

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NXP USA, INC. and QUALCOMM INCORPORATED.,
Petitioner,

v.

REDSTONE LOGICS LLC,
Patent Owner.

Case IPR2025-00485
Patent 8,549,339

PATENT OWNER'S PRELIMINARY RESPONSE

TABLE OF CONTENTS

I.	Introduction	1
II.	Background	2
	A. NXP and Qualcomm’s Petition.....	2
	B. Overview of the ’339 Patent	3
III.	Level of Ordinary Skill in the Art.....	5
IV.	Claim Construction	5
V.	The Petition’s Ground 1 Fails	5
	A. No Motivation to Combine White and Talwar	5
	B. First and Second Clock Signals as Input Are Not Disclosed In the Art	12
VI.	The Petition’s Ground 4 Also Fails.....	13
	A. Ground IV Fails to Disclose Sets of Processor Cores	14
	B. Ground IV Fails to Disclose First and Second PLLs.....	21
	C. No Motivation or Ability to Combine Ozer.....	23
VII.	Conclusion	25

TABLE OF AUTHORITIES

Cases

<i>Facebook v. Windy City Innovations, LLC</i> , 973 F.3d 1321 (Fed. Cir. 2020)	22
<i>Intel Corp. v. PACT XPP Schweiz AG</i> , 61 F.4th 1373 (Fed. Cir. 2023)	20, 23
<i>PersonalWeb Techs., LLC v. Apple, Inc.</i> , 917 F.3d 1376 (Fed. Cir. 2019)	13, 16
<i>Xerox Corp. v. Bytemark, Inc.</i> , IPR2022-00624, Paper 9 (PTAB August 24, 2022)	21, 22

PATENT OWNER'S EXHIBIT LIST

No.	Description
Ex. 2001	<i>Redstone v NXP USA, Inc.</i> , No. 7:24-cv-00028 ADA-DTG, Dkt. No. 36 Scheduling Order
Ex. 2002	<i>Redstone v Qualcomm, Inc., Qualcomm Technologies, Inc.</i> , No. 7:24-cv-00231 ADA, Dkt. No. 24 Scheduling Order
Ex. 2003	Motion to Stay before Judge Albright, Docket Navigator
Ex. 2004	Guidance on USPTO's recission of "Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation"
Ex. 2005	Western District of Texas Case Management Statistics
Ex. 2006	NXP's Preliminary Invalidity Contentions
Ex. 2007	NXP's Amended Invalidity Contentions
Ex. 2008	Qualcomm's Invalidity Contentions
Ex. 2009	Exhibit B to NXP's Amended Invalidity Contentions
Ex. 2010	Exhibit B to Qualcomm's Invalidity Contentions
Ex. 2011	3/26/2025 email from Eric Green, Redstone Logics LLC v. NXP USA, Inc., Case No. 7:24-cv-00028-ADA-DG, in the United States District Court for the Western District of Texas
Ex. 2012	3/26/2025 Email from Daniel Leventhal, Redstone Logics LLC v. Qualcomm Incorporated, et al., Case No. 7:24-cv-00231-ADA, in the United States District Court for the Western District of Texas
Ex. 2013	Interim Processes for PTAB Workload Management
Ex. 2014	<i>Redstone v NXP USA, Inc.</i> , No. 7:24-cv-00028 ADA-DTG, Dkt. No. 50 Claim Construction Order
Ex. 2015	<i>Redstone v MediaTek Inc., et al</i> , No. 7:24-cv-00029 ADA-DTG, Dkt. No. 40 Claim Construction Order
Ex. 2016	MediaTek's Amended Invalidity Contentions
Ex. 2017	US Patent No. 7,853,808 to Kim et al ("808 Kim")
Ex. 2018	US Patent Publication No. 2009/0138737 to Kim et al ("Kim")

No.	Description
Ex. 2019	US Patent Publication No. 2009/0106576 to Jacobowitz et al (“Jacobowitz”)
Ex. 2020	3/31/2025 Email from Joshua Scheufler, Re: Redstone Logics LLC v. Qualcomm Incorporated, et al., Case No. 7:24-cv-00231-ADA, in the United States District Court for the Western District of Texas
Ex. 2021	4/20/2025 Email from Daniel Leventhal, Re: Redstone Logics LLC v. Qualcomm Incorporated, et al., Case No. 7:24-cv-00231-ADA, in the United States District Court for the Western District of Texas
Ex. 2022	<i>MediaTek, Inc. et al v Redstone Logics LLC</i> , IPR2025-00085, Paper 7 (April 22, 2025)
Ex. 2023	NXP Motion to Stay Pending Settlement
Ex. 2024	MediaTek Motion to Stay Pending Settlement
Ex. 2025	Crouch, Dennis. “Estopple Gutted: A Pelican’s Guide to Patent Litigation.”
Ex. 2026	Crouch Dennis. “Reference Recycling and the Case for Sua Sponte Rehearing in Igenico”
Ex. 2027	Declaration of Dr. Michael C. Brogioli
Ex. 2028	<i>Cirriculum Vitae</i> of Dr. Michael C. Brogioli
Ex. 2029	<i>Redstone v Qualcomm Inc. et al.</i> , No. 7:24-cv-00231 ADA, Dkt. No. 32 Joint Claim Construction Statement
Ex. 2030	Jon Stokes, “Two billion-transistor beasts: Power7 and Niagara 3”
Ex. 2031	Jinuk Shin, <i>et al</i> “A 40nm 16-Core 128-Thread SPARC Soc Processor”

