219].

2. Sutardja ’785

Sutardja ’785 is the grandparent (through a continuation and continuation in part) of the application for Sutardja ’748. *See* Ex. 1007 at p. 1 (63) (identifying U.S. Patent Application No. 12/215,760, as a continuation of U.S. Patent Application No. 12/145,160, which is a continuation-in-part of U.S. Patent Application No. 11/523,996, Sutardja ’785). Sutardja ’785 describes a system on chip (“SOC”) that includes first and second processors, which each have active and inactive states, with the second processor consuming less power when operating in its active state than the first processor when operating in its active state. Ex. 1008 at [0009].

The SOC also includes a control module that communicates with the first and second processors and selectively transfers threads⁶ from the first processor to the second processor and selects the inactive state of the first processor. *Id.*

3. Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as claimed

To find a claim unpatentable for obviousness, even “single reference obviousness” as advanced by MB USA in its petition, *see* Pet. at 5 n.4, it is necessary to show that “the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art.” *Takeda Chem. Indus., Ltd. v. Alphapharm Pty, Ltd.*, 492 F.3d 1350, 1354-1355 (Fed. Cir. 2007) (quoting 35 U.S.C. §103(a)). This necessarily requires, among other things, determining the scope and content of the prior art. *KSR*, 550 U.S. at 406-407. And, while the

⁶ In programming, a “thread” is “a process that is part of a larger process or program.” Ex. 2007 at 518; *and see* Ex. 2006 at ¶ 43.

obviousness analysis is flexible, what is not flexible is the need to consider all of the limitation of a claim and find, based on a reasoned analysis, they are present, before concluding the claim is obvious. *See, e.g., Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1362 (Fed. Cir. 2016) (“[R]eferences to ‘common sense’—whether to supply a motivation to combine or a missing limitation—cannot be used as a wholesale substitute for reasoned analysis and evidentiary support, especially when dealing with a limitation missing from the prior art references specified.”).

In the present case, and contrary to MB USA’s allegations, Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as required by all of the challenged claims. Ex. 2006 at ¶¶ 44-52. Nor does MB USA’s petition provide any reasoned analysis that explains why a person of ordinary skill in the art would have modified Sutardja to provide such an operating system. Accordingly, MB USA has failed to demonstrate that any of the challenged claims are obvious in view of Sutardja.

Claim 1 of the ’080 patent requires:

power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a

drop in demand below a threshold without disabling any of the second plurality of cores, wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.

Ex. 1001 at 7:62 – 8:3. Independent claims 9 and 17 each recite the same “power management hardware” and “operating system” requirements. *Id.* at 8:45-53 (claim 9); 9:33-41 (claim 17).

With respect to this feature of the claims, the petition argues, “Sutardja teaches that the transition from the HP mode to the LP mode is performed using hardware (the claimed power management hardware). [A] ‘core switching module’ is used to transition the system between the HP and LP modes. *The hardware controlled by the core switching module corresponds to the claimed power management hardware*[, which] Sutardja explains [] ‘may be integrated with [a] hypervisor module 213,’ as illustrated in Figure 5.” Pet. at 32-33 (emphasis added and citations removed). The petition also provides a highlighted and annotated version of Fig. 5 from Sutardja ’748:




FIG. 5

Since both Sutardja and the petition discuss a hypervisor, it is helpful to review what a hypervisor is. In its most general form, a hypervisor is a component that abstracts one computer resource from another resource, allowing that other resource to operate in a self-contained fashion. Often, the hypervisor is a software component that abstracts underlying hardware (e.g., processors, memory, etc.) on which the hypervisor runs from other software resources. These are known as Type 1 or “bare metal” hypervisors. In other cases, hypervisors may be “hosted” by other operating systems that are running on underlying hardware. These Type 2 hypervisors run on top of the

“host” operating system, abstracting other software resources from that host.

Ex. 2006 at ¶¶ 46-47; *and see* Ex. 2002 at 19:25 – 21:14.

The petition notes that Sutardja describes both kinds of hypervisors: “Sutardja teaches that the hypervisor may ‘operate[] above the kernel’ or it ‘may operate below the kernel.’ Because the kernel is the lowest level of the operating system, a POSITA would have understood that a hypervisor that ‘operate[s] below the kernel’ is installed on underlying hardware.” Pet. at 33 (citations omitted). Therefore, relying on a scenario in which a hypervisor is installed below the kernel, the petition alleges, “[a] POSITA would have understood that *the hardware on which the disclosed hypervisor is installed corresponds to the claimed power management hardware.*” Pet. at 33-34 (citations omitted).

