

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

APPLE INC.,
Petitioner,

v.

REDSTONE LOGICS LLC,
Patent Owner.

Case IPR2025-01532
Patent 8,549,339

PATENT OWNER'S PRELIMINARY RESPONSE

TABLE OF CONTENTS

I.	Introduction	1
II.	Background	2
	A. Apple’s Petition	2
	B. Overview of the ’339 Patent	3
III.	Level of Ordinary Skill in the Art.....	4
IV.	Claim Construction	5
V.	The Petition’s Ground 1 Fails	5
	A. No Motivation to Combine White, Talwar, and/or Lee.....	5
	B. First and Second Supply Voltage Are Not Disclosed.....	13
	C. First and Second PLL and Clock Signals as Input Are Not Disclosed In the Art.....	15
VI.	The Petition’s Grounds 2 and 3 Also Fail.....	18
	A. Grounds 2 and 3 Fails to Disclose Sets of Processor Cores	18
	B. No First and Second Clock Signals or PLLs	22
VII.	Conclusion	24

TABLE OF AUTHORITIES

Cases

<i>PersonalWeb Techs., LLC v. Apple, Inc.</i> , 917 F.3d 1376 (Fed. Cir. 2019)	16, 20
---	--------

PATENT OWNER'S EXHIBIT LIST

No.	Description
Ex. 2001	<i>Redstone Logics LLC v. Apple Inc.</i> , No. 7:25-cv-00183-ADA, Dkt. No. 25 Scheduling Order
Ex. 2002	Motion to Stay before Judge Albright, Docket Navigator
Ex. 2003	<i>Redstone Logics LLC v. Apple Inc.</i> , No. 7:25-cv-00183-ADA, Dkt. No. 1
Ex. 2004	Time to Trial Statistics for Small Cases Before Judge Albright, Docket Navigator
Ex. 2005	Time to Trial Statistics for All Cases Before Judge Albright, Docket Navigator
Ex. 2006	Western District of Texas Case Management Statistics
Ex. 2007	<i>Redstone Logics LLC v. Apple Inc.</i> , No. 7:25-cv-00183-ADA, Dkt. No. 30, Claim Construction Order
Ex. 2008	Redstone's Preliminary Infringement Contentions
Ex. 2009	Apple's Preliminary Invalidity Contentions
Ex. 2010	RIN 0651-AD89 Revision to Rules of Practice before the Patent Trial and Appeal Board, Dkt. No. PTO-P02025-0025
Ex. 2011	Interim Processes for PTAB Workload Management
Ex. 2012	Jon Stokes, "Two billion-transistor beasts: Power7 and Niagara 3"
Ex. 2013	Jinuk Shin, <i>et al</i> "A 40nm 16-Core 128-Thread SPARC Soc Processor"

I. Introduction

The instant petition largely copies the last, now-denied petition, IPR2025-00485, to challenge the '339 Patent. Despite limited attempts to address the deficiencies identified in Redstone's last patent owner preliminary response, Apple's instant Petition still fails to show a reasonable likelihood of success. Rather, the Petition presumes any missing limitation to be obvious without making the requisite showing of obviousness 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 contradicts 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 26 (emphasis in original); 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 26 (emphasis added). 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 is not taught or even suggested by Elgebaly, with or without Lee.

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. Apple's Petition

Petitioner Apple Inc. (“Apple” or “Petitioner”) 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: White, Talwar, and Lee Render Obvious Claims 1–3, 5, 8, 10–11, 14, and 21
- Ground 2: Elgebaly and Lee Render Obvious Claims 1–3, 5, 8, 10–11, 14, and 21
- Ground 3: Elgebaly, Lee, and Vajda Render Obvious Claim 9

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 13–14. 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. 2007. This construction is helpful, though not necessary, to understanding both the Petition and this preliminary response.

V. The Petition’s Ground 1 Fails

Ground 1 fails to justify the central combination of White and Talwar and cannot support the further inclusion of Lee to support its contentions regarding PLLs and clock signals.