I. Introduction

Neither of the Petition's two grounds addressing the independent claims demonstrates a reasonable likelihood of prevailing in demonstrating invalidity of any claim of the '339 Patent. Rather, the Petition treats any missing limitation as an obvious, well-known technique in the art and proposes to combine art a POSITA would be unable or unwilling to combine.

The first set of grounds offer a motivation to combine that directly counters the teachings of the primary reference. White teaches a device that may “include *any number* of logic cores configured to operate at independent voltages and/or frequencies.” *See* Pet. at 25; Ex. 1006 at 6:56-60. But somehow White's design “would *need* to be adapted as the number of cores increases” based on real estate concerns. Pet. at 25. But even this “real estate” concern is unfounded as any increase in space needed to house voltage or clock control components is de minimis or readily handled off chip.

The second set of grounds similarly suffer in finding a basis to suggest the claimed “sets of cores” is obvious. Instead of relying on a reference disclosing sets of cores, the Petition argues the disclosure of a standard multi-core processor is somehow enough. But multiple sets of cores is a far leap from even multiple, multi-core processors. Such an arrangement requires different coherency and interfacing infrastructure which neither Elgebaly nor Ozer provide.

Because all grounds fail to disclose or render obvious key limitations of the independent claims, the Petition has not shown a substantial likelihood of success that any claim is invalid. The Petition should be denied.

II. Background

A. NXP and Qualcomm's Petition

Petitioners NXP USA, Inc.¹ (“NXP”) and Qualcomm Incorporated (“Qualcomm”) (collectively “Petitioners”) challenges claims 1-3, 5, 8-11, 14, and 21 of U.S. Patent No. 8,549,339 (the “’339 Patent”). Petitioner structures its asserted grounds for unpatentability as follows:

- Ground 1: Obviousness of claims 1-3, 5, 8, 10-11, 14, and 21 over White in view of Talwar
- Ground 2: Obviousness of claims 5, 14 over White and Talwar in view of Cho.
- Ground 3: Obviousness of claims 8-9 over White and Talwar in view of Vajda.
- Ground 4: Obviousness of claims 1-3, 5, 8, 10-11, 14, and 21 over Elgebaly in view of Ozer.
- Ground 5: Obviousness of claims 8-9 over Elgebaly and Ozer in view of Vajda.
- Ground 6: Obviousness of claims 1-3, 5, 8, 10-11, 14, and 21 over Elegabaly and Ozer in view of Williams.

¹ NXP has moved to withdraw though Qualcomm maintains the joint Petition.

- Ground 7: Obviousness of claim 8-9 over Elegabaly, Ozer, and Vajda in view of Williams.

Because all challenged dependent claims across all grounds depend from claim 1, this Preliminary Response need only address the Petition's arguments concerning claims 1 and 21. If the Petition does not show a reasonable likelihood of prevailing on claims 1 and 21, then it cannot prevail on any other challenged claim.

B. Overview of the '339 Patent

At a high level, the '339 Patent is directed generally to a novel operation of a multi-core processor enabling coordination of groups of multiple cores operating at different supply voltages and clock signals. *See* Ex. 1001 at 2:4–3:15.

The '339 Patent explains: a “multi-core processor [] may be divided into stripes... [where e]ach stripe may be associated with an independent power profile ... and/or may be associated with an independent clock domain defined by a clock signal received from a clock control block.” *Id.* at 2:25–31. However, “[t]o maintain the limited differential relationship discussed above, adjusting the supply voltage to one stripe may involve adjusting the supply voltages to the other stripes.” *Id.* at 3:8–10.

To address this, the '339 Patent describes an architecture to achieve this wherein “the stripe [] may be powered by a supply voltage received from a power

control block[] and/or may be associated with an independent clock domain defined by a clock signal received from a clock control block[].” *Id.* at 2:26–31.

This solution to the prior art problem of dynamic management of sets of processor cores is reflected, for example, in claim 1 of the ’339 Patent, which recites:

1. A multi-core processor, comprising:
a first set of processor cores of the multi-core processor, wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage and a first output clock signal of a first phase lock loop (PLL) having a first clock signal as input;
a second set of processor cores of the multi-core processor, wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage and a second output clock signal of a second PLL having a second clock signal as input, wherein the first supply voltage is independent from the second supply voltage, and the first clock signal is independent from the second clock signal; and
an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores, wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.

