

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

QUALCOMM INCORPORATED,
Petitioner,

v.

DAEDALUS PRIME LLC,
Patent Owner.

Case No. IPR2023-00567
U.S. Patent No. 10,049,080

**PATENT OWNER'S PRELIMINARY RESPONSE
PURSUANT TO 37 C.F.R. § 42.107(A)**

TABLE OF CONTENTS

I.	Introduction.....	5
II.	Technology Background.....	6
A.	Prior Art Failures.....	7
III.	Overview of the '080 Patent.....	9
IV.	Discretionary Factors.....	10
A.	<i>Becton, Dickinson</i> factor: “the similarities and material differences between the asserted art and the prior art involved during examination”	11
B.	<i>Becton, Dickinson</i> factor: “the cumulative nature of the asserted art and the prior art evaluated during examination”	16
C.	<i>Becton, Dickinson</i> factor: “the extent to which the asserted art was evaluated during examination, including whether the prior art was the basis for rejection”	19
D.	<i>Becton, Dickinson</i> factor: “whether petitioner has pointed out sufficiently how the examiner erred in its evaluation of the asserted prior art”	20
V.	Overview of Select Cited Art	22
A.	Overview of Sutardja '748	22
B.	Overview of Sutardja '785	23
C.	Overview of Mathieson.....	24
VI.	Ground 1 Fails	26
A.	Sutardja '748 and Sutardja '785 Are Not a Single Reference, and Petitioner Has Not Demonstrated They Would be Combined by a POSITA	26
B.	The Sutardja References do not render obvious Claim 1(a)(i)	29
C.	The Sutardja References do not render obvious Claim 1(a)(ii)	32
D.	The Sutardja References do not render obvious Claim 1(b)(i)	34
E.	The Sutardja References do not render obvious Claim 1(b)(ii).....	38
F.	The Sutardja References do not render obvious Claims 9 and 17	40
VII.	Ground 4 Fails	41

A.	A POSITA Would Not Have Been Motivated to Combine Mathieson with the Sutardja References.....	41
B.	Mathieson, Alone or in View of the Sutardja References, Does Not Render Obvious Claims 1, 3-4, and 7-8.....	43
VIII.	Conclusion	53

TABLE OF AUTHORITIES

Page(s)

Cases

<i>KSR Int’l Co. v. Teleflex Inc.</i> , 550 U.S. 398 (2007).....	31
<i>Callaway Golf Co. v. Acushnet Co.</i> , 576 F.3d 1331, 1346 (Fed. Cir. 2009)	27
<i>Personal Web Techs., LLC v. Apple, Inc.</i> , 848 F.3d 987, 993 (Fed. Cir. 2017)	42
<i>Securus Techs., Inc. v. Global Tel*Link Corp.</i> , 701 F. App’x 971, 977 (Fed. Cir. 2017)	42
<i>Stone Basket Innovations, LLC v. Cook Med. LLC</i> , 892 F.3d 1175, 1179 (Fed. Cir. 2018)	12
<i>Advanced Bionics v. Med-El Elektro. GmbH</i> IPR2019-01469, Paper 6.....	10-12
<i>Johns Manville Corp. v. Knauf Insulation, Inc.</i> , IPR2018-00827, Paper 9.....	42

Statutes

35 U.S.C. § 325(d)	5, 10
--------------------------	-------

TABLE OF EXHIBITS

Exhibit	Description
2001	U.S. Pat. Pub. No. 2008/0263324 to Sutardja et al (“Sutardja ’324”)
2002	U.S. Pat. Pub. No. 2006/0095807 to Grochowski et al. (“Grochowski”)
2003	U.S. Patent No. 8,984,523 to Vajda
2004	U.S. Patent No. 8,892,931 to Kruglick
2005	U.S. Patent No. 9,086,883 to Thomson
2006	U.S. Patent Publication No. 2009/0222654 to Hum

I. Introduction

The '080 Patent improved how power is managed in computing systems by disclosing use of two sets of CPU cores that can consume different amounts of power and offer different performance capabilities. Whereas previously, multi-core processors used identical cores that were not intended for improved power savings or capabilities, the inventors of the '080 Patent devised of a processor with a plurality of cores intended for each role. Even previous use of non-identical cores was insufficient, because they did not support the same instruction set, increasing the difficulty of moving computational processing between them. Using different cores that support the same instruction set, the multi-core processor of the '080 Patent can operate at higher performance using the performance cores, or at higher efficiency using the low power cores, depending on computing demand. This greatly increases the power efficiency of the '080 Patent's processor at all levels of demand.

The Petition asserts that the challenged claims of the '080 patent are invalid as obvious, but fails because it: (1) improperly relies on the *ipse dixit* of a declarant who merely parrots attorney argument; (2) should be denied under 35 U.S.C. § 325(d) for relying on substantially the same art and arguments previously presented to the Office; and (3) fails to demonstrate that the content of the cited art renders obvious any of the challenged claims.

The cited art disclosed by Qualcomm Incorporated (“Petitioner”¹) fails to render obvious the claims of the ’080 Patent, and the Petition should be denied in its entirety².

II. Technology Background

Modern computing devices can make use of multicore processors that combine a plurality of processor cores on a single semiconductor die. Ex. 1001, 1:22-25. Multiple processor cores can be used, for example, to provide separate functional units, which can increase the amount of computations that can be performed at one time.

The amount of power consumed by a processor is strongly correlated with its performance capabilities. Ex. 1001, 2:26-29. Because decreasing the processing performance of the system corresponds to power savings, processor power

¹ As filed, the Petition identified Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc. (collectively “Samsung”) along with Qualcomm Incorporated as Petitioners. However, on June 2, 2023, Patent Owner and Petitioner Samsung jointly moved to terminate as to Samsung. Paper No. 10. On June 27, 2023, The Board granted the motion and terminated the proceedings as to Samsung, and noted that Qualcomm Incorporated would remain as Petitioner. Paper No. 11.

² Petitioner challenges independent claims 1, 9, and 17 as allegedly obvious only in Grounds 1 and 4. Thus, this Preliminary Response primarily addresses those Grounds; each of the challenged dependent claims is non-obvious at least because it depends from a non-obvious independent claim. Patent Owner reserves the right to address all arguments made in the Petition, including those made for other Grounds, in further briefing.

management typically scales up the performance of the system as the system's workload increases, and scales down the performance as workload decreases. Ex. 1001, 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. Ex. 1001, 2:30-34.

A. Prior Art Failures

As set forth in the '080 Patent, prior art approaches to power management did not provide for interchangeable use of different types of processor cores. For example, the '080 Patent identifies a need in the industry for a suitable multi-core processor having non-identical cores supporting a same instruction set, and a power management scheme for the same that enables and disables the cores based on demand. See Ex. 1001, 2:19-20 ("Computer system power consumption is becoming more and more of a concern."), 3:34-49 ("[P]rior art multi core processor power management schemes have been implemented on processors whose constituent cores are identical. ... In other approaches, the cores are not identical but are radically different. ... A problem with this approach, however, is that it is difficult for system software to adjust switch operation between processor cores having different instruction sets.").

Publications in the art noted a need for better power management, but failed to solve it. *See* Ex. 2003, 1:39-65 (“Amdahl’s law, even it’s [sic] version for chip multi-processors, will put a limit to the amount of performance increase that can be achieved even for highly parallelized applications.”). For example, some prior art approaches merely adjusted the frequency (also called “clock”) and voltage at which the cores operated to adjust the power consumption/processing capability balance, but such approaches presented core management difficulties. *See* Ex. 1001, 2:30-34 (“A typical way to scale processing performance and power consumption with workload is to enable/disable entire cores and raise/lower their supply voltages and operating frequencies in response to system workload”); Ex. 2004, 1:31–37 (“Attempts to apply [Dynamic Voltage and Frequency Scaling] to individual cores within a multicore processor may encounter a number of difficulties as the number of cores increases.”); Ex. 2005 at 1:42-56 (“Conventional [Dynamic Clock and Voltage Scaling] solutions [for multi-core processors] exhibit a number of performance problems, and implementing an effective DCVS method ... for each core ... is an important and challenging design criterion.”). Other prior art approaches attempted to use low power states, but this required time and energy to raise the cores or processor out of the low power state when needed. *See* Ex. 2006 at ¶ 2 (“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. Furthermore, not all processes and tasks that are run on a processor require the full processing throughput of the processor.”).