A. No Motivation to Combine White, Talwar, and/or Lee

Ground 1 relies on the primary combination of White and Talwar, in particular, the architecture of White with the clusters taught by Talwar. *See* Petition at 28–30. But the motivation for such a combination enumerated by the Petition is a “need” that contradicts the express teachings of White destroying any advantage of following the teachings of White that the Petition supports with a serious misunderstanding of White’s applicability. Moreover, the Petition offers no basis to

add an even further reference to the already tenuous combination of White and Talwar.

The primary 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, and the importance to minimize overhead and devote as much of the chip as possible to additional cores.” *Id.* at 26 (emphasis added). But this “*need*” does not exist. In fact, White itself explains the device it teaches may already “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 already be configured to operate with their own independent voltages and frequencies, there is no *need* for further adapting to 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.

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.” 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.

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. Horst to create this need and rewrite White, but his support actually demonstrates why Talwar’s “clusters” would not be used. *See id.* at ¶¶85–92. While Dr. Horst 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 ¶88.

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 ¶88 (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. *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. *See* 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. 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. A POSITA would therefore not find any savings associated with grouping cores to motivate a change to White.

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. For example, the Niagara or UltraSPARC T3 announced at the time of the invention, *see* Ex. 2012, wastes no recognizable chip space on voltage and clock control:




Figure 5.2.6: Die Micrograph.

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). The other major components are interface circuitry like PCIe or the CLC and coherency components that address transactions within the chip. *See* Ex. 2013 at 2. In the real world, chip space is not a significant enough concern to govern the chosen DVFS mode.

Exemplary of the Petition's misunderstanding of the state of the art at the time of the invention is Petitioner's contentions that White and Talwar achieve the same

goal: “performing DVFS internal to the chip...” Pet. at 26–27. That is not White’s goal. White’s goal is protecting expensive components from damage due to overheating. *See* Ex. 1006 at 1:27–52, 2:29–44. Indeed, White is unconcerned with whether DVFS is performed internally or “off-chip.” *See id.* at 5:10–16 (“As with the voltage regulators and any accompanying signal dividing or adjusting circuitry discussed above, clock generators and accompanying signal dividing or translating circuitry may reside as part of, **external** to, or a combination of internal to and external to, the processor, according to various embodiments.”) (emphasis added); 5:17–24 (similar); 8:34–37(similar). Given the Petition offers no support for its contention as to White’s “goal” and indeed the Petition is wrong as to White’s goal, this offers no motivation to combine White and Talwar.

Lastly, the Petition 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. 1006 at 8:44–49. But White explains that the application to “any integrated circuit does not relate to cores or clusters of cores as would be required in combination with Talwar. White provides “[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. A POSITA would not be able to take this teaching of White to create the clusters referenced in Talwar.

Exacerbating the Petition's already unsupportable combination of White and Talwar is the Petition's addition of "Lee's further implementation details on clocking, synchronization, level shifting, and chip layout..." The Petition contends "[a] POSITA would look to Lee for specifics on generating clock signals for each cluster/domain, implementing the clocks with PLLs from independent clock inputs per cluster/domain, communicating between the cluster with synchronizers and level shifters, and laying out the clusters in rows on an integrated circuit." Incredibly, despite such a list of "specifics" a POSITA would derive from Lee, the Petition does not actually cite any teaching from Lee nor any motivation outside the Petition's self-serving expert declaration. *See* Pet. at 27–28. Indeed, Dr. Horst declaration likewise cites nothing from Lee when discussing the motivation to combine or indeed the actual combination itself. *See* Ex. 1002 at ¶¶85–92 (citing nothing from Lee). Without a single citation to Lee, the Petition cannot show a motivation to combine Lee with White and Talwar much less what that combination actually is. Because the Petition cannot support the combination of Ground 1, it does not show a reasonable likelihood of success on Ground 1.

Beyond the Petition's utter lack of any actual motivation to combine Lee with White/Talwar, a POSITA would be uninterested in the teachings of Lee in view of White/Talwar. Indeed, White already provides a clocking architecture. *See* Ex. 1006 at 7:38–45. Beyond meeting the first and second PLL limitations, the Petition does not identify any reason a POSITA would abandon the teachings of White.