Ex. 1001 at claim 1.

All other challenged dependent claims 2-6, 8-11, and 14 depend from claim 1 and claim 21 reads similarly to claim 1.

III. Level of Ordinary Skill in the Art

The Petition proposes that a person of ordinary skill in the art in the field of the '339 Patent “would have had at least (1) a bachelor’s degree in electrical engineering, computer engineering, computer science, or a similar field, and (2) at least two years of academic or industry experience designing or analyzing electronic circuits, semiconductors, processors, or power management, and related firmware and software, or the equivalent.” Paper 1 at 8. For purposes of this preliminary response, Patent Owner does not challenge that definition.

IV. Claim Construction

For purposes of this preliminary response, Patent Owner only notes in the related district court litigation the parties agreed “a first/second set of processor cores” should be construed as “a first/second group of two or more processor cores.” Ex. 2029. This construction is helpful, though not necessary to understanding both the Petition and this preliminary response.

V. The Petition’s Ground 1 Fails

Grounds 1-3 fail to justify the central combination of White and Talwar and makes unsupported assumptions to meet the PLL limitations.

A. No Motivation to Combine White and Talwar

Grounds 1-3 all rely on the same primary combination of White and Talwar, in particular, the architecture of White with the clusters taught by Talwar. *See*

Petition at 26-28. But the only motivation for such a combination enumerated by the Petition is a “need” that contradicts the express teachings of White and would have destroyed any advantage of following the teachings of White.

The only basis the Petition enumerates for combining White and Talwar is that a POSITA would supposedly “have understood [White’s multicore processor design] would *need* to be adapted as the number of cores increases due to chip space constraints associated with adding an additional voltage and clock control and supply circuitry for each additional core.” *Id.* at 25 (emphasis added). But this “*need*” does not exist. In fact, White itself explains the device it teaches may “include *any number* of logic cores configured to operate at *independent voltages and/or frequencies*.” Ex. 1006 at 6:56-60 (emphasis added). If any number of cores may be configured to operate with their own independent voltages and frequencies, there is no *need* for an architecture that provides grouped voltages and frequencies as the number of logic cores increases. A POSITA considering White would have no *need* to alter it, much less a motivation to achieve any particular configuration. Ex. 2027 at ¶30.

Moreover, “chip space” is not a constraint in White. White explains “a number of different voltages may be generated and/or regulated an input to integrated circuit 100 through various connections.” Ex. 1006 at 8:29-31. “Likewise, multiple clock signals at differing frequencies may be generated either off-chip or internally and

coupled to individual logic cores 120 and/or common bridge logic 110.” *Id.* at 8:34-37. White already has a solution to the “chip space” constraint issue, simply generate the voltages and clock signals outside the integrated circuit or “off-chip.” Ex. 2027 at ¶¶31-32. As White provides the solution that does not require any change to the alleged “need” underpinning the Petition’s motivation to change White, no motivation exists to combine White and Talwar. *Id.*

Moreover, even if White’s own teaching that chip space constraints do not motivate any change were ignored, the Petition cannot support this chip-space-constraint motivation.

First, the Petition relies on Dr. Annavaram to create this need and rewrite White, but his support actually demonstrates why Talwar’s “clusters” would not be used. *See id.* at ¶¶33-36. While Dr. Annavaram cites to references that acknowledge certain design considerations, they explain how a POSITA actually responds to those considerations and do not align with the Petition. *Id.* at ¶34.

While both Ex. 1019 and 1029 acknowledge there is some greater cost associated with multiple power and clock domains, each shows why that cost is not a driving consideration. *See* Ex. 1002 at ¶73 (citing Ex. 1029 at 3-4, Ex. 1019 at 52-53). Ex. 1019 explains that research considering “the benefits of per-core versus per-chip DVFS” found that “on a four-core CMP in which DVFS was being employed to avoid thermal emergencies (rather than simply to save power), a per core approach

had 2.5X better throughput than a per-chip approach.” Ex. 1019 at 53. Notably, temperature management is the primary concern White discusses with its DVFS system. Ex. 2027 at ¶35, *see e.g.* Ex. 1006 at 1:37-52 (“Since complex ICs such as microprocessors tend to be expensive, it is typically desirable to protect them from damage due to overheating.”); 2:29-44 (“Also, rather than simply shutting down a processor that is operating to warmly(as is generally done in the prior art), one or more of the logic cores in the processor may be adjusted to operate at a lower frequency and/or voltage so that the processor may generate less heat, and therefore may cool down to an acceptable temperature.”); Claims 2-10, 13-14, 16-24, 26-29, 31-32 (all claiming a temperature control limitation). Moreover, the per-core design of White along with temperature sensors integrated into the DVFS, would lead a POSITA to use this design in situations where temperature management is key over other considerations. Ex. 2027 at ¶35, Ex. 1006 at 12:16-58. A POSITA considering White and the potential 2.5X throughput advantage over chip-wide DVFS would not have been motivated to change White to make it more like a chip-wide model. Ex. 2027 at ¶35. Indeed, a POSITA would recognize, in light of this throughput advantage, it could simply reduce the number of cores it needs rather than trying to save on peripherals like voltage and clock sources. *Id.* A POSITA would therefore not find any savings associated with grouping cores to motivate a change to White. *Id.*