III. Overview of the '080 Patent

The '080 Patent describes a multicore processor that can make use of two sets of CPU cores that can consume different amounts of power and offer different performance capabilities. The CPU cores are thus referred to as “asymmetric.” The '080 Patent describes that its invention provides a way for a processor to make use of these two sets of asymmetric CPU cores 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.” Ex. 1001 at 2:19-29, 3:50-62, 4:20-46.

In particular, the '080 Patent describes that a multi-core processor according to its invention can include power management hardware to enable and disable the cores in different configurations, and to change the frequency and voltage at which they operate, based on changes in computational demand. Ex. 1001 at 4:54-6:24. For example, in one embodiment the power management hardware may disable all of a first set of cores but keep all of a second set of cores active, when there is a drop in demand below a threshold. Ex. 1001 at 4:65-5:6, 5:25-28.

IV. Discretionary Factors

It is well established that the Board has discretion to deny a petition where the art cited therein is merely cumulative of art previously considered during examination of the challenged patent. 35 U.S.C. § 325(d) (“the Director may take into account whether, and reject the petition or request because, the same or substantially the same prior art or arguments previously were presented to the Office”).

[U]nder § 325(d), the Board uses the following two-part framework: (1) whether the same or substantially the same art previously was presented to the Office or whether the same or substantially the same arguments previously were presented to the Office; and (2) if either condition of first part of the framework is satisfied, whether the petitioner has demonstrated that the Office erred in a manner material to the patentability of challenged claims.

Advanced Bionics v. Med-El Elektro. GmbH, IPR2019-01469, Paper 6, at 8 (PTAB Feb. 13, 2020) (precedential). In considering whether the same or substantially the same art or arguments were previously presented to the Office, “the *Becton, Dickinson* factors provide useful insight into how to apply the framework under 35 U.S.C. § 325(d).” *Id.* at 9; *id.* at n.10 (identifying the *Becton, Dickinson* factors).

Here, Petitioner's challenges all rely on the Sutardja references (Ex. 1007: Sutardja '748, and Ex. 1008: Sutardja '785), which were themselves considered during prosecution of the '080 Patent, or Mathieson, which is cumulative of at least the Grochowski reference that was considered during prosecution. The Board should thus exercise its discretion to deny the Petition in its entirety.³

A. *Becton, Dickinson* factor: “the similarities and material differences between the asserted art and the prior art involved during examination”

Each of Petitioner's cited references is materially similar to the art involved during examination, making this factor favor denial.

1. Sutardja '748 (Ex. 1007) was considered

The first of the asserted Sutardja references, U.S. Patent Pub. No. 2008/0288748 to Sutardja et al. (“Sutardja '748,” Ex. 1007) was itself considered during prosecution of the '080 Patent, making it perfectly similar to “the prior art involved during examination.” See IPR2019-01469 at n. 10 (factor (a)). For example, Sutardja '748 (Ex. 1007) is identified on the face of the '080 Patent as one

³ Petitioner only challenges independent claims 1, 9, and 17 as either allegedly obvious over the Sutardja references (Ground 1), or Mathieson in view of the Sutardja references (Ground 4). Because Petitioner does not assert that any of the other cited art affects its contentions with respect to the independent claims, a finding that Mathieson and the Sutardja references are cumulative of art considered during prosecution could dispense of the entire Petition.

of the “References Cited.” Ex. 1001 at References Cited. *See Stone Basket Innovations, LLC v. Cook Med. LLC*, 892 F.3d 1175, 1179 (Fed. Cir. 2018) (“when prior art is listed on the face of a patent, the examiner is presumed to have considered it”).

Additionally, the parent application of Sutardja ’748 (Ex. 1007)—Sutardja ’324 (Ex. 2001)—contains an identical disclosure and was also considered during examination of the ’080 Patent. *See* Ex. 1007; Ex. 2001; Pet. at 8-9. Specifically, the Sutardja ’748 (Ex. 1007) reference cited by Petitioner here is a continuation of Sutardja ’324 (Ex. 2001), which was considered during examination of the ’080 Patent. Ex. 1007 at (63); Pet. at 9. Because it contains a disclosure identical to one considered during examination, Sutardja ’748 (Ex. 1007) is perfectly similar to “the prior art involved during examination.” *See* IPR2019-01469 at n. 10 (factor (a)).

2. Sutardja ’785 (Ex. 1008) was considered under Petitioner’s positions⁴

The second of the asserted Sutardja references, U.S. Patent Pub. No. 2007/0083785 to Sutardja (“Sutardja ’785,” Ex. 1008) was considered during prosecution under Petitioner’s own logic. Petitioner asserts that Sutardja ’785 (Ex.

⁴ Although Patent Owner disagrees with Petitioner’s assertion that the two Sutardja references should be treated as a single reference for the reasons discussed below, if they could be treated as such, then Sutardja ’785 was considered during examination.

1008) “was not before the examiner” during the prosecution of the ’080 Patent, but Sutardja ’785 (Ex. 1008) is named by—and by Petitioner’s own assertion purportedly incorporated by reference into—each of Sutardja ’748 (Ex. 1007) and Sutardja ’324 (Ex. 2001).

For example, Petitioner asserts that “Sutardja ’748 incorporates by reference the disclosure of Sutardja ’785 (Ex-1008) in its entirety.” Pet. at 17 (citing Ex-1007, ¶1). Petitioner also repeatedly and throughout the Petition treats the two Sutardja references as a single reference because of the alleged incorporation of Sutardja ’785 (Ex. 1008) into Sutardja ’748 (Ex. 1007). *See, e.g.*, Pet. at 4 (identifying Ground 1 as based on “Sutardja alone (Ex-1007, incorporating Ex-1008”); n.2 (“this ground relies on Sutardja ’748 (Ex-1007) incorporating by reference Sutardja ’785 (Ex-1008), ***as a single reference***⁵ obviousness ground, referred to herein as the ‘combined Sutardja’ or simply ‘Sutardja’”); 17 (providing overview of the combined “Sutardja (Ex. 1007, Ex. 1008)” and describing “[b]oth Sutardja ’748 and Sutardja ’785 (collectively ‘Sutardja’)”). Thus, under Petitioner’s logic, by virtue of having considered Sutardja ’748 (Ex. 1007), the Examiner also considered Sutardja ’785 (Ex. 1008).

⁵ Emphasis added throughout unless otherwise noted.

3. Mathieson is materially similar to art that was involved in examination

Mathieson is materially similar to Sutardja '748, which was involved in examination. For example, Petitioner admits that Mathieson is materially similar to Sutardja '748, which was involved in examination. Petitioner asserts: “like Mathieson and the '080 Patent, Sutardja ['748] teaches that the power consumption differences between cores are due to transistor differences,” Pet., p.17 (citing Ex. 1007); “like Mathieson and the '080 Patent, Sutardja ['748] teaches operating the multi-core processor in a high-power mode (with the high-power cores and, optionally, the low-power cores enabled) and a low-power mode (with only the low-power cores enabled), based on system loading,” Pet., p.17 (citing Ex. 1007). Additionally, Petitioner asserts in Ground 4 that Mathieson discloses or suggests the same limitations that it asserts Sutardja '748 discloses in Ground 1, for similar reasons. These alleged similarities Sutardja '748 and Mathieson confirm that Mathieson is materially examination.

Additionally, Mathieson is materially similar to Grochowski, which was involved in examination. Petitioner asserts that the “Mathieson Ground, presented herein [is not] the same, substantially the same, or cumulative of Grochowski, relied upon by the Examiner in rejecting the independent claims of the '080 Patent,” but provides only a single allegedly distinguishing feature of Mathieson:

the Examiner relied on Grochowski's teaching of "parallelism in the program" and situations in which "parallelism increased" as teaching the claimed "demand" ... Nor is Mathieson, in Ground 4, based on "parallelism as Grochowski is. Mathieson instead includes express teachings to consider "workloads" against "threshold" values to determine the claimed "demand." None of these teachings relating to workloads, system loading parameters, and thresholds is cumulative of the Grochowski's "parallelism" considerations as applied by the Examiner because they operate in substantially different ways than Grochowski.

