

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

APPLE INC.,
Petitioner,

v.

REDSTONE LOGICS LLC,
PATENT OWNER.

Patent No. 8,549,339

Inter Partes Review No. IPR2025-01532

**PETITION FOR *INTER PARTES* REVIEW
UNDER 35 U.S.C. §§ 311-319 AND 37 C.F.R. § 42.100 *et seq.***

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. DISCRETIONARY DENIAL	3
III. OVERVIEW OF THE TECHNOLOGY.....	4
A. Dynamic Voltage and Frequency Scaling (DVFS).....	4
B. Multi-core Processors and DVFS.....	4
C. Clock Generation.....	5
D. Processor Cache Memory and Cache Coherency	9
IV. THE '339 PATENT	11
A. Summary	11
B. Prosecution History	12
C. Level of Ordinary Skill in The Art.....	13
D. Priority Date	14
V. CLAIM CONSTRUCTION	14
VI. PRECISE REASONS FOR RELIEF REQUESTED	15
A. Ground 1: White, Talwar, and Lee Render Obvious Claims 1-3, 5, 8, 10-11, 14, 21.....	16
1. White	16
2. Talwar	20
3. Lee.....	21
4. Motivation to Combine	26
5. Modification of White in View of Talwar and Lee	28
6. Element-by-Element Analysis	30
B. Ground 2: Elgebaly and Lee Render Obvious Claims 1-3, 5, 8, 10-11, 14, 21.....	51
1. Elgebaly	51
2. Motivation to Combine Elgebaly with Lee.....	53
3. Element-by-Element Analysis	59

C. Ground 3: Elgebaly, Lee, and Vajda Render Obvious Claim 9.....77

1. Vajda 77

2. Motivation to Combine Vajda with Elgebaly and Lee 78

3. Element-by-Element Analysis 79

VII. NOTICES AND STATEMENTS.....79

VIII. CONCLUSION.....82

Exhibit List for *Inter Partes* Review of USP 8,549,339

Exhibit Description	Exhibit #
U.S. Patent No. 8,549,339 (“339 Patent”)	1001
Declaration of Dr. Robert Horst (“Horst Declaration”)	1002
<i>Curriculum Vitae</i> of Dr. Robert Horst (“Horst CV”)	1003
U.S. Patent App. Pub. No. 2007/0096775 (“Elgebaly”)	1004
File History of U.S. Patent No. 8,549,339 (“339 Prosecution History”)	1005
U.S. Patent No. 7,263,457 (“White”)	1006
U.S. Patent App. Pub. No. 2009/0271646 (“Talwar”)	1007
Williams et al., “Low Latency Clock Domain Transfer for Simultaneously Mesochronous, Plesiochronous and Heterochronous Interfaces,” 13 th IEEE International Symposium on Asynchronous Circuits and Systems, 2007. (“Williams”)	1008
U.S. Patent App. Pub. No. 2008/0209133 (“Ozer”)	1009
Jeabin Lee et al., “Dynamic Voltage and Frequency Scaling (DVFS) Scheme for Multi-Domains Power Management,” IEEE Asian Solid-State Circuits Conference, 2007. (“Lee”)	1010
U.S. Patent App. Pub. No. 2012/0060170 (“Vajda”)	1011
RESERVED	1012
EX. 1002, Declaration of Murali Annavaram, NXP USA, Inc. and Qualcomm Incorporated v. Redstone Logics LLC, IPR 2025-00485	1013
IEEE Webpage for Williams et al., “Low Latency Clock Domain Transfer for Simultaneously Mesochronous, Plesiochronous and Heterochronous Interfaces,” 13 th IEEE International Symposium on Asynchronous Circuits and Systems, 2007. https://ieeexplonorthbridre.ieee.org/document/4137045 (accessed	1014

Exhibit Description	Exhibit #
August 30, 2025). (“Williams Paper IEEE Webpage”)	
IEEE Webpage for Jeabin Lee et al., “Dynamic Voltage and Frequency Scaling (DVFS) Scheme for Multi-Domains Power Management,” IEEE Asian Solid-State Circuits Conference, 2007. https://ieeexplore.ieee.org/document/4425705 (accessed August 30, 2025). (“Lee Paper IEEE Webpage”)	1015
IEEE Webpage for Jeabin Lee et al., “A Power Management Unit with Continuous Co-Locking of Clock Frequency and Supply Voltage for Dynamic Voltage and Frequency Scaling,” 2007 IEEE International Symposium on Circuits and Systems (ISCAS). https://ieeexplore.ieee.org/document/4253087 (accessed September 3, 2025). (“Lee II Paper IEEE Webpage”)	1016
RESERVED	1017
RESERVED	1018
S. Kaxiras, Computer Architecture Techniques for Power-Efficiency, 2008. (“Kaxiras Textbook”)	1019
U.S. Patent No. 8,312,225 (“Fryman”)	1020
U.S. Patent No. 7,584,369 (“Capps”)	1021
U.S. Patent No. 7,777,330 (“Pelley”)	1022
R. Ramanathan, Intel multi-core processors: Making the move to quad-core and beyond, Technology@Intel Magazine, December 2006. (“Intel Multi-Core Processor White Paper”)	1023
Y. Panov, Design Considerations for 12-V/1.5-V, 50-A Voltage Regulator Modules, Power Electronics, IEEE Transactions on. 16. 10.1109/APEC.2000.826080. (“Panov Paper”)	1024
U.S. Patent App. Pub. No. 2005/0034002 (“Flautner”)	1025
J.S. Liptay, Structural aspects of the System/360 Model 85, II. The cache, IBM Systems Journal, Vol. 7, No. 1, 1968. (“Liptay 1968”)	1026