Ex. 1029 reinforces this, explaining “depending on the heterogeneity of workload characteristics, per-core DVFS offers *substantial additional savings* compared to global DVFS schemes by better adapting to the different requirements of each core.” Ex. 1029 at 4 (emphasis added). Moreover, it explains that while older on-chip regulators were significant space concern, “[r]ecently, on-chip regulators have been proposed, integrated on the same die as the processor load... [h]ence, on-chip regulators offer an interesting solution that can supply multiple power domains in CMPS with per-core DVFS.” *Id.* at 6. This is more than a mere academic distinction, real world multi-core chips from the time of the invention had no issue with chip-space for DVFS. *Id.* at ¶36. For example, the Niagara or UltraSPARC T3 announced at the time of the invention, *see* Ex. 2030, wastes no recognizable chip space on voltage and clock control:




Figure 5.2.6: Die Micrograph.

Ex. 2027 at ¶37. As shown above, the vast majority of chip space is taken up with the cores and memory components (SPARC 0-15, L2D 0-15, L2T 0-7, Memory). *Id.* The other major components are interface circuitry like PCIe or the CLC and coherency components that address transactions within the chip. *Id.*, *see also* Ex. 2031 at 2. In the real world, chip space is not a significant enough concern to govern the chosen DVFS mode.

Lastly, while not included in the Petition as part of the motivation to combine, Dr. Annavaram misreads White's discussion of how its system may be applied. White explains that its "independently controlling operating voltages and/or frequencies may be applied to any integrated circuit or component..." See Ex. 1002 at ¶75 (citing Ex. 1006 at 8:44-49). But that application to "any integrated circuit" is explained and does not relate to cores. Ex. 2027 at ¶38. White explains "[f]or instance, a controller, bridge or other components configured to interface to, or communicate with multiple components may be configured to independently control the frequency and or voltages of those interfaces." Ex. 1006 at 8:49-53. This teaches applying domain crossing techniques to whatever circuitry requires it, not arranging cores in any manner. Ex. 2027 at ¶38. A POSITA would not be able to take this teaching of White to create the clusters referenced in Talwar. *Id.*

The Petition has failed to demonstrate any motivation to change White. Rather, its own support demonstrate that such a change would likely decrease the functionality of White for no advantage. As such the Petition cannot show sets of processor cores would have been obvious and there is no likelihood of success on Grounds 1-3.

B. First and Second Clock Signals as Input Are Not Disclosed In the Art

Even if the White/Talwar combination is considered, the Petition fails to show that the claimed PLLs would have been an obvious implementation of Talwar's clusters. Rather, the Petition merely contends "clock control circuitry that White describes at the core level (e.g., a PLL at the core and accompanying 'control registers' storing 'PLL configuration parameters') would operate at the cluster level." Petition at 27. But there is no explanation for why or how this could be done.

First, the Petition does not show such a use of PLLs is disclosed, expressly or inherently. Talwar makes no reference whatsoever to PLLs or a particular clock architecture. *See* Ex. 1007. White likewise shows no use of a PLL for multiple cores. Ex. 2027 at ¶40. None of the configurations proffered by the Petition disclose a PLL providing a final frequency to multiple cores. *See* Petition at 23. Rather, such a use would go against the teaching of White to configure "multiple logic cores [] to operate at independent voltage levels and frequencies." Ex. 1006 at 6:23-25, *see also id.* at 2:15-19 (same), 8:41-58 (same), 14:45-46 (same), 14:63-67 (same), Ex. 2027 at ¶40. Neither Talwar nor White expressly disclose the claimed PLLs.

To the extent the Petition argues the claimed PLLs are inherent, the Petition provides no support for that theory. A party suggesting inherency "must show that *the natural result flowing* from the operation as taught would result in the

performance of the questioned function.” *PersonalWeb Techs., LLC v. Apple, Inc.*, 917 F.3d 1376, 1382 (Fed. Cir. 2019) (emphasis original; internal citations and quotations omitted). This is the Petition’s burden but the Petition devotes only part of a parenthetical to identify these PLLs, much less explain that they are a “natural result” of anything taught. Petition at 27. Dr. Annavaram offers no better and merely concludes “frequency regulator 140 provides on clock signal to a local PLL shared by cores 120A and 120C and an independent clock signal to a separate, local PLL shared by cores 120B and 120D.” Ex. 1002 at ¶78. But White, which teaches the “local PLLs,” *only* uses a one-to-one arrangement of local PLLs and cores. Ex. 2027 at ¶41, Ex. 1006 at 9:62-10:10. Talwar, where a POSITA would more reasonably expect to derive a cluster architecture, provides *no* explanation for how to provide a cluster-wide frequency. Ex. 2027 at ¶41, *see* Ex. 1007. Only hindsight provides an explanation for why a cluster-wide PLL would be used, because it meets the limitation. The claimed PLLs are not inherent to either cited reference.