According to the petition then, in Sutardja it is both the hardware controlled by the core switching module and the hardware on which the disclosed hypervisor is installed that correspond to the power management hardware recited in the challenged claims. *Id.* at 32-34.

The petition goes on to allege, “Sutardja discloses an operating system executing on the multi-core processor.” *Id.* at 34. According to the petition, Sutardja’s Fig. 5 shows, “a functional block diagram,” and a POSITA would understand that the “PMS Module [shown in Fig. 5] is executed by the

kernel module (part of the Linux OS) to control the functioning of the multi-core processor.” *Id.* at 35 (Citations omitted). The petition then concludes, “Sutardja further discloses that *the operating system (kernel module and PMS module) controls the power management hardware (core switching module) based on the demand*: ‘the PMS module 242, and/or the kernel module 220 ... generate[s] a system call to a ... core switching module 212 to trigger a transition.’” *Id.* (emphasis added).

This argument is deeply flawed and is fatal to Petitioner’s case for obviousness of the challenged claims. Originally, the power management hardware was alleged to be not the core switching module but instead *the hardware controlled by the core switching module* and also *the hardware on which the disclosed hypervisor is installed*. *Id.* at 32-34. But, when seeking to map the operating system limitation, the petition backtracks and alleges that the power management hardware is only *the core switching module*. This is not merely an inconsistency. It is an admission that Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as claimed.

First, Sutardja’s core switching module is not “power management hardware.” In fact, it is not “hardware” at all. Ex. 2006 at ¶ 45. According to

Sutardja, the core switching module is software: “Depending on whether the core switching module 212 is integrated in the hypervisor module 213, the hypervisor module 213 or the core switching module 212 is loaded into the high-speed memory 210 from RAM or other storage after bootup.” Ex. 1007 at [0256]. Hardware cannot be “loaded into high speed memory from RAM or another storage device.” Therefore, the core switching module, whether it is integrated in the hypervisor module or is a stand-alone module is not hardware; it is software. Ex. 2006 at ¶¶ 45, 50. Consequently, the petition’s argument that Sutardja discloses an “*operating system (kernel module and PMS module) [that] controls the power management hardware (core switching module) based on the demand,*” fails.

Second, if one adopts the earlier allegation in the petition that the “power management hardware” recited in the claims is Sutardja’s hardware controlled by the core switching module and the hardware on which the hypervisor is installed, then, in keeping with the petition’s argument that the hypervisor “operates below the kernel,” it cannot be that the operating system controls this hardware. Ex. 2006 at ¶ 52. Consider that, a hypervisor operating “below the kernel” is the element that controls the hardware. *Id.*; and see Ex. 1028 at 0003-0004 (A bare metal hypervisor, which is a hypervisor that runs directly on hardware, “executes commands against the

hardware.”); Ex. 2002 at 20:9-10. Even Sutardja recognizes that, a hypervisor operating below the kernel is the thing that is “manag[ing] transfers between the processors and suspend[ing] operation of one asymmetric processor while the other asymmetric processor operates.” Ex. 1007 at [0255]. Indeed, the whole point of such a hypervisor is to abstract the underlying hardware from an operating system kernel running “above the hypervisor.” Ex. 2006 at ¶ 46. Accordingly, if, as alleged in the petition the “power management hardware” is the hardware on which the hypervisor is installed, then the operating system of Sutardja does not control that hardware and the petition has failed to demonstrate that Sutardja discloses an operating system that controls the power management hardware based on the demand, as claimed. *Id.* at ¶ 51.

Third, to the extent MB USA may try to argue that the “power management hardware” need not be the hardware on which the hypervisor is installed in Sutardja and instead need only be the hardware controlled by Sutardja’s core switching module, that argument too would fail. Sutardja is clear that, the kernel and the PMS module do not “control” the core switching module or the hardware on which it runs. Instead, the kernel and the PMS module only make system calls to the core switching module: “When the LP core 200 reaches its highest operating speed and additional

system loading occurs, the core profile module 240, the PMS module 242, and/or the kernel module 220 disable interrupts and generate a system call to a hyper visor module 213 or a core switching module 212 to trigger a transition. *The core switching module 212 may transition (switch) the operation from the LP core 200 to the HP core 204 in response to the system call and return control to the kernel module 220 after the transition (switching) is complete.*” Ex. 1007 at [0253] (emphasis added). Thus, it is the core switching module, not the operating system (kernel and PMS module), that controls the hardware and makes the switch from one core to another. Ex. 2006 at ¶ 52; Ex. 2002 at 28:22 – 29:5.