Exhibit Description	Exhibit #
U.S. Patent App. Pub. No. 2009/0063888 (“Gold”)	1027
U.S. Patent No. 7,538,625 (“Cesky”)	1028
W. Kim, System Level Analysis of Fast, Per-Core DVFS using On-Chip Switching Regulators, 14 th International Symposium on High Performance Computer Architecture, February 2008. (“Kim Paper”)	1029
P. Senstrom, A Survey of Cache Coherence Schemes for Multiprocessors, Computer, Volume 23, Issue 6, June 1990. (“Senstrom Paper”)	1030
Jeabin Lee et al., “A Power Management Unit with Continuous Co-Locking of Clock Frequency and Supply Voltage for Dynamic Voltage and Frequency Scaling,” 2007 IEEE International Symposium on Circuits and Systems (ISCAS). (“Lee II”)	1031
U.S. Patent No. 4,873,703 (“Crandall”)	1032
U.S. Patent App. Pub. No. 2004/0210738 (“Kato”)	1033
C. Tran et al., “Low-power high-speed level shifter design for block-level dynamic voltage scaling environment,” International Conference on Integrated Circuit Design and Technology, 2005. (“Tran Level Shifter Paper”)	1034
WO 2010/138031 (“Vajda PCT”)	1035
RESERVED	1036
U.S. Patent No. 10,482,952 (“Campbell”)	1037

Claim Listing

Claim	Limitation #	Limitation
1	[1.pre]	A multi-core processor, comprising:
	[1.1]	a first set of processor cores of the multi-core processor,
	[1.1.1]	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;
	[1.2]	a second set of processor cores of the multi-core processor,
	[1.2.1]	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,
	[1.2.2]	wherein the first supply voltage is independent from the second supply voltage,
	[1.2.3]	and the first clock signal is independent from the second clock signal; and
	[1.3]	an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,
	[1.3.1]	wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.
2	[2.1]	The multi-core processor of claim 1, wherein the interface block further comprises a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set of processor cores for a first signal traveling from the first set of processor cores to the second set of processor cores.
3	[3.1]	The multi-core processor claim 1, wherein the interface block further comprises a second level

		shifter that is referenced to the first supply voltage and adapted to translate second logic levels associated with the second set of processor cores to first logic levels associated with the first set of processor cores for a second signal traveling from the second set of processor cores to the first set of processor cores.
5	[5.1]	The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.
8	[8.1]	The multi-core processor of claim 1, wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor.
9	[9.1]	The multi-core processor of claim 8, wherein the first region and the second region are overlapping regions of the multi-core processor.
10	[10.1]	The multi-core processor of claim 8, wherein the first region and the second region are non-overlapping regions of the multi-core processor.
11	[11.1]	The multi-core processor of claim 8, wherein the first region corresponds to a first row of the multi-core processor, and wherein the second region corresponds to a second row of the multi-core processor.
14	[14.1]	The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores.
21	[21.pre]	A multi-core processor, comprising:
	[21.1]	a first set of processor cores of the multi-core processor,

	[21.1.1]	wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block;
	[21.2]	a second set of processor cores of the multi-core processor,
	[21.2.1]	wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage from the power control block and a second output clock signal from a second PLL having a second clock signal as input in the clock control block,
	[21.2.2]	wherein the first supply voltage is independent from the second supply voltage, and
	[21.2.3]	the first clock signal is independent from the second clock signal; and
	[21.3]	an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,
	[21.3.1]	wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.

I. INTRODUCTION


Petitioner Apple Inc., (“Petitioner” or “Apple”) petitions for inter partes review of claims 1-3, 5, 8-11, 14, and 21 (the “Challenged Claims”) of U.S. Patent No. 8,549,339 (“the ’339 Patent” (EX1001)) in accordance with 35 U.S.C. §§ 311-319 and 37 C.F.R. § 42.100 *et seq.* Each of the Challenged Claims is unpatentable in view of the prior art, as shown below and supported by the Declaration of Dr. Robert Horst (EX1002).

The Challenged Claims relate to a “multi-core processor” having a “first” and “second” “set of processor cores,” each “configured to dynamically receive a [first] / [second] supply voltage and a [first] / [second] output clock signal of a first phase lock loop (PLL) having a [first] / [second] clock signal as input.” EX1001, claim 1. They further recite that “the first supply voltage is independent from the second supply voltage, and the first clock signal is independent from the second clock signal.” *Id.* Both multi-core processors and applying Dynamic Voltage and Frequency Scaling (DVFS) to sets of cores in a multi-core processor were well known prior to the Challenged Claims’ earliest priority date.


White discloses a multi-core processor where each core “may be configured to operate at frequencies and voltages *independently of* the other logic cores.” EX1006, 2:15-44. Talwar discloses “divid[ing]” the cores of a multi-core processor “into clusters” where each “cluster receive[s] power from a single voltage source”

and “[e]ach cluster may have a different cluster frequency.” EX1007, [0011], [0018]-[0019]. Lee discloses a “DVFS scheme for Multi-power domains” each with independent power in which “the voltage and frequency of each power domain are dynamically scaled according to the performance requirement of each domain.” EX1010, 1. Further, Lee discloses usage of a domain converter with level shifter (LS) and synchronizer (Sync.) for communicating between the three different domains. EX1010, 2. Lee shows that a “PMU [power management unit] generates the clock and the supply voltage” to each of the domains, and that the clocks in the PMU are generated by a phase locked loop (PLL) structure. *Id.*, 2-3. Applying Talwar’s and Lee’s teachings to White therefore renders obvious the challenged independent claims, as detailed below.