The Petition’s Grounds 1-3 fail to demonstrate a POSITA would find a first and second set of processor cores or the claimed first and second PLLs were obvious. As such, Grounds 1-3 do no show a reasonable likelihood of success on the merits.

VI. The Petition’s Ground 4 Also Fails

Grounds 4-8 fair even worse. Where Grounds 1-3 present a tenuous motivation to combine references to achieve sets of cores, Grounds 4-8 conjures this

limitation from thin air. Likewise, the first and second PLLs are alleged from a single reference to unorganized PLLs. The Petition must allege these were “known design choices,” but there is no evidence that these design choices were ever used, let alone known. Last, the proposed combination of Elgebaly and Ozer would be inoperable as the Petition describes it. With at least three elements entirely missing from the Petition’s disclosure, the Petition does not show a reasonable likelihood of success under grounds 4-8 as to any claim.

A. Ground IV Fails to Disclose Sets of Processor Cores

The Petition asks the Board to find sets of processor cores disclosed from unexplained reservation of rights provisions in Elgebaly. But the suggestion that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores,” does not disclose the claimed set, let alone first and second sets, of processor cores. Instead, a POSITA would have understood this disclosure to refer to the well understood and long-established chip-wide DVFS mode. The Petition’s parroted expert declaration citing to unsupported design motivations and art that does not even make up this ground cannot create or render obvious sets of processor cores.

Elgebaly teaches an ASIC that “includes various processing units... include[ing] DSP cores 130_a and 130_b processor cores 130_c and 130_d, a clock generator 140, an internal memory 150, an external interface unit 160, and other

circuits 170.” Ex. 1004 at [0027]. While the reference further teaches “ASIC 120 may include fewer or more and/or different processing units” and “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores,” such arrangements are never explained in detail. *Id.* at [0030], [0033]. Rather they are disclosed as “other embodiments” that are never explained or explored. *See id.*

The Petition asks the Board to take this vagueness as an express disclosure of the invention or at least a suggestion that it is obvious. Petition at 54-56. But that is not what a POSITA would have understood.

Rather a POSITA would have understood “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores,” as the standard chip-wide DVFS mode. Ex. 2027 at ¶45. At the time of the invention, two modes for controlling voltage and frequency were well known in the art: Chip-wide DVFS and per-core DVFS. *Id.* (citing Ex. 1019 at 52-53). Elgebaly itself teaches a per-core model for its first embodiment: “the clock for *each processing core* 130 may be set independently by clock generator 140, and the supply voltage for *each processing core* 130 may be set independently by power management unit 180.” Ex. 1004 at [0033], Ex. 2027 at ¶45. Thus, when faced with “other embodiments” where “cores may share a common clock and/or a common

supply voltage,” a POSITA would naturally apply the other most common DVFS mode that aligns with “common” clock/voltage, a chip-wide DVFS. Ex. 2027 at ¶45.

To the extent the Petition intends to invoke inherency, a party suggesting inherency “must show that *the natural result flowing* from the operation as taught would result in the performance of the questioned function.” *PersonalWeb Techs., LLC v. Apple, Inc.*, 917 F.3d 1376, 1382 (Fed. Cir. 2019) (emphasis original; internal citations and quotations omitted). The Petition has made no showing that sets of cores in a “natural result” of cores sharing a common clock and/or supply voltage. *See* Petition at 54-56. Indeed, it is not the natural result. Ex. 2027 at ¶46. As noted above, a POSITA would recognize chip-wide DVFS as the standard common clock/supply voltage DVFS option. As such chip-wide DVFS is the natural result, if any, from the disclosure of Elgebaly and sets of processor cores is no inherent.

In contrast, the Petition piles assumption upon assumption to reach its tenuous conclusion that a POSITA would ever implement multiple sets of processor cores. To even have enough cores for multiple sets², the Petition first assumes a POSITA would have implemented Elgebaly with four cores. Petition at 54. While the Petition fails to actual identify a motivation for this change, it seems to be simply an increase

² The Petitioner and Patent Owner agree that “set of processor cores” requires a group of two or more cores, as such to have a first and second set of cores, a multi-core processor must have a minimum of four cores.

in processing power. The Petition argues “each additional processor core provides a roughly linear increase in processing power, whereas the voltage and frequency flexibility of setting cores differently may give diminishing returns at the cost of valuable chip space for additional power generation and control circuitry, AVS units, and clock generation circuitry.” Petition at 55. But each additional processor core absolutely does not provide a linear increase in processing power. Ex. 2027 at ¶47. In reality, processing throughput is dependent on the computational resources, voltage, operating frequency, and the computational task provided at a given time. *Id.* Indeed, the Petition’s own supporting documents show that in some cases, a per-core DVFS arrangement can provide more than twice the processing power as a chip-wide DVFS. *Id.* (citing Ex. 1019 at 53). With that in mind, there is no clear motivation to take the two core, per-core DVFS embodiment taught by Elgebaly, and double the number of cores, but arrange them into sets.