Accordingly, as to the challenges to independent claims 1, 9, and 17 that depend on Sutardja as a primary reference, the petition fails to make out a prima facie case of obviousness, *CFMT, Inc. v. Yieldup Intern. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003); and, hence, fails to satisfy its obligation to “specify where each element of the claim is found in the prior art.” 37 C.F.R. § 42.104(b)(4).

B. Claims Depending from Claims 1, 9, and 17 are Not Obvious in View of Sutardja

Challenged claims 2-4, 7, and 8 depend from claim 1, claims 10-12, 15, and 16 depend from claim 9, and claims 18-20, 23, and 24 depend from

claim 17. Each respective dependent claim is not unpatentable in view of Sutardja for at least the same reasons as its respective independent claim. *Hartness Int'l. Inc. v. Simplimatic Engineering Co.*, 819 F.2d 1100 (Fed. Cir. 1987).

C. The Remaining Challenges that Rely on Sutardja as a Primary Reference Likewise Fail

Claims 5 and 6 depend from claim 1, claims 13 and 14 depend from claim 9, and claims 21 and 22 depend from claim 17. Each respective dependent claim is challenged as being unpatentable in view of the teachings of Sutardja and Rychlik. Pet. at 5. In these challenges, Petitioner relies on Sutardja for teaching the limitations of the respective independent claims concerning the power management hardware and the operating system. *See* Pet. at 27-35, 45-51. Rychlik is cited for teaching disabling cores according to workload (demand). *Id.* at 46-48. Even if true, however, and even if a person of ordinary skill in the art were to apply such teachings to Sutardja, as alleged in the petition, it would remain the case that a Sutardja-Rychlik combination would fail to teach or suggest all of the requirements of claims 1, 9, and 17 for the same reasons as given above. Therefore, dependent claims 5, 6, 13, 14, 21, and 22 are not unpatentable in view of Sutardja-

Rychlik for at least those same reasons. *Hartness Int'l.*, 819 F.2d 1100; *CFMT, Inc.*, 349 F.3d at 1342.

Claim 7 depends from claim 1, claim 14 depends from claim 9, and claim 23 depends from claim 17. Each respective dependent claim is challenged as being unpatentable in view of the teachings of Sutardja and Carmack. Pet. at 5. In these challenges, Petitioner relies on Sutardja for teaching the limitations of the respective independent claims concerning the power management hardware and the operating system. See Pet. at 27-35, 51-53. Carmack is cited for teaching when cores are operating with different power consumption properties, one set of cores may be operating at their peak frequency. *Id.* at 51-52. Even if true, however, and even if a person of ordinary skill in the art were to apply such teachings to Sutardja, as alleged in the petition, it would remain the case that a Sutardja-Carmack combination would fail to teach or suggest all of the requirements of claims 1, 9, and 17 for the same reasons as given above. Therefore, dependent claims 7, 14, and 23 are not unpatentable in view of Sutardja-Carmack for at least those same reasons. *Hartness Int'l.*, 819 F.2d 1100; *CFMT, Inc.*, 349 F.3d at 1342.

D. Mathieson-Sutardja Fails to Teach or Suggest the Invention Recited in the Challenged Independent Claims

MB USA's challenges based on the combined teachings of Mathieson and Sutardja also fail to demonstrate that the challenged claims are unpatentable. In this set of Mathieson-Sutardja challenges, MB USA relies on the teachings of Mathieson⁷ for teaching the power management hardware and operating system limitations of independent claims 1, 9, and 17. Pet. at 59-67, 79-80. Sutardja's teachings are cited for providing "implementation details," to the extent a person of ordinary skill in the art would have required them. *Id.* at 53.

1. Overview of Mathieson

Mathieson describes a "technique for reducing the power consumption required to execute processing operations" using a processing complex such as "a CPU or a GPU" that includes "a first set of cores comprising one or


⁷ Mathieson is alleged to incorporate by reference the disclosure of Carmack. See Pet. at xxx. The petition separately references the Carmack teachings in connection with the power management hardware limitations, Pet. at 63-64. Therefore, this argument is addressed separately in this Patent Owner Response.

more fast cores and a second set of cores comprising one or more slow cores.” Ex. 1005 at Abst. The technique disclosed in Mathieson determines a “processing mode” based on “one or more of the workload characteristics performance characteristics ... power characteristics ... and operating conditions of the first and second sets of cores ... [and] a controller causes the processing operations to be executed by either the first set of cores or the second set of cores.” *Id.*

More particularly, the system disclosed by Mathieson includes a “central processing unit (CPU) 102,” which includes “one or more ‘fast’ cores 130 and one or more ‘shadow’ or slow cores 140.” *Id.* at [0020].