Pet. at 9-10 (citations omitted) (referring to U.S. Pat. Pub. No. 2006/0095807 to Grochowski et al., Ex. 2002). However, Petitioner does not explain how Mathieson allegedly works differently than Grochowski, or how Mathieson allegedly presents different considerations than Grochowski did.

For example, Petitioner does not dispute that the examiner considered whether Grochowski's disclosure of increased parallelism corresponds to an increased "demand." Grochowski states:

the amount of parallelism may be the number of simultaneous threads supported. In other embodiments, other metrics may be used to express the amount of parallelism, such as the aggregate number of instructions

retired per second, or the number of branch instructions that may support speculative multi-threaded execution.

Ex. 2002 at [0024]. Nor did the applicant distinguish Grochowski on the ground that it described parallelism considerations. Therefore, Mathieson's disclosure is materially similar to the disclosure of Grochowski that was previously considered during examination.

B. *Becton, Dickinson* factor: “the cumulative nature of the asserted art and the prior art evaluated during examination”

Petitioner's cited art is also collectively cumulative of the prior art evaluated during examination, making this factor favor denial.

1. The Sutardja References are cumulative of Grochowski in view of Sutardja '324

During examination of the '080 Patent, the Examiner issued an office action based on Grochowski in view of Sutardja '324. *See* Ex. 1004 at 0067-0072. Petitioner asserts that the Sutardja references are not cumulative of Grochowski, but ignores that (1) several of Grochowski's disclosures mirror those that the Petition identifies in the Sutardja references, and (2) the prior art evaluated during examination included Grochowksi in combination with Sutardja '324. Pet. at 10.

For example, Petitioner asserts the following with respect to Grochowski:

the Examiner relied on Grochowski's teaching of “parallelism in the program” and situations in which

“parallelism increased” as teaching the claimed “demand.”

But Grochowski discloses **enabling** high-power (A) cores and **disabling** low-power (B) cores in response to decreasing “parallelism.” In contrast, Sutardja discloses demand in the form of “system loading,” and its multi-core processor exhibits the opposite behavior—high-power cores are **disabled** and low-power cores remain **enabled** in response to decreasing loading.

Pet. at 10 (citations omitted) (emphasis in original). However, Grochowski is not limited to disabling low power cores, the only thing allegedly distinguishing it from the Sutardja references. For example, as the Examiner cited, Grochowski explains that “[e]ach time the throttle module 210 makes the determination of the amount of parallelism in the program, it may initiate **powering the A cores and B cores up or down** using signal lines 224 through 266.” Ex. 2002 at [0034]; *see also* Ex. 1004 at 0068 (citing Grochowski at [0034]). Because the disclosures in the Sutardja references mirror those of Grochowski, they are cumulative of the art evaluated during examination.

Additionally, Grochowski was not considered in isolation during examination—it was specifically considered in combination with Sutardja ’324. *See* Ex. 1004 at 0067-0072. As explained above, Sutardja ’324 (Ex. 2001) contains the same disclosure as Sutardja ’748 (Ex. 1007) because it is the continuation-parent

thereof. And yet, Petitioner cites to Sutardja '748 (Ex. 1007) as allegedly disclosing something not considered during examination: “system loading” and high-power cores being disabled. Pet. at 10 (citing Ex. 1007). This is nonsensical; the same disclosure cannot now disclose something that it did not during examination.

2. Mathieson is cumulative of Grochowski in view of Sutardja '324

Petitioner asserts that Mathieson is not cumulative of Grochowski, but ignores that the prior art evaluated during examination included Grochowksi in combination with Sutardja '324. For example, Petitioner asserts that Mathieson is not “based on ‘parallelism’ as Grochowski is,” and that Mathieson’s disclosures “relating to workloads, system loading parameters, and thresholds” are non-cumulative of Grochowski. Pet. at 10-11. As to parallelism, Petitioner ignores that the art considered during examination—including Sutardja '324 (Ex. 2002)—was not exclusively directed to parallelism. As to disclosures of workloads, as explained above, Sutardja '324 (Ex. 2001) contains the same disclosure as Sutardja '748 (Ex. 1007) because it is the continuation-parent thereof. And yet, Petitioner cites to Sutardja '748 (Ex. 1007) as allegedly disclosing “a demand in the form of ‘system load’ and/or ‘power/performance demand.’” See Pet. at 26 (citing Ex. 1007). In fact, as explained above, Petitioner asserts that Mathieson is similar to Sutardja '748 (Ex. 1007): “like Mathieson and the '080 Patent, Sutardja ['748] teaches that the power

consumption differences between cores are due to transistor differences,” Pet., p.17 (citing Ex. 1007); “like Mathieson and the ’080 Patent, Sutardja [’748] teaches operating the multi-core processor in a high-power mode (with the high-power cores and, optionally, the low-power cores enabled) and a low-power mode (with only the low-power cores enabled), based on system loading,” Pet., p.17 (citing Ex. 1007). These alleged similarities between Sutardja ’748 and Mathieson confirm that Mathieson is merely cumulative of the collective prior art evaluated during examination.

C. *Becton, Dickinson* factor: “the extent to which the asserted art was evaluated during examination, including whether the prior art was the basis for rejection”

The disclosures of Sutardja references were evaluated during examination, including by being used as the basis for rejection, making this factor favor denial.

Specifically, Sutardja ’324 (Ex. 2001) formed the basis of an office action during examination of the ’080 Patent, and contains an identical disclosure to Sutardja ’748 (Ex. 1007). Ex. 1004 (’080 Patent FH Excerpts) at 0070-0072; Pet. at 9. This demonstrates that the disclosure of Sutardja ’748 (Ex. 1007) was thoroughly considered during examination.

Further, Sutardja ’748 (Ex. 1007) itself was considered during examination of a parent to the ’080 Patent. During examination of the ’278 patent, three separate office actions were issued relying on Sutardja ’748 (Ex. 1007). *See* Ex. 1010 (’278

patent FH Excerpts) at 0049-0053, 0110-0116, 0153-0159; Pet. at 10, n.5. Indeed, the '278 patent was examined by the same Examiner as the '080 Patent—Eric Coleman—demonstrating the consideration of that disclosure that was undertaken during examination. *See, e.g.*, Ex. 1010 at 0009; Ex. 1001 at cover page.

D. *Becton, Dickinson* factor: “whether petitioner has pointed out sufficiently how the examiner erred in its evaluation of the asserted prior art”

Petitioner has failed to demonstrate that the examiner erred in evaluating the asserted art, and this factor favors denial.

Petitioner’s only assertion of examiner error is as follows:

even if deemed presented, it was material error not to apply Sutardja '324 or '748 incorporating Sutardja '785 (Ex-1008) in a rejection because fully-incorporated Sutardja discloses and teaches the claims pending during prosecution. The incorporated Sutardja '785 discloses what the applicant claimed was missing from the prior art: the plurality of low power cores.

Pet., pp.11-12. However, Petitioner omits that arguments regarding failure to disclose the plurality of low power cores were ***not made during prosecution of the '080 patent at issue here***. Rather, Petitioner cites to Ex. 1010, which is the “Prosecution History of U.S. Patent No. 9,569,278.” *See* Pet., p.12 (citing to “supra n.5,” which cites Ex. 1010); Pet., p.x (identifying Ex. 1010). Petitioner has not compared the claims that were pending during evaluation of the '278 patent to those at issue here from the '080 patent—evaluation of prior art with respect to different

claims is inapposite here. For example, the challenged claims of the '080 patent and the pending claims of the '278 patent are reproduced here:

Pending claims of the '278 Patent	Challenged claim of the '080 Patent
A multi-core processor comprising a plurality of cores that support a same instruction set, wherein at least two of the plurality of cores consume less power, for a same applied operating frequency and supply voltage, than at least one other of the plurality of cores.	[1pre] A multi-core processor comprising: [1a] 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 [1b] 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.

Compare Ex. 1010 at 0125, 0135 *with* Ex. 1001 at cl. 1. In glossing over this distinction, Petitioner ignores that different arguments were made during evaluation of the '080 Patent. For example, the cited art—including Sutardja '748 (Ex. 1007)—was distinguished based on limitation [1b] of the '080 Patent, which is distinct from

arguments regarding failure to disclose the plurality of low power cores that Petitioner points to as alleged error. *See* Ex. 1004, pp.0054-0057.