Beyond, the Petition’s motivation to add cores being suspect, the proscribed motivation to arrange any cores into sets is also faulty. The Petition contends space constraints would naturally lead a POSITA to sets of cores. Petition at 55-56. But, while “chip space” is somewhat limited, it hardly motivates a POSITA to implement multiple sets of processor cores. Ex. 2027 at ¶48, *see also id.* at ¶¶34, 36-37. Rather, the known design techniques of the time pushed in different directions. Moreover,

the AVS units taught by Elgebaly are not compatible with the proposed clusters and would not achieve any of the benefits alleged by the Petition.

The first of these alternatives to address “chip space” concerns is to simply remove the voltage and clock systems from the chip. The entire motivation offered by the Petition to group the cores taught by Elgebaly is “that chip space is limited.” *See* Petition at 55. While cost is mentioned, the principle argument is that “the voltage and frequency flexibility of setting cores differently may give diminishing returns at the cost of valuable chip space...” *Id.* But while Elgebaly discloses its two core embodiment on a single ASIC, there is no reason power and voltage circuitry would not simply be implemented off chip. Ex. 2027 at ¶49. To begin, Elgebaly already discloses implementing its power management unit 180 off the ASIC chip. *See* Ex. 1004 at Fig. 2, [0029]. Moreover, this is the solution offered by White and discussed by Dr. Annavaram in his declaration. Ex. 1002 at ¶35, Ex. 1006 at 1:14-26. Space outside of the chip is far less precious when designing devices as the manufacturing techniques are less expensive.

Beyond White, as noted above, “chip space” is not and was not at the time of the invention, a real world concern that might override DVFS concerns. By the time of the patent, “the bulky off-chip inductors and capacitors can be reduced in size and moved onto the package and die respectively.” Ex. 1029 at 6. Real world chips implemented this with no discernable chip space wasted on frequency and voltage

components. *See* Ex. 2027 at ¶50, Ex. 2030, *see also* Ex. 2027 at ¶¶34, 36-37. In the real world, chip space is not a significant concern to govern the chosen DVFS mode. Ex. 2027 at ¶50.

Another alternative was to simply accept the downsides of either a chip-wide or per-core DVFS system. While the Petition notes “chip space” as a design consideration, it does not actually show a POSITA would ever weigh it more highly than the processing and power advantages of a particular DVFS. *See* Petition at 55-56. In reality, a POSITA would consider the expected processing load. Ex. 2027 at ¶51. For example, for some processing loads “[p]er-core DVFS achieves similar power savings as chip-wide DVFS for both extremes (all memory-bound or all cpu-bound).” Ex. 1029 at 4, Ex 2028 at ¶51. Alternatively, under other conditions, power saving can differ by nearly 20%, *id.*, or processing power can differ by greater than 200%, Ex. 1019 at 53. Ex 2028 at ¶51. In light of this, it is not chip space that governs the DVFS design choice, but expected processing load. *Id.* As the Petition fails to even consider, this dominant design consideration, it does not show the claimed first or second sets of processor cores were obvious.

In contrast, the Petition has not met the legal test for a “known technique” argument. The Petition contends “this type of core grouping was a well-known approach to balancing competing design factors” but cannot show that to be true. Petition at 55. While the Federal Circuit has accepted that “if a technique has been

used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill,” it has placed requirements on this “known-technique” rationale. *Intel Corp. v. PACT XPP Schweiz AG*, 61 F.4th 1373, 1380 (Fed. Cir. 2023). Centrally, the Federal Circuit asks “if there’s a known technique to address a known problem using ‘prior art elements according to their established functions.’” *Id.* But “core grouping” was not a known technique to address any problem, much less “balancing competing design factors.” Ex. 2027 at ¶52. The Petition can show only one reference that considers multiple groups of cores in a single multi-core processor, Talwar³. Petition at 55-56. One example made public less than four months before the ’339 Patent hardly establishes a known technique. *Id.*, see Ex. 1007. Likewise, Dr. Annavaram’s expert declaration on this point is nearly identical to the Petition. Dr. Annavaram’s and the Petitions “Sets of cores” sections (pages 54-55 and ¶¶144-45) differ by a handful of phrasing changes. Because, as noted above, those arguments presuppose sets of cores when none are taught or made obvious, Dr. Annavaram’s repetition of conclusory attorney argument should be given no weight. See *Xerox Corp. v. Bytemark, Inc.*, IPR2022-

³ Talwar is not raised as a formal part of this ground and the Petition offers no explanation for why it would be obvious to combine its teachings with Elgebaly or Ozer. See Petition at 53-56.