Likewise, Elgebaly discloses “the clock for each processing core 130 may be set *independently*” “and the supply voltage for each processing core 130 may be set *independently*.” EX1004, [0033]. Elgebaly further discloses that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” *Id.* Accordingly, Elgebaly (along with the various teachings of Lee) renders obvious the challenged independent claims as detailed below.



EX1010, FIG. 2 (modified).



EX1004, FIG. 2 (modified).



EX1006, FIG. 2 (modified).

II. DISCRETIONARY DENIAL

Pursuant to the guidance provided in the USPTO's "FAQs for Interim Processes for PTAB Workload Management", Petitioner omits discussion here of issues related to discretionary denial. (See <https://www.uspto.gov/patents/ptab/faqs/interim-processes-workload-management>, accessed September 10, 2025 ("The petition should not address discretionary issues. A petitioner should raise any discretionary issues in its opposition to a

patent owner’s discretionary denial brief”).)

III. OVERVIEW OF THE TECHNOLOGY

A. Dynamic Voltage and Frequency Scaling (DVFS)

The amount of work a processor core can accomplish in a given period of time depends on clock rate, with higher clock rates resulting in more work. EX1002, ¶46. But a higher clock rate comes with costs—voltage must be increased to support the clock rate, increasing both power drain and heat. *Id.* DVFS, “dynamic adjustments to supply voltage, clock frequency, or both,” was well known to increase clock rate and voltage when needed for higher workloads, but also decrease them when possible to reduce power consumption and heat. EX1019, 35; EX1002, ¶¶46-47. By 2010, DVFS had already become “[t]he most popular technique for CPU power management.” EX1007, [0002].

B. Multi-core Processors and DVFS


Multi-core processors were also well known before the priority date of the ’339 Patent. EX1002, ¶34; EX1019, 52-53. By 2008, multi-core processors were already “becom[ing] the predominant general-purpose, high-performance microprocessor platform.” EX1019, 52; . For example, dual- and quad-core processors were commercially available as early as 2006 and processors with 8, 16, and 32 cores were being designed and studied throughout the 2000s. *See e.g.*, EX1023, 2; EX1007, FIG. 1, [0014]; EX1020, 1:16-21; EX1021, 2:53-57, FIG. 1; EX1022, FIG. 1. Multi-core processors “are typically used to gain higher

performance by arranging the different processor cores to execute respective data processing operations in parallel” by “shar[ing] access to shared memory.” EX1009, [0004]; EX1002, ¶35.

By 2010, “[i]ndependent DVFS for each core,” cluster or “domain”-based DVFS, and “per-chip DVFS” had already become well known. EX1019, 52-53; EX1006, 2:15-44; EX1004, [0032]; EX1002, ¶36. These options all presented known trade-offs: per-core DVFS enables the most precise adjustments and most power and heat efficiencies, but “per-core DVFS is considered more expensive; it requires more than one power/clock domain per chip, and synchronizer circuits are required to avoid metastability between domains.” EX1019, 52-53; EX1002, ¶36. In addition to such a “[clock] synchronisation module,” it was well known that a “voltage level shifter” is required “to deal with clock synchronisation issues and the different supply voltage levels (voltage [signaling] levels) between the two domains.” EX1025, [0048]; see e.g., EX1037, FIGs. 4-5, 5:45-57, EX1034, 1-2, Figures 2(a), 2(b), 3; EX1002, ¶49. Balancing costs of component and limited die space against the desire to set operating conditions granularly, it was well known to group cores into a “cluster” or “domain” that “share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1007, [0011]; EX1004, [0033]; EX1002, ¶36.

C. Clock Generation

The basic design principles for processor clock signals were also widely known by 2010. EX1002, ¶39. A clock signal may be provided via one or more of: (a) an external reference oscillator (e.g., crystal oscillator), (b) a phase locked loop (PLL), and (c) “clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals.” EX1028, 2:61-3:12; EX1006, 7:23-45; EX1002, ¶¶39-42.



EX1028, FIG. 1 (annotated).

There were known alternatives to create multiple clock signals at different frequencies. EX1002, ¶40; *see e.g.*, EX1004, [0028]; EX1006, 7:23-45; EX1008, 1-2; EX1027, [0005]-[0010]. For example, it was known that “[t]he simplest way to implement multiple clock signals is to provide multiple oscillators.” EX1027, [0005]; EX1002, ¶40. For example, a 2007 publication by Williams illustrates such a configuration with two external crystal oscillators each fed into a separate PLL:




Figure 1: Generic clock domain diagram

EX1008, FIG. 1 (annotated); EX1002, ¶¶41, 62.

As another example, “clock trees or other circuitry” such as “a plurality of dividers” may be “configured to split or divide a clock signal into multiple individual signals” at different frequencies. EX1006, 7:23-45; EX1027, [0027]; EX1002, ¶42.




Fig. 3

EX1027, FIG. 3.

In addition, generating clocks from multiple PLLs, each PLL receiving an independent clock at a frequency determined by a divider was also a known technique. EX1031, FIG. 2; EX1002, ¶¶43-44. For example, in a 2007 publication by Lee, (“Lee II”), the clock signal fed to the Phase Frequency Detector is the target frequency divided by a number from 5 to 283, based on the Divider Divide ratio specified to be “5~283” in the figure below.




Figure 2. Total block diagram of power management unit for dynamic voltage and frequency scaling

EX1031, FIG. 2; EX1002, ¶¶43-44, 64.

D. Processor Cache Memory and Cache Coherency

Processor designers have long used memory storage hierarchies, otherwise known as caches, to compensate for slow memory access speeds which may limit the performance of applications that depend on the rate at which data can be accessed. EX1002, ¶50; *see* EX1026, 1; EX1030, 1. By 2010, it was well known that multiple levels of cache was a desirable configuration for multi-core processors, balancing the trade-off between memory size and speed. EX1002, ¶50; *see e.g.*, EX1006, FIGs. 1, 8. As shown in the figure below, each core may have a local cache that is unique to each core and not directly accessible by other cores. EX1030, 1; EX1002, ¶50; *see also* EX1009, FIG. 1A.