V. Overview of Select Cited Art

A. Overview of Sutardja '748

Sutardja '748 is directed to a “core switching system” that includes a “mode switching module that receives a switch signal to switch operation between a first mode and a second mode.” Ex. 1007 at Abstract. During the first mode, “instructions associated with applications are executed by a first asymmetric core, and a second asymmetric core is inactive.” Ex. 1007 at Abstract. During the second mode, “the instructions are executed by the second asymmetric core, and the first asymmetric core is inactive.” Ex. 1007 at Abstract.

The multi-core processor disclosed in Sutardja '748 includes a high power (“HP”) core and a low power (“LP”) core in which “the multi-core processing system dynamically switches between the LP and HP cores depending on the type of applications being executed.” Ex. 1007 at Fig. 2, ¶ 212. Sutardja '748 uses what it calls “core morphing,” which “includes adaptive and dynamic switching” between the LP core and the HP core. Ex. 1007 at ¶ 216. This switching is done “securely and transparently to the OS and/or the application,” resulting in the OS and the applications being “unaware of which core is currently executing the processing load.” Ex. 1007 at ¶ 216.

For example, in connection with its Figure 5, Sutardja '748 describes that “when the LP core 200 reaches its highest operating speed and additional system loading occurs,” the system “generate[s] a system call to a hypervisor module 213 or core switching module 212 to trigger a transition.” Ex. 1007 at Fig. 5, ¶ 253. Once the Power Management System (“PMS”) module triggers the transition, the “core switching module 212 copies the state of the LP core 200, activates the HP core 204, and initializes the HP 204 core with the state of the LP core 200.” Ex. 1007 at ¶ 258. Subsequently, the core switching module deactivates the LP core, the HP core resumes executing the application, and the interrupts are enabled by the PMS module. *See* Ex. 1007 at ¶ 258. When the HP core(s) are executing tasks and system load decreases, the core switching module may transition operating to the LP core in a manner consistent with the transitioning of tasks from the LP core to the HP core. *See* Ex. 1007 at ¶ 259.

B. Overview of Sutardja '785

Sutardja '785 relates to “[a] system on chip (SOC) that includes first and second processors and a control module,” and which “has active and inactive states, wherein the second processor consumes less power when operating in the active state than the first processor operating in the active state.” Ex. 1008 at Abstract. Sutardja '785 describes a high-power “HP Processor” and a low-power “LP Processor.” *See, e.g.,* Ex. 1008 at Fig. 31A, [0208]-[0209].

Sutardja '785 describes that “[w]hen transitioning between the high-power and low-power modes, the high-power and low-power processors selectively transfer one or more program threads to each other.” Ex. 1008 at ¶ 207. According to Sutardja '785, the low-power processor and high-power processor “may temporarily operate in the active state at the same time when transitioning from the low-power mode to the high-power mode and when transitioning from the high-power mode to the low-power mode,” but “in the low-power state the low-power processor 1308 operates in the active state and the high-power processor 1304 is in the inactive state.” Ex. 1008 at ¶ 209.

C. Overview of Mathieson

U.S. Patent Pub. No. 2011/0213950 to Mathieson et al. (“Mathieson”) is directed to 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 more fast cores and a second set of cores comprising one or more slow cores.” Ex. 1005 at Abstract. 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” Ex. 1005 at Abstract.

Mathieson refers to a computer hardware system. Ex. 1005 at ¶ 2. The system disclosed by Mathieson comprises a “central processing unit (CPU) 102,” which includes “one or more ‘fast’ cores 130 and one or more ‘shadow’ or slow cores 140. Ex. 1005 at ¶ 20. Mathieson cautions against cores having the same functionality, noting “trade-offs in usage” for providing CPU cores 210 and 220 with “shared functionality.” Ex. 1005 at ¶ 25.

Mathieson’s controller is configured to evaluate a “least workload associated with one or more operations.” Ex. 1005 at ¶ 30. “Based on the evaluated workload, the controller is able to configure the CPU 102 to operate in a first mode ... or a second mode of operation,” where the first mode of operation enables the first set of cores and disables the second set of cores and the second mode of operation enables the second set of cores and disables the first set of cores. Ex. 1005 at ¶ 30. “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 ¶ 43.

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.” Ex. 1005 at ¶ 31.

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.” Ex. 1005 at ¶ 75. When referring to Fig. 6, power consumption (axis 604) is shown as a function of the operating frequency (axis 602) of the “fast” cores (line 606) and “slow” cores (line 608).

The controller 240 (see Mathieson at Fig. 2) 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. Ex. 1005 at ¶ 74. The controller then executes the tasks on the set of cores that would consume less power at that frequency. See Ex. 1005 at ¶¶ 74-75.

VI. Ground 1 Fails

The Petition asserts in Ground 1 that Claims 1-5 and 7 are rendered obvious by the Sutardja references.

A. Sutardja ’748 and Sutardja ’785 Are Not a Single Reference, and Petitioner Has Not Demonstrated They Would be Combined by a POSITA

Petitioner treats Sutardja ’748 (Ex. 1007) and Sutardja ’785 (Ex. 1008) as a single reference, which it refers to as “Sutardja.” Petitioner in many places does not

specify whether the discussion relates to Sutardja '748 (Ex. 1007) or Sutardja '785 (Ex. 1008). This makes the analysis of the proposed combination of the two Sutardja references insufficiently clear, which fails to show a reasonable likelihood of success. Indeed, there is no specific obviousness analysis in the Petition, for example by discussing combining teachings from any specific embodiments whether from Sutardja '748 (Ex. 1007) or Sutardja '785 (Ex. 1008). Petitioner has thus failed to meet its burden.

Petitioner asserts that Sutardja '748 incorporates by reference the disclosure of Sutardja '785, but has not demonstrated that the limited statement purporting to do so is sufficiently clear. "To incorporate matter by reference, a host document must contain language clearly identifying the subject matter which is incorporated and where it is to be found; a mere reference to another application, or patent, or publication is not an incorporation of anything therein." *Callaway Golf Co. v. Acushnet Co.*, 576 F.3d 1331, 1346 (Fed. Cir. 2009). Here, Sutardja '748 (Ex. 1007) does not identify with particularity what material from Sutardja '785 (Ex. 1008) it incorporates. *See* Ex. 1007, [0001]. For example, Sutardja '748 (Ex. 1007) does not identify any specific disclosure allegedly relevant to any of its embodiments or components. This is particularly problematic where, as here, the references are directed to two different subjects: high and low power **processors** and transferring sets of thread between them (Sutardja '785, Ex. 1007, at ¶ 9) vs. high and low power

cores and switching between the cores by copying states (Sutardja '748, Ex. 1008, at ¶¶ 212, 228).

Petitioner also asserts that “[t]o the extent that any relied upon teaching is argued to have not been fully incorporated from Ex-1008 into Ex-1007, it would have been obvious to combine such feature/teaching from Ex-1008 with Ex-1007, for the same reasons provided herein.” Pet., n.2. However, Petitioner does not explain any alleged motivation to combine the teachings of the two Sutardja references, despite asserting that it would have been obvious to combine certain teachings. Petitioner nowhere analyzes the traditional considerations of a motivation to combine analysis. For example, Petitioner does not allege that Sutardja '785 provides a solution to a problem identified in Sutardja '748. Nor does the proposed combination represent a substitution of any known elements from Sutardja '785 for an element in Sutardja '748. Nor is the proposed combination the use of a technique disclosed in Sutardja '785 to improve Sutardja '748. Nor does Petitioner identify any market or other factors creating incentive for a variation of Sutardja '748. Nor does Petitioner identify how either of the Sutardja References would allegedly be altered based on the other. Petitioner has thus failed to provide any cogent reason why a POSITA would have been motivated to combine the two Sutardja references.

Further, Petitioner is incorrect that “[b]oth Sutardja '748 and Sutardja '785 [] disclose the same subject matter of a multi-core processor with both high-power and

low-power cores.” Pet., p.17. Rather than being directed to multi-**core** processors, Sutardja ’785 discloses a system with high and low power **processors** and GPUs. Sutardja ’785 does not ever describe the use of a processor containing multiple cores, or even mention the word “core” anywhere in its disclosure. Additionally, the two Sutardja references do not contain similar elements. For example, Sutardja ’785 relates to dynamically transferring threads between processors, while Sutardja ’748 relates to copying the state of one core to another core of a different type. Petitioner’s reasoning therefore fails for at least this reason as well.