Mathieson’s controller (which can be included within the CPU, e.g., as software executed by the CPU) is configured to

evaluate a “least workload associated with one or more operations to be executed by the CPU 102.” *Id.* at [0030]. Based on the evaluated workload, the controller “is able to configure the CPU 102 to operate in a first mode of operation or a second mode of operation,” where in the first mode of operation the first set of cores is enabled and the second set of cores is disabled and in the second mode of operation the first set of cores is disabled



and the second set of cores is enabled. *Id.* “The transition from the first mode to the second mode, and vice versa, can be done transparently to high level software, such as the operating system.” *Id.* at [0043].

In discussing whether the first and the second cores can process tasks simultaneously, Mathieson states that any simultaneous operation of a first and second set of cores is limited to the “transition period when transitioning between the first mode and second mode, or vice versa” and notes that this potential for simultaneous operation is a net negative on the CPU’s power efficiency because “operating in both the first and second modes simultaneously may result in lower overall power efficiency.” *Id.* at [0031].

Mathieson states that “the determination of which type of cores [e.g., the first set or second set of cores] to use when executing the processing operations may be based on operating frequency, as shown in Fig. 6.” *Id.* at [0075]. When referring to Fig. 6 (reproduced below), power consumption (axis 604) is shown as a function of operating frequency (axis 602) for both the “fast” cores (line 606) and “slow” cores (line 608).




FIG. 6

The controller determines whether to execute processes on the “fast” or “slow” cores by referencing a “particular operating frequency threshold 610” that signifies the operating frequency at which the “slow” cores no longer consume less power than the “fast” cores. *Id.* at [0074]. The controller then executes the tasks on the set of cores that would consume less power at that frequency. *See id.* at [0074]-[0075].

2. Mathieson-Sutardja does not teach or suggest an operating system executing on a multi-core processor to monitor a demand for the multi-core processor and control power management hardware based on the demand, as claimed

As noted above, the challenged independent claims each recite an operating system executing on a multi-core processor to monitor a demand

for the multi-core processor and control power management hardware based on the demand. Ex. 1001 at 7:62 – 8:3 (claim 1); 8:45-53 (claim 9); 9:33-41 (claim 17). The petition alleges, that Mathieson describes a controller within a CPU that monitors demand. Pet. at 65. According to the petition, because the controller is implemented in software, is executed by the CPU 102, and both manages processing operations executing on CPU 102 and selects the underlying hardware (e.g., cores 210 and/or 220) on which those operations are executed, a POSITA would recognize that the controller performs the functions of an operating system, and thus would view controller 240 as being part of an operating system executing on CPU 102. *Id.* at 66.

This contention is demonstrably false. Mathieson states,

When execution of the processing operations switches from the first set of cores to the second set of cores, in some embodiments, the controller 240 is configured to transfer the processor state from the first set of cores to the second set of cores. In one embodiment, the controller saves the processor state to the shared resource 230, triggers a hardware mechanism that stops and powers off the first set of cores 210, and boots the second set of cores 220. The second set of cores 220 then restores the processor state from the shared resource 230 and continues operation at the lower speed associated with the second set of cores 220. In other embodiments,

the processing state may be stored in any memory unit when transferring execution of the operations between the two sets of cores. In still further embodiments, the processing state may be directly transferred to the other set of cores via a dedicated bus, where the processing state is not stored in any memory unit with switching between the two sets of cores. ***The transition from the first mode to the second mode, and vice versa, can be done transparently to high level software, such as the operating system.***

Ex. 1005 at [0043] (emphasis added). A person of ordinary skill in the art reading Mathieson would appreciate that if the transition between modes, which involves the controller transferring processor state from one set of cores to the other set of cores, is to be done *transparently* to the operating system, then it is not the operating system that is effecting the transition. Ex. 2006 at ¶¶ 58-61. If it were the operating system that was effecting the transition, then that transition would not be transparent to the operating system. *Id.* at ¶ 58. Dr. Brogioli explains that, “In the context of computing systems, a process is ‘transparent’ or is done ‘transparently’ to the operating system if the operating system is not aware of the process being performed.” *Id.* A person of ordinary skill in the art would further understand that for such transparency to exist, the transition between modes must be provided

by a mechanism at a layer *below the operating system*. *Id.* at ¶ 59. *and see* Ex. 2002 at 31:16 – 32:1 (Dr. Horst indicating that to be done transparently to the operating system, the process must take place below the operating system). Indeed, Mathieson itself confirms this understanding by emphasizing that the transition is done transparently to *higher level software*, such as the operating system. Ex. 1005 at [0043].