Figure 2. An example of a multiprocessor with private caches.¹

EX1030, FIG. 2. But, to allow the multiple processor cores to execute data processing operations in parallel, it was known for the local caches to “share access to shared memory.” EX1009, [0004]; EX1002, ¶50; *see e.g.*, EX1023 5; EX1006, FIGs. 1, 8. Additionally, shared memory allows multiple “low-cost microprocessors [to be] economically interconnected with shared memory modules,” making this design particularly “cost-effective.” EX1030, 1.


It was “important to ensure that those processors will access the up-to-date data when seeking to access the associated address in shared memory.” EX1009, [0006]; EX1030, 1; EX1002, ¶51. Ensuring all processors operate using up-to-date data, cache coherency, was a known design requirement for processor designs using a shared memory. EX1002, ¶51; *see* EX1009, [0006]-[0009]; EX1030, 1. Cache coherency schemes for multi-core processors have been well known since at least

1990. EX1002, ¶¶52-53; *see* EX1030, 1-2.

IV. THE '339 PATENT

A. Summary

The '339 Patent relates to “power management for a multi-core processor” and “techniques for handling communication between processor cores.” EX1001, 1:58-60, Abstract. Specifically, the '339 Patent teaches that power consumption in a multi-core processor can be reduced using “dynamic supply voltage and clock speed control” by “operat[ing] at high power and high clock frequency when needed and at low power when the computing requirements are reduced.” *Id.*, 1:10-14. Further, dynamic voltage and clock speed control can be applied separately to different groups of processor cores, or “stripes,” (green, red) within the multi-core processor. *Id.*, 2:24-31, FIG. 1.



Id., FIG. 1 (annotated).

The '339 Patent also discloses an “interface block” “to facilitate communications between the first set of processor cores and the second set of processor cores” through the use of “level shifters” to translate between “supply voltages” and “synchronizers” to translate between clock rates of different “clock signals.” *Id.*, FIGs. 2-3, Abstract, 3:27-4:17.




FIG. 2




FIG. 3

EX1001, FIGs. 2-3 (annotated).

B. Prosecution History

Patent Owner filed Application No. 12/713,220, which resulted in the '339 Patent, on February 26, 2010, and was granted an allowance after Patent Owner amended the claims in response to the single Office Action issued by the Examiner.

EX1005, 30-33 (Notice of Allowance), 90-100 (Non-Final Office Action rejecting claims on 101, 112, 102, and 103 grounds).

To overcome the 102 and 103 rejections in the Office Action, Patent Owner argued in an interview with the Examiner that the prior art relied on by the Examiner “disclosed using *a single reference clock*,” which was different from FIG. 3 of the ’339 Patent where “clock signals 1 through 3 were *different/independent clock signals* input to the PLLs.” EX1005, 72 (Applicant-Initiated Interview Summary following the Office Action). In the same interview, “[t]he examiner agreed that Jacobowitz discloses using *a single reference clock*,” but disagreed the current claims excluded such an arrangement. *Id.* “Applicant’s representative stated that he would discuss amends [sic] to the claims.” *Id.* Two days later, Patent Owner filed a response amending the independent claims to add the phrase “the first clock signal *is independent* from the second clock signal.” *Id.*, 76, 78-80. The Examiner allowed the claims following this claim amendment without further comment. *Id.*, 30-33.

C. Level of Ordinary Skill in The Art

Based on the disclosure of the ’339 Patent and the priority date, a person of ordinary skill in the art (“POSITA”) 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. EX1002, ¶32. The prior art evidences this level of ordinary skill. *Id.* Additional education can substitute for some experience, and vice versa. *Id.*

Here, the background technology described in Section III and the prior art described demonstrate that a POSITA would have been familiar with the underlying principles of multi-core processor design, including known techniques for generating and managing distinct clock signals, power management techniques such as dynamic voltage and frequency scaling, intra-chip communication techniques, efficient memory hierarchy designs such as caches and cache coherency protocols, and common semiconductor die layouts. *Id.*, ¶33.

D. Priority Date

For the purposes of this petition, Petitioner assumes a priority date of February 26, 2010. The '339 Patent was filed on this date and does not claim priority to any other document.

V. CLAIM CONSTRUCTION

Petitioner does not believe any terms require a formal construction for the purposes of this Petition.¹ As set forth herein, the challenged claims are unpatentable

¹ The '339 Patent is the subject of co-pending district court action as identified in Section VII. Petitioner reserves the right to address construction of any terms raised

under any reasonable interpretation of the claim terms.

VI. PRECISE REASONS FOR RELIEF REQUESTED

Pursuant to 37 C.F.R. § 42.104(b), Petitioner requests cancellation of the challenged claims based on the following references (EX1002).

Prior Art Reference	Abbreviation (Exhibit No.)
U.S. Patent No. 7,263,457	“White” (EX1006)
US 2009/0271646	“Talwar” (EX1007)
Jeabin Lee et al., “Dynamic Voltage and Frequency Scaling (DVFS) Scheme for Multi-Domains Power Management” (IEEE Asian Solid-State Circuits Conference, 2007)	“Lee” (EX1010)
US 2012/0060170	“Vajda” (EX1011)
US 2007/0096775	“Elgebaly” (EX1004)

The statutory grounds for the challenge of each claim are set forth below. All of the statutory citations are pre-AIA.

at trial, or in the co-pending district court action. Petitioner further reserves the right to raise other invalidity issues, such as invalidity under 35 U.S.C. § 101 or § 112, in the co-pending district court action because such issues cannot be raised in an IPR petition.