B. The Sutardja References do not render obvious Claim 1(a)(i)

Petitioner has failed to demonstrate that the Sutardja references, even in combination, disclose or suggest “a first plurality of cores and a second plurality of cores that support a same instruction set” as claimed.

For example, Petitioner aligns Sutardja ’748’s “high-speed, high-power (HP) core” with the claimed “first plurality of cores,” and Sutardja ’748’s “low-speed, low-power (LP) core” with the claimed “second plurality of cores.” Pet., p.20. However, as Petitioner implicitly acknowledges, Sutardja ’748 describes a system with just a **single** “low-speed, low-power (LP) core.” See Pet., p.20; Ex, 1007 at [0212] (describing a system “with **a** low-speed, low-power (LP) core”), [0213] (referring to “**the** LP core 12”). In fact, whereas Sutardja ’748 describes use of

multiple HP cores, *see* Sutardja '748 at [0219] (“multiple HP cores may be used”), it only ever describes the use of a single LP core.

Petitioner asserts that a POSITA would have been motivated to use multiple LP cores, but use of a single LP core is beneficial in the context of Sutardja '748's disclosure, which is focused on “mobile computing devices [that] currently use only one single-core processor.” Ex. 1007, [0211]. Sutardja '748 provides an improvement to such devices “by using a multi-core processing system with a low-speed, low-power (LP) core and a high-speed, high-power (HP) core” that uses the single LP core for processing “to increase power savings.” Ex. 1007, [215]. Thus, there would be no need or motivation to add additional LP cores: the system is designed to replace single-core processors with more power-efficient setups, but adding multiple LP cores would thwart that purpose. Additionally, Petitioner fails to explain how the system of Sutardja '748 could be altered to incorporate multiple LP cores. For example, Petitioner has not explained whether a POSITA would understand when and how to transition between LP and HP mode, given multiple of each type of core. In contrast to having a single LP core, for which Sutardja '748 explains that a demand exceeding the LP core capacity prompts transitioning processing to an HP core, Ex. 1007 at [0253], there is no teaching—and no explanation provided by Petitioner—regarding what should happen if multiple LP cores were used.

Petitioner also asserts that the second Sutardja reference, “Sutardja ’785 discloses a multi-processor system that includes... multiple secondary processors, in the form of a low-power CPU and a low-power GPU.” Pet., p.22. But Petitioner glosses over the difference between **cores** and **processors**: none of Sutardja ’785’s processors is multi-core—none discloses a processor using multiple cores at all, let alone a first plurality and a second plurality of cores. Petitioner also does not explain how or why the Sutardja references would be combined in the manner claimed, rather than, for example, the more straightforward combination: using the processors of Sutardja ’748 (each with a single LP core) in the multi-processor system of Sutardja ’785.⁶ See *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 405 (“A patent composed of several elements is not proved obvious merely by demonstrating that each element was, independently, known in the prior art.”).

Moreover, even ignoring the distinction between **cores** and **processors**, Petitioner has not demonstrated that a POSITA combining the Sutardja references would implement their teachings such that the processors (e.g., the HP CPU and GPU or the LP CPU and GPU) “support a same instruction set” as required of the

⁶ Petitioner’s position is akin to arguing that a reference disclosing a vehicle with two wheels and a reference disclosing multiple vehicles would be combined to disclose a four-wheeled vehicle, rather than the simpler combination of multiple two-wheeled vehicles.

claimed first and second “plurality of cores.” Indeed, Sutardja ’748 is focused on asymmetric cores having *different* instruction sets:

In other features, the first asymmetric core uses a first instruction set architecture (ISA). The second asymmetric core uses a second ISA. The first ISA is compatible with the second ISA. A first set of instructions of the first ISA is a superset of a second set of instructions of the second ISA. The first set includes more instructions than the second set.

Ex. 1007, [0014].

C. The Sutardja References do not render obvious Claim 1(a)(ii)

Petitioner has failed to demonstrate that the Sutardja references, even in combination, disclose or suggest “the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores.”

Although Sutardja ’748 discloses that its HP and LP cores “have different properties and capabilities (e.g., different processing speeds, power consumption, etc.),” Ex. 1007 at [0213], it does not disclose or suggest that the LP core consumes less power than the HP core “*for a same applied operating frequency and supply voltage*” as required. Instead Petitioner uses hindsight, relying on the ’080 Patent’s own disclosure to assert that Sutardja’s LP core consumes less power “for a same applied operating frequency and supply voltage.” See Pet., p.24 (“the ’080 Patent explains...”). However, Sutardja ’748 explains just the opposite—the LP core consumes less power *because it operates at a lower frequency* (Hz): “the low-speed

processor may operate at speeds that are less than or equal to 500 MegaHertz (MHz) while the high-speed processor may operate at speeds that are greater than or equal to 1 GigaHertz (GHz).” Ex. 1007, [0210]; *see also id.* at [0234]. Indeed, Sutardja ’748 names the cores by their frequencies: the “**a low-speed**, low-power (LP) core and a **high-speed**, high-power (HP) core,” with “speed” referring to operating frequency. Ex. 1007, [0212]; [0210] (referring to “speeds” in units of Hz); [0252] (same).

Petitioner asserts that Sutardja ’748 discloses that “transistors 106 of the HP core 24 tend to consume more power during operation in the active state than the transistors 110 of the LP core 28,” because “[g]ates of the transistors 106 may be larger than gates of the transistors 110.” Pet., p.24 (citing Ex. 1007 [0230]). However, Sutardja ’748 does not connect the transistor sizes to the power consumption at a given frequency and voltage—Petitioner adds the “because” based on hindsight using the ’080 Patent’s disclosure. Pet., p.24. By contrast, Sutardja ’748 explicitly states that “operation at higher speeds may also be associated with higher power consumption.” Ex. 1007, [0252]. Sutardja ’748 also makes clear that its asymmetric cores (HP and LP) have **different operation frequency**, not power consumption **at the same frequency and voltage** as required by the claims:

a first maximum speed of the first asymmetric core is
greater than a second maximum speed of the second

asymmetric core. The first asymmetric core operates at frequencies greater than a predetermined frequency. The second asymmetric core operates at frequencies less than the predetermined frequency.

Ex. 1007, [0013]; *see, e.g., id.* at [0044], cls.3, 17, 25.

Petitioner does not assert that the second Sutardja reference, Sutardja '785, discloses or suggests this limitation. *See* Pet., p.24 (never citing or referring to Ex. 1008).

D. The Sutardja References do not render obvious Claim 1(b)(i)

Petitioner has failed to demonstrate that the Sutardja references, even in combination, disclose or suggest “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” as claimed.

First, Sutardja '748 does not disclose or suggest the claimed “power management hardware.” For example, Petitioner asserts Sutardja '748’s “core switching module” is the element that “is used to transition the system between the HP and LP modes.” Pet., p.29. Petitioner thus identifies “hardware on which the disclosed hypervisor [which includes the core switching module] is installed” as the claimed “power management hardware.” Pet., p.30. As Petitioner implicitly

acknowledges, the “core switching module” is software and therefore cannot itself serve as the “power management **hardware**.” See Pet., pp.29-30 (describing it as installed on hardware); Ex. 1007, [0257] (describing the core switching module as something that “executes”), [0256] (“the core switching module 212 is loaded into the high-speed memory”). However, merely running power management software on, for example, a CPU core of a processor is insufficient to satisfy this limitation. The ’080 Patent specifically distinguishes “power management software that is running on the multi-processor (or another multi-processor or separate controller, etc.)” from a “power management process flow [that] can be executed entirely in hardware on the multi-processor or by some combination of such hardware and software.” Ex. 1001, 4:46-53. Sutardja ’748 does not describe what hardware the hypervisor and its core switching module runs on, but the only hardware described by Sutardja ’748 that can execute software is either the HP core or LP core, the very type of processors that the ’080 patent distinguishes.