The Petition has failed to demonstrate any motivation to change White either with Talwar or Lee. Rather, the Petition's 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 or the claimed PLLs would have been obvious and there is no likelihood of success on Ground 1.

B. First and Second Supply Voltage Are Not Disclosed

Claim 1 claims “wherein each processor core from the [first/second] set of processor cores is configured to dynamically receive a [first/second] supply voltage...” The Petition fails to identify where the first/second supply voltage that each core of the first/second set of processor cores is dynamically configured to receive is taught, rendered obvious, or is inherent from the prior art.

For these limitations, the Petition cites White's and Lee's disclosure of clock and voltage distribution with no mention of Talwar beyond “multiple cores per cluster.” But neither White nor Lee teach providing a single supply voltage to a set

of cores as claimed much less a configuration to dynamically receive said voltage.¹ Rather, White expressly teaches “when adjusting the operating characteristics for logic core 120A, power management logic 200 may be configured to ... instruct voltage regulator 120 to supply a particular voltage to logic core 120A.” Ex.1006, 10:15–20. In White, logic core 120A is a singular core. *Id.* There is no teaching in White of providing a singular, dynamic, voltage to multiple cores. Indeed, the logic taught by White is only applicable to controlling the voltage to the entire circuit or individual cores, not a particular set of cores, and no configuration is taught for a set of cores. Ex. 1006 at 13:24–63.

Likewise, Lee teaches only providing a voltage to a “domain” based on “the performance requirement of each domain.” Ex. 1010 at 1. There is no reference in Lee as to the “domains” containing multiple cores or explanation for how such cores could be “configured to dynamically receive a [first/second] supply voltage.” Indeed, Lee does not teach any configuration of any core or set of cores. The Petition seems to concede this and argues “[a] POSITA would understand those domains as disclosed by Lee could include *first set of processor cores* such as White’s ‘any

¹ For that matter, though not cited at all for this matter, Talwar teaches using an MPM to provide voltage to its “clusters” and provides for a different power and voltage management scheme than White or Lee. *See e.g.* 1007 at [0018]-[0020]. The Petition offers no discussion that the divergent techniques between Talwar and White or Lee are compatible.

number of processor cores’ or Talwar’s cluster cores.” Pet. at 34 (citing Ex.1002, ¶105). But that conclusion fails to show *how* a POSITA would implement the supposedly obvious combination. Even considering Dr. Horst’s discussion, he offers the same conclusion with no explanation of how a POSITA might have configured multiple cores within a set of cores to dynamically receive the same voltage.

Because the Petition cannot cite any reference that discloses multiple cores within a set of cores configured to dynamically receive a single voltage and makes no effort to show that such would be obvious, the Petition fails on this limitation.

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

Even if the White/Talwar/Lee 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 “Lee’s PLL design supplied additional detail on how to generate a clock for each cluster.” Pet. at 29. But Lee never considers any clusters. Moreover, the Petition contends similar of White, that the “clock control circuitry that White describes that the core level (e.g., a PLL at the core and accompanying ‘control registers’ storing ‘PLL configuration parameters’) would operate at the cluster level.” *Id.* at 30. 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. *See* Ex. 1006. The same is true of Lee, which only considers vague “domains” with no discussion of multiple cores within said domain.² *See* Ex. 1010. None of the configurations proffered by the Petition disclose a PLL providing a final frequency to multiple cores. *See* Petition at 28-30. Rather, such a use would go against the teaching of the base reference, 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). Neither Talwar nor White nor Lee 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

² While the Petition and Dr. Horst both conflate a cluster with a domain, neither supports this conflation. *See* Pet. at 29 (“a multi-cluster (multi-domain) DVFS”); *see also* Ex. 1002 at ¶98 (same). The circuitry needed to support different kinds of domains is different, that is the entire subject of the ’339 Patent, yet the Petition would ignore this.