Ground	35 U.S.C. §	Claims	Prior Art Reference(s)
1	103	1-3, 5, 8, 10-11, 14, 21	White, Talwar, Lee
2	103	1-3, 5, 8, 10-11, 14, 21	Elgebaly, Lee
3	103	9	Elgebaly, Lee, Vajda

Below, Petitioner discusses why the challenged claims are unpatentable under the statutory grounds raised, including how and where the prior art satisfies each element of each challenged claim, as required by 37 C.F.R. § 42.104(b)(4).² Petitioner’s showing establishes that there is a reasonable likelihood that it will prevail on each ground of unpatentability to each challenged claim.

A. Ground 1: White, Talwar, and Lee Render Obvious Claims 1-3, 5, 8, 10-11, 14, 21

1. White

White published on August 28, 2007, and is thus prior art under at least pre-AIA §102(b). EX1006; EX1002, ¶65. White discloses a multi-core processor in which each core “may be configured to operate at frequencies and voltages independently of the other logic cores.” EX1006, 2:15-44; EX1002, ¶65. A POSITA

² Petitioner also shows how each claim preamble is satisfied, but does not concede that any preamble is limiting for any ground in the Petition.

would understand that White’s adjustment of frequency and voltage is the same as the well-known technique known as dynamic frequency and voltage scaling (DVFS) as discussed above. EX1002, ¶66, *supra* Section III.A.

Figure 2A illustrates the disclosed structure for varying frequency and voltage of the cores:




FIG. 2A

EX1006, FIG. 2A (annotated); EX1002, ¶67.

“[Orange] power management logic 200” “control[s] the operating characteristics, such as the operating voltage and/or operating frequency, of each of multiple logic cores 120.” EX1006, 8:59-65; EX1002, ¶68. “As with power management logic 200,” “temperature control logic 210” may also be configured to “adjust the voltages and frequencies supplied to logic cores 120A and 120B.” *Id.*; EX1006, 13:27-43. “[Blue] voltage regulator 130 . . . may be configured to supply

independent voltages to each logic core,” each of which can be “different from the voltage levels at which other logic cores are operated.” *Id.*, 7:23-45, 10:31-32, FIG. 2A; EX1002, ¶68. Similarly, “[purple] frequency regulator 140” may “supply multiple, independent clock signals” to the [green, red] logic cores 120A and 120B. EX1006, 6:41-49, 7:23-45; EX1002, ¶68. “[Orange] power management logic 200 may be configured to communicate ... with [blue] voltage regulator 130 to instruct voltage regulator 130 to supply a particular voltage to [green] logic core 120A” and a particular voltage to [red] logic core 120B. *Id.*, EX1006, 10:11-24. “[purple] frequency regulator 140” “instruct[s] frequency regulator 140 to supply a particular frequency or clock signal to [green] logic core 120A” and a particular clock signal to [red] logic core 102B. *Id.*; EX1002, ¶68.

White discloses multiple options for generating and changing clock signals using a combination of elements at frequency regulator 140 and the cores. *Id.*, ¶69. Frequency regulator 140 includes “various sorts of clock sources, such as an oscillator or phase locked loop (PLL) to generate a reference clock” and uses those “multiple clock sources in order to supply multiple, independent clock signals of different frequencies.” *Id.*, EX1006, 7:38-45. Frequency regulator 140 may also include “clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals.” *Id.*, 7:40-43; EX1002, ¶69. Additionally, each “logic core may generate its internal clock using a phase-locked loop (PLL) locked to an

externally provided clock signal” from “frequency regulator 140.” *Id.*, EX1006, 9:62-65.

White also discloses the use of “common bridge logic 110, such as a Northbridge” “configured to communicate with the logic cores as well as with other elements or peripherals . . . such as . . . [a] memory subsystem 170.” *Id.*, 7:1-7, 17:29-32, FIGS. 2A, 2B; EX1002, ¶70. “Common bridge logic 110” can “interface with a logic core using the same logic levels and/or clock timing as the logic core.” *Id.*, EX1006, 7:64-8:10. Common bridge logic further includes an interface “to translate the communication, such as the signals on one or more control lines . . . between the different frequencies.” *Id.*, 8:21-26, FIG. 3; EX1002, ¶70. A POSITA would understand that the common bridge logic disclosed by White includes known designs for voltage-level shifters and clock synchronization blocks to translate signals between the voltage levels and clock domains of each logic core. *See supra* Section III.B, EX1010, FIG. 3; EX1002, ¶70.

Finally, while Figure 2A illustrates two logic cores, White explains that its teachings “may be applied to systems or devices that include ***any number of logic cores*** configured to operate at independent voltages and/or frequencies.” EX1006, 6:56-60, 6:63-67, claims 1, 25; EX1002, ¶71.

White is analogous art to the claimed invention because it is directed to the same field of endeavor—power management of computer processors. (EX1001, 1:6-

14; EX1006, 2:15-44.) By disclosing a multi-core processor in which “[t]he operating voltage and frequency of the logic cores ... may be adjusted,” including for power management, White is reasonably pertinent to a problem faced by the ’339 patent inventors—dynamic control of core voltage and clock speed to optimize “power consumption management.” (EX1001, 1:6-14; EX1006, 2:15-44.)