Second, Sutardja ’748 does not disclose, “from a state where ***the first plurality of cores and the second plurality of cores are enabled***, disabling all of the first plurality of cores ***for a drop in demand below a threshold***” as claimed. Instead, Sutardja ’748 only describes that applications are processed by ***either*** the HP core ***or*** the LP core. For example, Petitioner points to Sutardja ’748’s HP mode as the claimed “state where the first plurality of cores and the second plurality of cores are

enabled.” Pet., p.25. However, in Sutardja ’748’s “HP mode,” only “the HP core 24 is in the active state **and processes threads.**” Ex. 1007, [0223]-[0224]. Although Sutardja ’748 describes that “[t]he LP core 28 may be in the active state during all or part of the HP mode,” it does not disclose or suggest that the LP core is enabled to **processes threads** during HP mode. Moreover, as Sutardja ’748 explains, the HP core and LP core are only ever—temporarily—both able to process threads when the system is transitioning between HP and LP modes, demonstrating that they are not normally operating at the same time. Ex. 1007, [0225] (“Both HP and LP cores 24 and 28 may **temporarily** operate in the active state at the same time when transitioning”). Petitioner implicitly acknowledges this shortcoming, asserting that a “POSITA would have been motivated to enable both Sutardja’s HP and LP cores to operate together,” for example “when workload is large,” but ignores that Sutardja ’748 teaches a different approach for such a scenario: adding additional HP cores. See Ex. 1007, [0219] (“When applications demand still higher performance than that provided by one HP core, multiple HP cores may be used”).

That the HP core and the LP core are not simultaneously enabled to process threads before a “drop in demand below a threshold” is evidenced by Sutardja ’748’s teachings about the purposes of the cores. For example, Sutardja ’748 teaches that HP mode is only used when the processing capacity of the LP core is exceeded—if the HP mode is active, it is because the LP core does not have sufficient capacity.

Ex. 1007, [0226]. However, as Petitioner acknowledges, “during the transition [between HP mode and LP mode]... the low power cores receive the ‘threads’ and other ‘information ... from the HP cores.” Pet., p.29. The LP core cannot receive threads during the transition if it is already at capacity, meaning that it is not processing threads before the transition, *i.e.*, in HP mode.

Petitioner does not point to the transitioning between states as the claimed “state where the first plurality of cores and the second plurality of cores are enabled,” and for good reason. The transition is not a state that the system is in before “drop in demand below a threshold” as claimed. For example, it is only after the threshold is crossed that the system begins the transition. Ex. 1007, [0241]. Pointing to the transitioning as the claimed state from which the first plurality of cores are disabled flips the claim on its head and eviscerates the requirement that the system be in a first state before a threshold, and then disable the first plurality of cores to arrive at a second state. Further, Sutardja ’748 explains that the transitions between modes involve the switch of all operation from one core to the other, not to the sharing of operations between both: “[w]hen the HP core 24 begins operating at a first processing speed below a first predetermined threshold, the core profile module 120 may indicate a need to **switch operation to** the LP core 28.” Ex. 1007, [0241]; see *also id.* at [0219] (“execution is switched from one core to another”). And as Petitioner acknowledges, even in the reverse direction, “the core switching module

212 may transition operation *from the HP core 204 back to the LP core 200.*” Pet., p.27 (citing Ex. 1007, [0259]).

The second Sutardja reference, Sutardja '785, also does not disclose or suggest this limitation. Petitioner does not describe any disclosures of Sutardja '785 allegedly relevant to this limitation, and only cites it once in its discussion thereof. See Pet., pp.29 (citing “Ex. 1008, ¶237”). However, this citation is merely for the proposition that “during the transition... the low-power cores receive the ‘threads’...” *Id.* However, this proposition is insufficient for the same reasons discussed above—that the HP and LP core are both powered during the switch from HP mode to LP mode does not mean that both were enabled, nor that a drop in demand below a threshold is detected or determining while both cores are enabled.

E. The Sutardja References do not render obvious Claim 1(b)(ii)

Petitioner has failed to demonstrate that the Sutardja references, even in combination, disclose or suggest “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” as claimed.

Petitioner asserts that Sutardja '748 “discloses that a core profile module included in a ‘power management system (PMS) module’ (illustrated in Figure 5 []) monitors this demand.” Pet., p.32. Petitioner alleges that the “core profile module” is part of the operating system—as required by the claims under Petitioner’s

reading—because it is illustrated in FIG. 5, a “functional block diagram.” Pet., p.32.

For example, Petitioner asserts that a “POSITA would have understood that functional block diagrams like Figure 5 are used to depict how various software modules function together,” but ignores that Figure 5 also illustrates hardware as well: the “LP Core” and the “HP Core” are shown at reference numbers 200 and 204, as are “I/O Devices” (Input/Output devices) at reference number 234.

Contrary to Petitioner’s assertion, neither the power management system nor its core profile module are part of the operating system of Sutardja ’748. Indeed, Petitioner falls short of saying that either is: Petitioner asserts that “the PMS Module is **executed by** the kernel module (part of the Linux OS).” Pet., p.32. However, merely being “executed by” the operating system does not make an application or module **part of** the operating system such that the system contains “an operating system ... to monitor a demand for the multi-core processor and control the power management hardware” as claimed. Software executed by an operating system (or a component thereof) is not part of the operating system; if it were, every application running on a computer would be part of the operating system.

The PMS module is repeatedly illustrated as being separate from the operating system in Sutardja ’748. For example, in FIGs. 9A and 9B, the PMS Module 402 is separate from the “OS Services Module” 416, which contains the Kernel IOCTL Module 422. Indeed, the remainder of Figure 5 indicates that the PMS module is not

part of the operating system. For example, Sutardja '748 depicts its hypervisor module 213 as having the same level of connection to the kernel module as the PMS module, and Petitioner asserts that the hypervisor module is **not** part of the OS. *See* Pet., p.30 (“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”).

Moreover, Petitioner’s assertion that Sutardja '748’s “operating system... is to monitor a demand for the multi-core processor and control the power management hardware” as claimed ignores the repeated teaching of Sutardja '748 that its switching of cores is “transparent” or “invisible” to the operating system. *See* Ex. 1007 at [0012], [0022], [0043], [0053], [0073], [0083], [0104], [0114], [0137], [0151], [0165], [0179], [0216], [0219], [0256]-[0257], [0269], [0298], [0307], cls. 8, 22, 36, 50. For example, Sutardja '748 describes that “the switching of cores is transparent or invisible to the OS and applications, which continue to run as if the system comprises a single core instead of multiple cores.” Ex. 1007, [0219]. Because core switching is invisible to the OS, the switching module cannot be part thereof.

F. The Sutardja References do not render obvious Claims 9 and 17

Petitioner merely cites to the arguments made with respect to claim 1 for its analysis of independent claims 9 and 17. *See* Pet., pp.40-41. The Petition therefore

fails to demonstrate that claims 9 and 17 are obvious for the same reasons that it fails with respect to claim 1.

VII. Ground 4 Fails

The Petition asserts in Ground 4 that Claims 1-4, 7-12, 15-20, and 23-24 are rendered obvious by Mathieson in view of the Sutardja references.

A. A POSITA Would Not Have Been Motivated to Combine Mathieson with the Sutardja References

Petitioner asserts that a “POSITA would have been motivated to combine the teachings of Mathieson and Sutardja... because each relates to the same well-known subject matter of managing power in multi-core processors,” and that “both references contemplate that power consumption differences between cores may be achieved through differences in constituent transistors.” Pet., p.50. However, relying on the references’ alleged use of “differences in the cores’ constituent transistors” is not a sufficient motivation to combine these references, and improperly relies on hindsight. For example, merely because references could be compatible or relate to the same subject matter does not provide a motivation to combine them. *Johns Manville Corp. v. Knauf Insulation, Inc.*, IPR2018-00827, Paper 9 at 10 (Designated Informative) (citing *Securus Techs., Inc. v. Global Tel*Link Corp.*, 701 F. App’x 971, 977 (Fed. Cir. 2017) and *Personal Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 993 (Fed. Cir. 2017)). Neither Mathieson nor Sutardja ’748 is directed toward

a problem with transistors, or a means of improving transistors within processor cores—the fact that they both use transistors does not relate to solving any articulated problems or motivate a POSITA to combine them. Moreover, the use of allegedly different transistors does not relate at all to the motivation to combine Mathieson with Sutardja '785, which does not even disclose the use of cores at all, let alone the properties of the transistors within those cores. Nor does Petitioner provide any analysis with respect to Sutardja '785 in connection with this combination.