00624, Paper 9 at 15-17 (PTAB August 24, 2022) (designated precedential) (expert declaration that merely repeats “the conclusory assertion for which it is offered to support” “is entitled to little weight”). Because Elgebaly does not disclose multiple sets of cores and any relevant design considerations point to other solutions, the Petition does not show a first and second set of processor cores.

B. Ground IV Fails to Disclose First and Second PLLs

Much like its failure to disclose sets of processor cores above, Ground IV fails to disclose the First and Second PLLs associated with those sets of processor cores. Even assuming the first and second sets of processor cores, as the Petition does, the Petition cannot support the claimed PLLs are “known techniques.”

Elgebaly does not teach any particular clock/frequency architecture. Instead, Elgebaly provides an amorphous block, clock generator 140, that “may include one or more phase locked loops (PLLs) that control the oscillators used to generate the clocks.” Ex. 1004 at [0028]. Elgebaly teaches that each core provides a target frequency to the clock generator and the clock generate provides a clock signal at that frequency. *Id.* at [0032]. Nothing more is explained as to the structure of Elgebaly’s clock generator, just a function and a handful of components. Ex. 2027 at ¶54. Despite this, the Petition contends a reference to “oscillators” and PLLs means the claimed architecture is rendered obvious. Petition at 52-53. Not so. Ex. 2027 at ¶54.

The Petition offers that “multiple PLLs each receiving a distinct clock signal from a different external oscillator crystal and, in turn outputting one of the desired clock signals” would be a “known clock distribution channel” or an “obvious design choice.” Petition at 53. To support these conclusions, the Petition relies on the Declaration of Dr. Annavaram, *id*, but all Dr. Annavaram provides is a near copy of the petition. He offers some minor phrasing changes, “two PLLs” instead of “multiple PLLs,” “the clock for each core ‘may be set independently’” instead of “clocks ‘may be set independently’ for each core,” but it is still a mere parroting in substance. *Compare* Petition at 53 *with* Ex. 1002 at ¶¶141-42. Dr. Annavaram’s declaration is entitled to little weight. *See Xerox Corp. v. Bytemark, Inc.*, IPR2022-00624, Paper 9 at 15-17 (PTAB August 24, 2022) (designated precedential) (expert declaration that merely repeats “the conclusory assertion for which it is offered to support” “is entitled to little weight”). Conclusory attorney argument based on a conclusory expert declaration does not support an obviousness argument. *See also, Facebook v. Windy City Innovations, LLC*, 973 F.3d 1321, 1340-41 (Fed. Cir. 2020) (affirming expert testimony “did not ‘add materially to Petitioner’s unpersuasive attorney argument’ because it ‘merely repeated Petitioner’s argument, nearly verbatim, without citation to the basis for his testimony.’”)

Even if Dr. Annavaram’s “known technique” argument is considered, it is legally deficient. A “known technique” argument requires several showings. There

must be a “known technique to address a known problem using ‘prior art elements according to their established functions.’” *Intel Corp. v. PACT XPP Schweiz AG*, 61 F.4th 1373, 1380 (Fed. Cir. 2023). Neither the Petition nor Dr. Annavaram address any of these requirements. In particular, applying a single PLL to multiple cores is not an established function. Ex. 2027 at ¶55. The only evidence either the Petition or Dr. Annavaram can cite to is a Figure from a tertiary reference not part of any ground. Ex. 1002 at ¶¶141-42. But this Figure only provides how two clock domains might be created, not how they would be implemented or be applied to any arrangement cores. Ex. 2027 at ¶55, *see Id*, Ex. 1008, Fig. 1. But this is critical, the claims mandate a single PLL providing an output clock signal to multiple cores. But because this feature is absent from every single reference cited in the Petition, the Petition must rely on Dr. Annavaram to provide it in every ground. The Petition’s Grounds 4-8 cannot show the claimed first and second PLLs are obvious and thus cannot show any claim is invalid.

C. No Motivation or Ability to Combine Ozer

Additionally, the Petition proposes combining a heavily modified Elgebaly having to sets of cores operating at independent frequencies with Ozer to meet the “interface block” limitation. But, while Ozer discloses a caching circuitry that supports communication between cores, it would not allow for communication across clock domains. Specialized circuitry is needed to transmit information

between different core domains, but neither Elgebaly nor Ozer disclose any such circuitry. Ex. 2027 at ¶56. Without that circuitry, the combination of Ozer and Elgebaly would fail, a POSITA would have no expectation of success, and the combination is not obvious.