2. *Talwar*

Talwar was filed October 31, 2008, and published October 29, 2009, and is thus prior art under at least pre-AIA §§102(a), (e). EX1007; EX1002, ¶72. Talwar discloses a multi-core processor design in which “[t]he cores and voltage sources are divided into clusters, whereby multiple cores in a cluster receive power from a single voltage source” and “[e]ach cluster may have a different cluster frequency,” but “all the cores in a cluster [operate] at the same frequency.” *Id.*, EX1007, [0011], [0018]-[0019]. The values of: (a) the “single voltage source” for the cluster, and (b) “the frequency of all cores in a cluster” can be adjusted “[d]uring power management.” EX1007, [0011]-[0012], [0018-0019], [0021-0025], [0030] (“The cores in a cluster ... can be independently controlled from other clusters.”); EX1002, ¶72. Providing one voltage and clock source per cluster enables “maximum power savings” by efficiently performing DVFS on each cluster in “stark contrast to conventional multi-core chips” with only a single voltage source. EX1007, [0012], [0021]; EX1002, ¶74.

Talwar is analogous art to the claimed invention because it is directed to the same field of endeavor—power management of computer processors. (EX1001, 1:6-14; EX1007, Abstract, [0011]-[0012].) By disclosing a multi-core processor system in which “power management is performed in [a] multi-core system based on one or more of core utilization and a management policy,” Talwar is reasonably pertinent to a problem faced by the ’339 patent inventors—dynamic control of core voltage and clock speed to optimize “power consumption management.” (EX1001, 1:6-14; EX1007, Abstract.)

3. *Lee*

“Dynamic Voltage and Frequency Scaling (DVFS) Scheme for Multi-Domains Power Management” by J. Lee et al. (“Lee”) is a paper prepared for presentation and distributed at the 2007 IEEE Asian Solid-State Circuits Conference held in Jeju, Korea, on November 12-14, 2007. EX1010; EX1002, ¶77. Lee was published by IEEE and added to IEEE Xplore on January 7, 2008. EX1015, 1; EX1002, ¶63. Lee was publicly accessible and would have been searchable and obtainable by the public by at least January 7, 2008, and is thus prior art under at least 35 U.S.C §§ 102(a) and (b). EX1015, 1; EX1002, ¶63.

Lee discloses a “DVFS scheme for Multi-power domains.” EX1002, ¶77. As illustrated in Figure 2, Lee shows three domains with independent power and in which “the voltage and frequency of each power domain are dynamically scaled

according to the performance requirement of each domain.” EX1010, 1; EX1002, ¶77.




Fig. 2 Triple Domains Power Management

EX1010, FIG. 2; EX1002, ¶77.

Lee discloses that “**each domain has different clock and supply voltage**” and further discloses techniques for communicating between the domains: “[t]he level shifter (LS) and synchronizer (Sync.) ... change the voltage level and convert the clock frequency of transferred data, respectively.” EX1010, 2 (emphasis added); EX1002, ¶78.




Fig. 3 (a) Level shifter (b) Synchronizer

EX1010, FIG. 3.

For another example, as shown below in Figure 8, Lee discloses a chip layout for domains arranged in rows (stripes):




Fig. 8 Chip micro-photograph

Id., FIG. 8; EX1002, ¶82.

As shown in Figures 8 and 2 above, Lee discloses that the domains are independent. *Id.*, ¶83. Each of the PLLs (in each of the PMUs) receives both a reference clock and a divider output signal “Div_Clk.” derived from a division of the clock for that domain. *Id.* As shown below in Lee’s Figure 4 (modified), the PLL compares the Div_Clk with the Reference clock and adjusts the voltage-controlled oscillator until Div_Clk’s frequency and phase aligns with the Reference Clock. *Id.* The division of clock is based on a divisor value set by the PMU based on the “Performance Index”. *Id.*




Fig. 4 The DVFS scheme for multi-domains power management

EX1010, FIG. 4 (modified with a signal label); EX1002, ¶83.

Further, as each of the domain has its own separate and independent reference clock input, Lee also discloses separate independent reference clocks to each of the domains. *Id.*, ¶84. Further, Figure 8 illustrates that the PMUs are located at different

corners of a die, and thus, a POSITA would understand that the reference clocks may be supplied to each of the domain from independent off-chip sources. *Id.*

Lee further discloses a system where the performance of the total system is determined by a RISC, which is known to a POSITA as a type of microprocessor. EX1010, 2; EX1002, ¶79. The performance of the system is communicated to the PMUs (power management unit) in one domain, and the other domains automatically adapt to produce/consume data at a matching rate. *Id.* Lee discloses a PMU for each domain that generates that domain’s “clock and supply voltage to adjust the speed of target domain.” EX1010, 2; EX1002, ¶80.

Additionally, Lee discloses that the clocks in the PMU are generated by the phase locked loop (PLL) structure. *Id.*, ¶81. Lee discloses that the DVFS scheme “uses the **phase locked loop (PLL)** structure to find supply voltage level and clock frequency level on the MEP.” EX1010, 3 (emphasis added), FIG. 4; EX1002, ¶81.

Lee is analogous art to the claimed invention because it is directed to the same field of endeavor—power management of computer processors. (EX1001, 1:6-14; EX1010, Abstract, I.) By disclosing processors having multiple different power domains, for which clock frequency and supply voltage can be separately adjusted for “power reduction,” Lee is reasonably pertinent to the problem faced by the ’339 patent inventors—dynamic control of core voltage and clock speed to optimize “power consumption management.” (EX1001, 1:6-14; EX1010, Abstract, I.)