Thus, a person of ordinary skill in the art would understand that, contrary to the allegation in the petition, the controller does *not* perform the functions of an operating system in effecting the transition between modes. Instead, the controller operates at a level below the operating system in order to do so (and therefore allows the transition to be done transparently to the operating system). Ex. 2006 at ¶ 60. Accordingly, and as Petitioner admits it is Mathieson's controller that controls power management hardware based on demand, Pet. at 66-67, in Mathieson there is no teaching or suggestion of an operating system that monitors demand for the multi-core processor and controls the power management hardware based on the demand, as required by the challenged claims. Ex. 2006 at ¶ 61.

Consequently, the challenges to independent claims 1, 9, and 17 that depend on Mathieson-Sutardja fail to make out a prima facie case of obviousness, *CFMT, Inc.*, 349 F.3d at 1342; and, hence, fail to satisfy

Petitioner's obligation to "specify where each element of the claim is found in the prior art." 37 C.F.R. § 42.104(b)(4).

3. Mathieson-Sutardja does not teach or suggest a first plurality of cores and a second plurality of cores that support a same instruction set, wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores, as claimed

In addition to the above, the challenged independent claims require that "for a same applied operating frequency and supply voltage," "the second plurality of cores consume less power . . . than the first plurality of cores." Ex. 1001 at 7:57-61 (claim 1); 8:39-44 (claim 9); 9:27-32)claim 17). The petition refers to an annotated version of Mathieson's Fig. 6 (reproduced below) and argues, "the power consumed by Mathieson's second set of cores (blue) is less than the power consumed by Mathieson's first set of cores (red) for *the same operating frequency* (e.g., "F₀")." Pet. at 58 (emphasis in original).




FIG. 6

The challenged claims require more, however. It is not merely for the same operating frequency, but for the same applied operating frequency *and supply voltage* that the second plurality of cores consume less power than the first plurality of cores in the challenged claims. Mathieson’s Fig 6 does *not* support an argument that the second plurality of cores consume less power, for a same *applied supply voltage*. Ex. 2006 at ¶¶ 62-63. As Dr. Brogioli explains, Mathieson’s Fig. 6 appears to be characterizing trends for various groups of cores, and the figure shows no curves that represent power consumed as a function of applied supply voltage. *Id.*

Recognizing the deficiency the petition goes on to allege, “Mathieson explains that for ‘particular operating characteristics,’ including ‘operating

frequency [and] voltage,’ either an operating mode in which the first set of cores is enabled or an operating mode in which the second set of cores is enabled may be ‘more optimal,’ with ‘the CPU 102 [] us[ing] less power to operate in [that] mode.’” Pet. at 58-59 (alterations in original). But this characterization of Mathieson is misleading. In fact, in the cited passage from which the purported quotation is taken Mathieson actually states,

Based on the measured and/or estimated characteristics, one or more fuses may be set on the CPU 102 to characterize the performance data and/or the power data of the CPU 102 based on various characteristics, such as operating frequency, voltage, temperature, throughput, and the like. In some embodiments, the one or more fuses may comprise the data 214 and 224 shown in FIG. 2. Accordingly, the controller 240 may be configured to read the data 214, 224 and determine which mode of operation is most optimal based on the particular operating characteristics at a particular time.

Ex. 1005 at [0037]. A person of ordinary skill in the art would understand that the “fuses” Mathieson is referring to are elements used to store configuration information such as operating frequencies, or perhaps voltage settings. Ex. 2006 at ¶¶ 64-65. They are a form of non-volatile storage. *Id.*; and see Ex. 2003. Thus, the referenced passage in the petition does **not** tie an “optimal” mode of operation to reduced power consumption for a given

applied operating frequency and supply voltage. At best, what determines an “optimal” mode of operation is a set of unspecified “operating characteristics.”