quotations omitted). This is the Petition's burden but the Petition devotes only part of a paragraph to identify these PLLs, much less explain that they are a "natural result" of anything taught. Petition at 29-30. Dr. Horst offers little better. He contends because his recreation of claim 1 has "each core in a cluster operat[ing] at the same voltage, there is a corresponding ideal frequency for the entire cluster" and "a POSITA would have understood that the internal clock signal for each core, as disclosed by White, should be shared among the cluster cores operating at the same voltage level to avoid the need for additional, redundant clock circuitry." Ex. 1002 at ¶94. First, this should be ignored as not part of the Petition and because Dr. Horst offers no support for his contention that such a design constraint existed. Second, this explanation fails to demonstrate that a POSITA would actually resolve any competing design constraints on providing the claimed PLL. Rather, Dr. Horst makes no showing or even suggestion that sharing a PLL between cores was plausible at the time or worth the alleged savings. Indeed, the *only* showing in any of the art cited by the Petition always uses a 1-to-1 arrangement of cores and PLLs. *See* Ex. 1006 at 9:62–65. Moreover, Talwar, where a POSITA would more reasonably expect to derive a cluster architecture, provides *no* explanation for how to provide a cluster-wide frequency. *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 Ground 1 fails to demonstrate a POSITA would find a first and second set of processor cores, the claimed voltage, or the claimed first and second PLLs were obvious. As such, Ground 1 does not show a reasonable likelihood of success on the merits.

VI. The Petition's Grounds 2 and 3 Also Fail

Grounds 2 and 3 fair even worse. Where Ground 1 presents a tenuous motivation to combine references to achieve sets of cores, Grounds 2 and 3 conjure this limitation from thin air. Likewise, Grounds 2 and 3 fail to identify clock signals as the first and second clock signals and instead identifies data signals. Last, the first and second PLLs are alleged obvious from Lee and Elgebally, but neither disclose or render obvious a single PLL feeding a set of cores. With at least three elements entirely missing from the Petition's disclosure, the Petition does not show a reasonable likelihood of success under Grounds 2 and 3 as to any claim.³

A. Grounds 2 and 3 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

³ Ground 3 only asserts invalidity of dependent claim 9 based on its argument in Ground 2, as the Petition cannot show a reasonable likelihood of success as to the limitations of the independent claims, Ground 3 falls with Ground 2.

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 52–58. 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. 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. *See* 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]. 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.

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 52–58, 59–60. Rather, the Petition and Dr. Horst simply conclude “Elgebaly may be extended to ‘multiple processing cores’ in which the clock signal and supply voltage are ‘jointly set for these cores’” or that from Elgebaly’s statement that it “may include ... any number of processor cores” “[a] POSITA would have

understood this disclosure to mean that Elgebaly's voltage and clock frequency control methods could be duplicated to support adaptive voltage and frequency control of up to N processors on an ASIC chip, where N is limited only by other chip design constraints." *Id.* at 56, 59. Such bald conclusions do not meet the Petitioner's burden. Moreover, it is not the natural result. 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 even have the necessary components to reach its tenuous conclusion that a POSITA would ever implement multiple sets of processor cores. To even have enough cores for multiple sets⁴, the Petition first assumes a POSITA would have implemented Elgebaly with four cores. *See* Petition at 56–57. But the Petition offers no reason why additional cores would be used. *See id.* Instead, the Petition merely cites portions of Elgebaly stating that the ASIC “may include ... any number of processor cores.” *Id.* at 56 (citing Ex. 1004 at [0030]). The only design motivation cited by the Petition is energy saving to extend battery life. *Id.* But this does not motivate a POSITA to add

⁴ 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.

cores to a device. Indeed, additional cores would generally increase, not decrease, energy consumption.

To the extent the Board would consider a motivation to add additional processing power while reducing power consumption, though not proposed by the Petition, such is not availing. Indeed, a lay person may assume the processing throughput of a device is directly correlated with the number cores in the device. This is not so. In reality, processing throughput is dependent on the computational resources, voltage, operating frequency, and the computational task provided at a given time. 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. *See* Ex. 1019 at 53. Applied here, there could be no gain whatsoever in taking the two core, per-core DVFS embodiment taught by Elgebaly, and double the number of cores, but arrange them into sets. By failing to consider the real-world intricacies involved in adding cores into a DVFS, the Petition fails to demonstrate multiple sets of processor cores would be obvious.