4. *Motivation to Combine*

A POSITA would be motivated to apply Talwar’s DVFS “clustering” of processor cores to White’s DVFS implementation in order to achieve a goal common to White and Talwar—DVFS in a multi-core processor internal to a chip—while meeting practical chip space constraints. EX1002, ¶85. Specifically, while White illustrates two logic cores, it explains that a device may “include *any number of logic cores* configured to operate at independent voltages and/or frequencies.” EX1006, 6:56-60; EX1002, ¶85. Following this disclosure, a device may include four—or eight or sixteen—cores, each with separate voltage and clock sources. EX1006, 6:56-60; EX1002, ¶89. However, a POSITA would have understood such a 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.*, ¶¶88-89 (citing EX1029, 3-4, EX1019, 52-53). Talwar discloses achieving the same goal as White of performing DVFS internal to the chip by grouping cores into “clusters” and achieving “maximum power savings”—in “stark contrast to conventional multi-core chips” with only a single voltage source (EX1007, [0021], [0012])—but simultaneously reducing the costs of the “more expensive” “per-core DVFS”

(EX1019, 52-53) by having “less voltage sources than cores on the chip” (EX1007, [0014]). EX1002, ¶87. Using Talwar’s clustering with White’s DVFS would therefore advantageously conserve chip space in multi-core processors. A POSITA would have been further motivated to apply Talwar’s DVFS “clustering” of processor cores to White’s DVFS implementation because White explicitly states that its teachings are not limited to per-core voltage and clock signals, but may more generally “be applied to any integrated circuit or component that includes multiple areas or sections of circuitry configured to operate at independent operating voltages and/or frequencies.” EX1006, 8:41-49; *id.*, 3:37-52, 6:63-67; EX1002, ¶90.

A POSITA would be further motivated to combine White/Talwar with Lee to benefit from adding Lee’s further implementation details on clocking, synchronization, level shifting, and chip layout to the known concept of DVFS that White and Talwar already explicitly disclosed. *Id.*, ¶86. Further, Talwar explicitly teaches clusters, each set to its own voltage and frequency, which corresponds with Lee’s disclosure of DVFS domains. *Id.*, ¶91. 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. EX1002, ¶86, 91-92. These specific implementation details would have added beneficial functionality to the White/Talwar combination.

Further, the POSITA would have had a reasonable expectation of success in making the combination, at least because of the specific implementation detail (e.g., specific chip layouts) provided by White, Talwar, and Lee, in view of the common components (e.g., clock and voltage supply circuitry) and advanced state of the art as of the priority date.

5. *Modification of White in View of Talwar and Lee*

As discussed above, White explicitly discloses applying its disclosures to a device that “include[s] *any number of logic cores* configured to operate at independent voltages and/or frequencies.” EX1006, 6:56-60; EX1002, ¶93. Following this disclosure, White’s FIG. 2 embodiment, including logic cores 120A-120B can be modified to include four logic cores 120A-120D. EX1002, ¶77. White’s common bridge logic would have a connection to each additional logic core. EX1002, ¶95. Then, applying Talwar’s clustering disclosure, [green] logic cores

120A and 120C constitute a first cluster and [red] logic cores 102B and 120D constitute a second cluster as shown below:




FIG. 2A

Combination of EX1006, FIG. 2A and EX1007, FIG.1 (annotated); EX1002, ¶93

Further, Lee supplied White and Talwar with the detail logic design for implementing clock generation and synchronizers required for a multi-cluster (multi-domain) DVFS system. EX1002, ¶98. For example, Lee's PLL design supplied additional details on how to generate a clock for each cluster. EX1002, ¶93. In accordance with Talwar's disclosure providing a "single voltage source" for the cluster and the same frequency for "all cores in a cluster" (EX1007, [0011], [0019]), which a POSITA would understand refers to a single clock source for the cluster,

each cluster (rather than each core) would receive a voltage signal from voltage regulator 130 and a single clock signal from frequency regulator 140. EX1002, ¶¶93-94. Further, 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. *Id.*, ¶94. Frequencies and voltages may be assigned to the clusters/domains in the White/Talwar/Lee combination through either Lee’s disclosure of having a microprocessor setting the voltage and frequency of one domain and the other domains adapt based on data rates, Talwar’s power management policies, or White’s temperature limits. The ’339 Patent does not specify the method of assignment and a POSITA could determine the method to use based on system requirements. *Id.*, ¶96.

With those modifications, the system would perform the same power management, or DVFS methods as White, but would do so on a cluster-by-cluster basis. *Id.*, ¶¶93, 97.

6. *Element-by-Element Analysis*

a. Claim 1

[1.pre] A multi-core processor, comprising:

To the extent the preamble may be limiting, White/Talwar/Lee renders obvious Limitation [1.pre]. *Id.*, ¶100. As discussed in Section VI.A.1, White discloses methods for independently adjusting supply voltages and clock frequencies in “a multi-core integrated circuit,” or “[a] processor or other integrated circuit [that]

may include multiple logic cores.” EX1006, 2:15-22, 4:31-35, 6:41-46, 8:41-49, 14:39-54, 15:42-64; EX1002, ¶100.

[1.1] a first set of processor cores of the multi-core processor,

White/Talwar/Lee renders obvious Limitation [1.1]. EX1002, ¶101. As discussed in Section VI.A.5 and shown in the White/Talwar version of Figure 2A, White/Talwar implements Talwar’s teaching that “[t]he cores and voltage sources are divided into clusters.” EX1007, [0011], claim 1; EX1006, claims 1, 15; EX1002, ¶101. Cores 120A and 120C are a first cluster (*first set of processor cores*) and cores 120B and 120D are a second cluster (*second set of processor cores*) of integrated circuit 100 (the *multi-core processor*). EX1002, ¶102.




FIG. 2A

EX1006, FIG. 2A (modified to show Talwar's multiple cores per cluster); EX1002, 101.