Indeed, as to power management of processor cores, Mathieson and Sutardja '748 actually teach divergent methods. For example, Sutardja '748 uses a core switching module in connection with non-OS software to determine whether to operate the low power core or the high power core based on present workload. *See* Ex. 1007, [0257]. By contrast, Mathieson manages power through a “hardware mechanism that stops and powers off the first set of cores 210, and boots the second set of cores 220.” Ex. 1005, [0043]. Thus, Sutardja '748 is explicitly carried out via software, whereas a dedicated hardware structure is used in Mathieson. Petitioner has not demonstrated that such different approaches to power management would be compatible, or whether they would require undue experimentation to combine.

Petitioner nowhere analyzes the traditional considerations of a motivation to combine analysis. For example, Petitioner does not allege that Sutardja '748 provides a solution to a problem identified in Mathieson. Nor does the proposed

combination represent a substitution of any known elements from Sutardja '748 for an element in Mathieson. Nor is the proposed combination the use of a technique disclosed in Sutardja '748 to improve Mathieson. Nor does Petitioner identify any market or other factors creating incentive for a variation of Mathieson. Nor does Petitioner identify how either of the references would allegedly be altered based on the other in combining them. Petitioner has thus failed to provide any cogent reason why a POSITA would have been motivated to combine Mathieson with the Sutardja references.

B. Mathieson, Alone or in View of the Sutardja References, Does Not Render Obvious Claims 1, 3-4, and 7-8

1. Mathieson and the Sutardja References do not render obvious Claim 1(a)(i)

Petitioner has failed to demonstrate that Mathieson, alone or in combination with the Sutardja References, discloses or suggests “a first plurality of cores and a second plurality of cores that support a same instruction set” as claimed.

For example, Petitioner does not assert that Mathieson discloses the first and second pluralities of cores supporting the same instruction set, and instead asserts that “a POSITA would have understood that a straightforward way to” implement Mathieson would be to do so. Pet., p.53. Petitioner improperly relies on hindsight in asserting that a POSITA “would [] implement the same ISA (and therefore support the same instruction set) in each core,” because neither Petitioner nor its declarant

identify any motivation or justification for doing so. Pet., p.53. For example, Mathieson does not describe ISAs, and Petitioner does not point to any disclosures of Mathieson or either Sutardja reference disclosing the use of ISAs at all, let alone use of the **same** ISA for **different** cores. Rather, Petitioner relies on Mathieson's statement that the sets of cores "may implement substantially the same functionality," but having "substantially the same functionality" does not disclose or suggest that such cores support the same ISA. Different cores can support substantially the same functionality using different instruction sets.

Petitioner also argues that a "POSITA would have understood that a straightforward way ensure the ability to directly transfer the processing state is for both sets of cores to support the same ISA." Pet., p.53. However, Petitioner is merely using the teachings of the '080 Patent—textbook impermissible hindsight—to provide the alleged motivation to use the same instruction. For example, the '080 Patent identifies a problem that its invention solves: in systems in which "the cores are not identical ... [a] problem ... is that it is difficult for system software to adjust switch operation between processor cores having different instruction sets." Ex. 1001, 3:34-49. By contrast, Mathieson neither identifies this as a problem, nor discloses or suggests the claimed solution. In fact, Mathieson explains that each given set of cores may implement a specialized function, suggesting each would have a particularized instruction set. See Ex. 1005, [0026] ("each given set of cores

210, 220 may implement a particular functional block of the CPU 102, such as an arithmetic and logic unit, a fetch unit, a graphics pipeline, a rasterizer, or the like”).

Petitioner also argues that Mathieson would be combined with the Sutardja references for this element. For example, Petitioner asserts that a POSITA “would have been motivated to search for implementation details in similar references” and that “Sutardja is one such reference.” Pet., p.54. Setting aside that Petitioner has not demonstrated that Mathieson and either Sutardja reference are “similar references,” neither Sutardja reference provides any alleged “implementation details” that Petitioner asserts a POSITA would seek out. For example, Petitioner does not identify any aspect of Mathieson that would be realized according to “implementation details” in either Sutardja reference. Nor does Petitioner identify any disclosures of either Sutardja reference that provide “implementation details” explaining how to go about making cores support *any* instruction sets, let alone *the same instruction set on different cores*. The Sutardja references are simply at a different level of disclosure—architecture level, rather than granular—and do not provide any “implementation details” that allegedly would have motivated modification of Mathieson.

Contrary to Petitioner’s assertion that a POSITA would have looked to Sutardja for “implementation details” regarding how to have different cores support the same instruction set, Sutardja ’748 focuses its disclosure on asymmetric cores

having *different* instruction sets. *See, e.g.* Ex. 1007, [0014] (“the first asymmetric core uses a first instruction set architecture (ISA). The second asymmetric core uses a second ISA”); [0213] (“An instruction set architecture (ISA) of the LP core 12 may differ from the ISA of the HP core”). Although Sutardja ’748 gives brief mention that “both may use the same ISA’s,” it does not provide any “details” regarding such alternative that a POSITA would allegedly have sought out. *See* Ex. 1007, [0213]. Nor does Sutardja ’748 identify any benefits of using the same ISA for different cores. *See id.* This is in contrast to the identified benefits of using different instruction sets that Sutardja ’748 discloses—different performance and capabilities between the asymmetric cores:

An instruction set architecture (ISA) of the LP core **12** may differ from the ISA of the HP core **16**. For example, the HP core **16** may execute floating point instructions while the LP core **12** may not. Additionally, the LP and HP cores (hereinafter cores) **12** and **16** have different performance and capabilities (e.g., different processing speeds, power consumption, etc.).

Ex. 1007, [0213]; *see also* Pet., p.24 (citing same). Petitioner ignores Sutardja ’748’s focus, and does not provide any explanation regarding why a POSITA seeking to implement Mathieson in view of either Sutardja reference would choose to rely on an unexplained alternative of Sutardja ’748 as opposed to its main teaching.

2. Mathieson and the Sutardja References do not render obvious Claim 1(b)(i)

Petitioner has failed to demonstrate that Mathieson, alone or in combination with the Sutardja References, discloses or suggests “a state where the first plurality of cores and the second plurality of cores are enabled” to perform operations as claimed.

For example, Petitioner relies on a disclosure from Mathieson that “the CPU 102 may operate in both the first mode and the second mode simultaneously *during a transition period* when transitioning between the first mode and second mode, or vice versa,” but ignores that Mathieson describes this is only in a transition state and does not disclose or suggest that the cores are enabled to perform operations during the transition. *See* Ex. 1005, [0031]; Pet., p.57 (citing same). Operating in both *modes* simultaneously does not indicate that both sets of cores are “enabled” to perform operations as claimed. Mathieson does not disclose or suggest the cores are “enabled” to perform operations as claimed during the transition period. In fact, Mathieson counsels against simultaneous operation of multiple sets of cores: “operating in both the first and second modes simultaneously, may result in lower overall power efficiency.” Ex. 1005, [0031].

Petitioner annotates step 402 of FIG. 4A, but ignores that is labeled within the figure as (and limited to) a step to “process one or more operations *on a first set of*

cores.” See Pet., p.57; FIG. 4A. Processing operations on a first set of cores does not disclose or suggest “a state where the first plurality of cores and the second plurality of cores are enabled” as claimed. In fact, Mathieson explains that very step as having the second set of cores disabled:

As shown, the method 400A begins at step 402, where a controller included in the processor causes one or more operations to be executed by a first set of cores. In one embodiment, ***when processing the one or more operations using the first set of cores, the cores included in the second set of cores are disabled and powered off.***

In alternative embodiments, when processing the one or more operations using the first set of cores, the cores included in the second set of cores are clock gated and/or power gated.

Ex. 1005, [0055] (with clock or power gating turning off the cores)

Petitioner has also failed to demonstrate that Mathieson, alone or in combination with the Sutardja References, discloses or suggests “disabl[ing] all of the first plurality of cores for a drop in demand below a threshold” as claimed. That is, even if Mathieson’s transition period serves as “a state where the first plurality of cores and the second plurality of cores are enabled” as claimed, it does not disclose or suggest disabling the first set of cores from that transition period “for a drop in demand below a threshold” as claimed. For example, Mathieson describes switching

from the first operating mode to the second operating mode upon detecting a decreased workload, but does not disclose or suggest starting from a transition state and then disabling the first set of cores in response to a change in demand. *See* Ex. 1005, [0043]. Any alleged drop in demand is detected while operating in the first mode (with just a single set of “cores” active), not during the transition period—the decision is made before the transition begins.