B. No First and Second Clock Signals or PLLs

The Petition identifies "Target Freq 3" and "Target Freq 4" as the "first clock signal as input" and the "second clock signal as input." *See* Pet. at 63, 65. However, as taught by Elgebaly, a "target frequency" is not a clock signal, but rather it is data about a frequency. For example, Elgebaly teaches "[e]ach core 130 provides to clock

generator 140 and to its AVS unit a *target frequency...*” Ex. 1004 at [0032] (emphasis added). A core does not create a clock signal. If it did, there would be no need for the clock generator to receive this target. Indeed, in response to receiving this target, “clock generator 140 *generates the clock* at the target frequency...” *Id.* (emphasis added). If the “target frequency” were a clock signal itself, this would not make grammatical sense. The clock generator would generate the clock signal at a clock signal. Instead, this statement makes sense only if “target frequency” is data about the frequency required by the core, so the clock generator generates the clock at the frequency required by the core.

Likewise, the Petition makes no showing that “each processor core ... is configured to dynamically receive ... a [first/second] output clock signal of a [first/second] phase lock loop (PLL).” Instead, the cited PLLs are disclosed as providing an output clock signal to a single core or “domain.” *See* Ex. 1004 at [0032], Fig. 2; Ex. 1010 at 2, Figs. 2, 4. Indeed, the Petition offers no support that each core is configured to dynamically receive a clock signal from a single PLL. *See* Pet. at 64. Rather, it contends because Elgebaly teaches “[e]ach core 130” provides a target frequency to a clock generator to receive a clock signal, it teaches the “configured to dynamically receive” limitation. *Id.* But such a teaching is entirely incompatible with multiple cores, as the parties agree is required by a set of cores, being dynamically configured. How would a POSITA address conflicting target

frequencies between multiple cores? How would a POSITA derive an acceptable target frequency for differing needs across cores? Why would a POSITA be motivated to arrive at a particular solution? The Petition has no answers for these questions and cannot show this limitation of one PLL to multiple cores is obvious.

Because what the Petition identifies as the first and second clock signal are not even clock signals and the Petition ignores that a set of cores involves multiple cores, the Petition cannot show limitation 1.1.1, 1.2.1, or 1.2.3 are anticipated or rendered obvious.

VII. Conclusion

For the reasons set forth above, Patent Owner respectfully requests that the Board deny institution of the Petition. Grounds 1, 2, and 3 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: December 30, 2025

Respectfully submitted,

/ Reza Mirzaie

Reza Mirzaie, Reg. No. 69,138

Russ, August & Kabat

12424 Wilshire Blvd., 12th Floor

Los Angeles, CA 90025

Attorney for Patent Owner

CERTIFICATION REGARDING WORD COUNT

Pursuant to 37 C.F.R. §42.24(d), Patent Owner certifies that there are 5,488 words in the paper excluding the portions exempted under 37 C.F.R. §42.24(a)(1).

Date: December 30, 2025

Respectfully submitted,

/ Reza Mirzaie

Reza Mirzaie, Reg. No. 69,138

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 December 30, 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:

Alex S. Yap, Reg. No. 60,609
AYap@mofo.com
Mehran Arjomand, Reg. No. 48,231
MARjomand@mofo.com
Jean Nguyen, Reg. No. 71,051
JNguyen@mofo.com
MORRISON & FOERSTER LLP
707 Wilshire Boulevard, Suite 6000
Los Angeles, CA 90017-3543

Bryan Blumenkopf, Reg. No. 73,282
BBlumen@mofo.com
MORRISON & FOERSTER LLP
12531 High Bluff Drive, Suite 200
San Diego, CA 92130

Apple-Redstone-IPR@mofo.com

Date: December 30, 2025

Respectfully submitted,

/ Reza Mirzaie

Reza Mirzaie, Reg. No. 69,138
Russ, August & Kabat
12424 Wilshire Blvd., 12th Floor
Los Angeles, CA 90025

Attorney for Patent Owner