The petition then alleges, “Mathieson explains that the difference in power consumption between cores 210 and 220, for the same operating characteristics (e.g., frequency and voltage), is due to transistor differences between the cores.” Pet. at 59. This characterization is also misleading. The actual referenced passage from Mathieson reads,

According to various embodiments, the cores 212 included in the first set of cores 210 comprise "fast" cores and the cores 222 included in the second set of cores 220 comprise "slow" cores. For example, the cores 212 may be manufactured using faster transistors that have significant static leakage. In some embodiments, when the computing needs and/or workload of the first set of cores 210 are lowered, then the clock speed is lowered to reduce power. The static leakage is not a significant issue at the high clock speeds required for peak performance. However, at slower clock speeds, the static leakage of the fast transistors can dominate the overall power consumption. According to various embodiments, the first set of cores includes N cores and the second set of cores includes M cores. In one embodiment, N is not equal to M. In other embodiments,

N is equal to M. In some embodiments, the first set of cores 210 may include multiple cores, e.g., four cores, and the second set of cores 220 may include a single core 222. In other embodiments, the first set of cores 210 may include a single core and/or the second set of cores 220 may include multiple cores.

Ex. 1005 at [0028]. Notably, in this passage Mathieson is **not** comparing power consumption of fast vs. slow cores for a same applied operating frequency and supply voltage. Rather, Mathieson is discussing power consumption of cores **only on the basis of operating frequency**. Nothing is said concerning the applied supply voltage. Indeed, as Dr. Brogioli explains, in this passage Mathieson, is just describing characteristics of a given CPU core, and not even **comparing** power consumption of those cores. Ex. 2006 at ¶ 66. For instance, the statement “static leakage current is not a significant issue at the higher clock speeds required for peak performance. However, at slower clock speeds the static leakage of the fast transistors can dominate the overall power consumption” is making a general observation concerning static vs. dynamic power consumption / leakage current for a given device itself. *Id.*

Hence, the petition fails to demonstrate that Mathieson-Sutardja, teaches or suggests, “wherein the second plurality of cores consume less

power, for a same applied operating frequency and supply voltage, than the first plurality of cores.” At best Mathieson, and hence a Mathieson-Sutardja combination, would be understood as describing a second plurality of cores that consume less power for a same applied operating frequency than a first plurality of cores. But this is insufficient to meet the requirements of the challenged claims. Accordingly, independent claims 1, 9, and 17 are not unpatentable in view of Mathieson-Sutardja. *CFMT, Inc.*, 349 F.3d at 1342.

4. The Challenged Dependent Claims are Also Not unpatentable in View of Mathieson-Sutardja

Challenged claims 2-4, 7, and 8 depend from claim 1, claims 10-12, 15, and 16 depend from claim 9, and claims 18-20, 23, and 24 depend from claim 17. Each respective dependent claim is not unpatentable in view of Mathieson-Sutardja for at least the same reasons as its respective independent claim. *Hartness Int’l.*, 819 F.2d 1100.

E. The Remaining Challenges that Rely on Mathieson-Sutardja Likewise Fail

Claims 5 and 6 depend from claim 1, claims 13 and 14 depend from claim 9, and claims 21 and 22 depend from claim 17. Each respective dependent claim is challenged as being unpatentable in view of the teachings

of Mathieson-Sutardja and Rychlik. Pet. at 6. In these challenges, Petitioner relies on Mathieson-Sutardja for teaching the limitations of the respective independent claims concerning the power management hardware and the operating system. *See* Pet. at 59-67. Rychlik is cited for teaching disabling cores according to workload (demand). *Id.* at 82. Even if true, however, and even if a person of ordinary skill in the art were to apply such teachings to Sutardja, as alleged in the petition, it would remain the case that a Mathieson-Sutardja & Rychlik combination would fail to teach or suggest all of the requirements of claims 1, 9, and 17 for the same reasons as given above. Therefore, dependent claims 5, 6, 13, 14, 21, and 22 are not unpatentable in view of Mathieson-Sutardja & Rychlik for at least those same reasons. *Hartness Int'l.*, 819 F.2d 1100; *CFMT, Inc.*, 349 F.3d at 1342.

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

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

Respectfully submitted,

/Tarek N. Fahmi/

Tarek N. Fahmi
Reg. No. 41,402

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 8,977 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. 16.84, 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
Damon M. Lewis
Jiaxing (Kyle) Xu
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
m
nicholas.rotz@hoganlovells.com
damon.lewis@hoganlovells.com
kyle.xu@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