When different components operate at different clock frequencies they fundamentally cannot communicate. Ex. 2027 at ¶57, Ex. 1019 at 52 (“per core DVFS ... requires more than one power/clock domain per chip, and synchronizer circuits are required to avoid metastability between domains.”). To borrow a turn of phrase, they are operating on different wavelengths. Ex. 2027 at ¶57. The ’339 Patent’s method to address this is “synchronizers” in the interface blocks. *See* Ex. 1001 at 4:1-17, *see also id.* Cl 4. But the Petition can point to nothing similar and in fact does not challenge claim 4.

Indeed, neither Ozer nor Elgebaly teach any similar tool. As the Petition admits, Elgebaly “provides little detail on the implementation of shared memory.” Petition at 56. So little, in fact, that the concept that the memory is “shared” is not even disclosed. Ex. 1004 at [0028] (“Internal memory 150 stores data and program codes used by the processing units within ASIC 120.”) Elgebaly makes no suggestion that any data or program would be used across multiple cores at all much less across clock domains. Ex. 2027 at ¶58. In Ozer, there is simply no mention of

domains, clock frequencies, or any circuitry needed for domain crossings. Ex. 2027 at ¶58, *see* Ex. 1013.

The Petition offers no modification of Ozer or Elgebaly to address this. Instead, the Petition merely contends Elgebaly would be modified so “each Processor Cores 1-4 include local caches, [and] cache coherency fabric and protocols are implemented in ASIC 120 to facilitate sharing of data between and for use by Processor Cores 1-4 as detailed in Ozer.” Petition at 62. There is no support in the Petition, Ozer, or Dr. Annavaram’s declaration for “cache coherency fabric and protocols... detailed in Ozer” being able to cross clock domains and facilitate communication between cores. Ex. 2027 at ¶59. Without clock domain crossing, the combination has no chance of success and any speed gain or energy savings in Ozer could not be recognized. *Id.* The combination is not obvious.

The central combination of Grounds 4-7 fails to disclose sets of processor cores or the claimed PLLs and would simply not function as proposed by the Petition. Without these limitations or a combination that a POSITA would have any expectation of success in, the Petition has not shown a reasonable likelihood of success on Grounds 4-7 as to any claim.

VII. Conclusion

For the reasons set forth above, Patent Owner respectfully requests that the Board deny institution of the Petition. Both Grounds 1 and 4 offer faulty motivations

to combine and fail to disclose the claimed sets of processor cores or claimed PLLs.

Without these elements, the Petition has not shown a reasonable likelihood of success as to any claim and institution should be denied.

Date: June 12, 2025

Respectfully submitted,
/ Neil A. Rubin/
Neil A. Rubin, Reg. No. 67,030
Russ, August & Kabat
12424 Wilshire Blvd., 12th Floor
Los Angeles, CA 90025

Counsel for Patent Owner

CERTIFICATION REGARDING WORD COUNT

Pursuant to 37 C.F.R. §42.24(d), Patent Owner certifies that there are 5,873 words in the paper excluding the portions exempted under 37 C.F.R. §42.24(a)(1).

Date: June 12, 2025

Respectfully submitted,
/ Neil A. Rubin/
Neil A. Rubin, Reg. No. 67,030
Russ, August & Kabat
12424 Wilshire Blvd., 12th Floor
Los Angeles, CA 90025

Attorney for Patent Owner

CERTIFICATE OF SERVICE (37 C.F.R. § 42.6(e)(1))

The undersigned hereby certifies that the above document was served on June 12, 2025, by filing this document through the Patent Trial and Appeal Board Cast Tracking System as well as delivering a copy via electronic mail upon the following attorneys of record for the Petitioner:

Eagle H. Robinson
eagle.robinson@nortonrosefulbright.com
Norton Rose Fulbright US LLP
98 San Jacinto Blvd., Ste. 1100
1550 Lamar Street, Ste. 2000
Austin, TX 78746
Houston, TX 77010-4106
512.536.3083 (telephone)
713.651.8360 (telephone)

Daniel S. Leventhal
daniel.leventhal@nortonrosefulbright.com
Eric C. Green
eric.green@nortonrosefulbright.com
G. Duncan Waldrop
duncan.waldrop@nortonrosefulbright.com
Norton Rose Fulbright US LLP
98 San Jacinto Blvd., Ste. 1100
Austin, TX 78746
512.536.3009 (telephone)

Richard S. Zembek
richard.zembek@nortonrosefulbright.com
Norton Rose Fulbright US LLP
1550 Lamar Street, Ste. 2000
Houston, TX 77010-4106
713.651.5283 (telephone)

Date: June 12, 2025

Respectfully submitted,
/ Neil A. Rubin/
Neil A. Rubin, Reg. No. 67,030
Russ, August & Kabat
12424 Wilshire Blvd., 12th Floor
Los Angeles, CA 90025

Attorney for Patent Owner