Petitioner also relies on another reference, Carmack, for this limitation. Pet., p.60. For example, Petitioner asserts that “Carmack [is] incorporated by reference into Mathieson,” but does not explain how, or demonstrate which parts, if any, of Carmack are properly incorporated into Mathieson. For example, much like with its analysis of the Sutardja references, Petitioner does not analyze the alleged incorporation by reference: “To incorporate matter by reference, a host document must contain language clearly identifying the subject matter which is incorporated and where it is to be found; a mere reference to another application, or patent, or publication is not an incorporation of anything therein.” *Callaway Golf*, 576 F.3d at 1346. Mathieson does not identify any specific disclosure from Carmack allegedly relevant to any of its embodiments or components.

Petitioner also asserts that “[t]o the extent that any relied upon teaching is argued to have not been fully incorporated from Ex-1006 into Ex-1005, it would have been obvious to combine such teaching from Ex-1006 with Ex-1005, for the

same reasons provided herein.” Pet., n.2. However, Petitioner does not explain any alleged motivation to combine the teachings of the Mathieson with Carmack. Petitioner nowhere analyzes the traditional considerations of a motivation to combine analysis. For example, Petitioner does not allege that Carmack provides a solution to a problem identified in Mathieson. Nor does the proposed combination represent a substitution of any known elements from Carmack for an element in Mathieson. Nor is the proposed combination the use of a technique disclosed in Carmack to improve Mathieson. Nor does Petitioner identify any market or other factors creating incentive for a variation of Mathieson. Nor does Petitioner identify how either Mathieson or Carmack would allegedly be altered based on the other. Petitioner has thus failed to provide any cogent reason why a POSITA would have been motivated to combine Mathieson with Carmack. This failure is further compounded by the fact that this Ground relies not just on Mathieson, but Mathieson in combination with “Sutardja”—itself a combination of Sutardja ’748 with Sutardja ’785—meaning that the Ground is rooted in combination of **four** references, none of which has Petitioner analyzed the inter-reference interactions for.

Even looking to Carmack, Petitioner has failed to demonstrate that it cures the deficiencies of Mathieson (even in view of Sutardja ’748 and Sutardja ’785 as this Ground proposes). For example, Petitioner cites to “Carmack’s second cores 120” and “Carmack’s first cores 110” as the claimed first and second pluralities of cores.

Pet., p. 60. However, much like Mathieson, Carmack merely teaches “switch[ing] between the first 110 and second core 120 and vice versa,” not “a state where the first plurality of cores and the second plurality of cores are enabled” to perform operations as claimed. Ex. 1006, [0016]. Carmack discloses use of either core as mutually-exclusive alternatives:

At 220, a first core 110 of the integrated circuit 100 is utilized and a second core 120 is idled if the performance parameter is within a first predetermined range. At 230, the second core 120 is utilized and the first core 110 is idled if the performance parameter is within a second predetermined range.

Ex. 1006, [0014]. Moreover, Petitioner ignores that Carmack teaches away from “disabl[ing] all of the first plurality of cores for a **drop in demand** below a threshold” as claimed, teaching instead the opposite approach. For example, Carmack explicitly teaches switching cores in response to an **increase** in workload:

When the workload **exceeds** a specified level, power to the second core 120 could be turned on and the context of the first core 110 transferred to the second core 120. Thereafter the power to the first core 110 may be turned off.

Ex. 1006, [0018]. Petitioner does not evaluate this aspect of Carmack, or how it would impact the combination of Mathieson in view of Carmack.

3. Mathieson and the Sutardja References do not render obvious Claim 1(b)(ii)

Petitioner has failed to demonstrate that Mathieson, alone or in combination with the Sutardja References, discloses or suggests “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” as claimed.

Petitioner points to Mathieson’s “controller 240” as allegedly satisfying this limitation, but fails to demonstrate that component is part of “an operating system” as claimed. *See* Pet., p.62. Recognizing that Mathieson does not explicitly disclose control by the operating system, Petitioner instead asserts that a “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.” Pet., p.63. However, Mathieson teaches that the controller 240 is hardware within the CPU 102. *See, e.g.*, Ex. 1005, [0030] (“a controller 240 included within the CPU 102”). Indeed, “controller 240” is shown as a subcomponent of CPU 102 in Figure 2 of Mathieson. Ex. 1005, Fig. 2. Petitioner’s annotations notwithstanding, there is no “software” illustrated in FIG. 2. *See* Pet., p.63. Petitioner does not address the additional disclosures of Mathieson that make clear that “controller 240” is a hardware component. For example, Mathieson discloses that controller 240

transmits a signal to the multiplexor 332 (Ex. 1005, [0051]), transfers processor states (Ex. 1005, [0043]), and reads performance and power data from fuses on the CPU (Ex. 1005, [0037]).

Even in the limited “embodiments [in which], the controller is implemented in software and is executed by the CPU 102,” Mathieson’s controller 240 is not “an operating system to execute on the multi-core processor” as claimed. *See* Ex. 1005, [0030]. For example, Petitioner does not address Mathieson’s valuable feature that controlling power management “can be done transparently to high level software, such as the operating system.” Ex. 1005, [0043]. Rather, Mathieson teaches away from having “an operating system ... monitor a demand” as claimed by expressly discouraging the involvement of the operating system in the disclosed core switching.

VIII. Conclusion

For at least the foregoing reasons, the Petition does not meet its burden of demonstrating a reasonable likelihood that Petitioner would prevail with respect to any of the challenged claims. Thus, Patent Owner respectfully requests that the Board deny institution.

Dated: August 14, 2023

Respectfully submitted,

/Peter F. Snell/

Peter F. Snell (Reg. No. 52,235)

MINTZ, LEVIN, COHN, FERRIS,
GLOVSKY AND POPEO, P.C.

919 Third Avenue

New York, New York 10022

Telephone: 212-935-3000

Facsimile: 212-983-3115

E-mails:

PFsnell@mintz.com

CERTIFICATE OF WORD COUNT

The undersigned certifies pursuant to 37 C.F.R. § 42.24(d) that the foregoing Patent Owner Preliminary Response, excluding any table of contents, table of authorities, certificates of service or word count, or appendix of exhibits or claim listing, contains 11,454 words according to the word-processing program used to prepare this paper (Microsoft Word). Patent Owner certifies that this Patent Owner Preliminary Response does not exceed the applicable type-volume limit of 37 C.F.R. § 42.24(a).

Dated: August 14, 2023

/Peter F. Snell/

Peter F. Snell (Reg. No. 52,235)

CERTIFICATE OF SERVICE

I certify that a copy of Patent Owner's Preliminary Response is being served
by electronic mail on the following counsel of record:

Lead Counsel

William M. Fink, (Reg. No. 72,332)
O'MELVENY & MYERS LLP.
1625 Eye Street, NW
Washington, DC 20006
Telephone: (202) 383-5300
Facsimile: (202) 383-5414
Email: tfink@omm.com

Backup Counsel

Daniel Leventhal (Reg. No. 59,576)
Richard Zembek (Reg. No. 43,306)
Darren Smith (Reg. No. 64,261)
NORTON ROSE FULBRIGHT
1301 McKinney, Suite 5100
Houston, TX 77010
Telephone: (713) 651-5151
Facsimile: (713) 651-5246
Emails:
daniel.leventhal@nortonrosefulbright.com
richard.zembek@nortonrosefulbright.com
Darren.smith@nortonrosefulbright.com

Eagle H. Robinson (Reg. No. 61,361)
NORTON ROSE FULBRIGHT
98 San Jacinto Blvd., Suite 1100
Austin, TX 78701
Telephone: (512) 474-5201
Facsimile: (512) 536-4598
Email:
eagle.robinson@nortonrosefulbright.com

*Counsel for Petitioner Qualcomm
Incorporated*

Dated: August 14, 2023

/Peter F. Snell/

Peter F. Snell (Reg. No. 52,235)