[1.1.1] 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;

White/Talwar/Lee renders obvious Limitation [1.1.1]. EX1002, ¶103. White teaches that independent supply voltages (*first supply voltage*) and clock signals (*first output clock signal*) may be provided to each processor core in a multi-core processor. *Id.* For example, White discloses “power management logic 200 may be configured to communicate ... with voltage regulator 130 to instruct voltage regulator 130 to supply a particular voltage to logic core 120A,” and “similarly, power management logic 200 may be configured to communicate with frequency regulator 140 to instruct frequency regulator 140 to supply a particular frequency or clock signal to logic core 120A.” EX1006, 10:11-24; EX1002, ¶103. Such supply voltages and clock signal frequencies may be “adjusted” as needed, for example to support “power management” or “temperature management” requirements. EX1006, 2:15-44; EX1002, ¶104.





FIG. 2A

EX1006, FIG. 2A (modified to show Talwar’s multiple cores per cluster and labels added corresponding to those in Lee); EX1002, ¶106.

Lee discloses separate PMUs (power management units) for three different domains, where each PMU has a PLL (*first phase lock loop (PLL)*) that receives a (*first clock signal*). EX1010, FIG. 4; EX1002, ¶105.

Lee discloses, with reference to Figure 4, that the PMU has multiple unique instances corresponding to the number of different frequency/voltage domains. EX1002, ¶106. The purple shaded areas of Lee’s Figure 4 diagram include a frequency generator, and the blue shaded areas included a voltage generator. *Id.* Lee’s divider outputs “Div_Clk” which sets the frequency and is the *clock signal*

input to the PLL. *Id.* Each instance of the PLL receives an independent Div_Clk (clock signal input) and produces an independent output clock signal for that domain. *Id.*



first/second clock signal as input
Div_Clk_AC or Div_Clk_BD

EX1010, FIG. 4 (modified with labels added); EX1002, ¶106.

A POSITA would understand that those domains as disclosed by Lee could include *first set of processor cores* such as White’s “any number of processor cores” or Talwar’s clusters of cores. EX1002, ¶105.

[1.2] a second set of processor cores of the multi-core processor,

White/Talwar/Lee renders obvious Limitation [1.2] for the reasons discussed for Limitation [1.1]. EX1002, ¶107.

[1.2.1] 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,

White/Talwar/Lee renders obvious Limitation [1.2.1] for the reasons discussed in for Limitation [1.1.1]. EX1002, ¶108. As discussed for Limitation [1.1.1], White teaches that independent supply voltages (*second supply voltage*) and frequencies (*second output clock signal*) may be provided to each processor core in a multi-core processor. *Id.* In addition, White also discloses such supply voltages and clock signal frequencies may be “adjusted” as needed, for example to support “power management” or “temperature management” requirements. EX1006, 2:15-44; EX1002, ¶104.

[1.2.2] wherein the first supply voltage is independent from the second supply voltage,

White/Talwar/Lee renders obvious Limitation [1.2.2]. EX1002, ¶109. As explained in Sections VI.A.1 and VI.A.2, both White and Talwar teach that the supply voltages to individual cores (White) or clusters of cores (Talwar) can be independently controlled. EX1006, 10:11-11:23; EX1007, [0011]-[0012]; EX1002, ¶109. Further, Lee discloses a system with three domains in which “each domain has different clock and supply voltage” and that “the voltage and frequency of each power domain are dynamically scaled according to the performance requirement of each domain.” EX1010, 1; EX1002, ¶109.

Therefore, in the White/Talwar/Lee combination, the first supply voltage to the first cluster/domain is independent from the second supply voltage to the second cluster/domain. EX1002, ¶109.

[1.2.3] and the first clock signal is independent from the second clock signal; and

White/Talwar/Lee renders obvious Limitation [1.2.3]. EX1002, ¶110. A POSITA would have understood that Lee discloses multiple, independent clock signals as inputs to the PLLs at each core/domain. *Id.* First, Lee discloses that the Div_Clk input to each PMU is a *clock signal* independent of other Div_Clk signals. *Id.* A POSITA would understand that the signals need to be independent because the domains may run at different frequencies, and the way those frequencies are selected is via the divider in each PMU that generates a Div_Clk signal at the desired frequency. *Id.* Second, Lee discloses that each PMU also receives a Reference Clock. EX1002, ¶111. The location of the PMUs in the Lee chip diagram (EX1010, FIG. 8, at the corners of the die) suggests to a POSITA that the Reference Clocks come from an off-chip source. EX1002, ¶111. The reference clocks can also be considered *independent* input clock signals because they may be driven by different oscillators or clock dividers. *See* Section III.C. EX1002, ¶¶110-111.

For example, as shown in Figure 2, Lee discloses three power management units (highlighted in yellow, blue and pink), which are independent. EX1002, ¶112. As described in Section VI.A.3, Lee discloses that each of the three PMUs has one of the PLLs shown in Figure 4. *Id.* Hence, there are three independent “Reference Clocks”, one for each PMU. *Id.* The “Div_Clk” in the first (yellow) PMU is the *first clock signal which is independent from the second clock signal* in the second (blue)

PMU. *Id.* Alternatively, each PMU has a Reference Clock in the first (yellow) PMU which is a *first clock signal* which is *independent from the second clock signal* in the second (blue) PMU. EX1002, ¶113. A POSITA would understand these clocks to be independent, whether from separate oscillators, or from independent outputs from dividers of a single oscillator. EX1002, ¶¶112-113.




Fig. 2 Triple Domains Power Management

EX1010, FIG. 2 (emphasis added); EX1002, ¶111.




Fig. 4 The DVFS scheme for multi-domains power management

EX1010, FIG. 4 (emphasis and labels added); EX1002, ¶111.

Further, White explicitly teaches that the frequency regulator “may include multiple clock sources in order to supply multiple, independent clock signals of different frequencies.” EX1006, 7:38-45. A POSITA would have understood, from such explicit teachings and the disclosures of independent reference clock as inputs to PLLs at each domain in Lee, that White/Talwar/Lee discloses the first clock signal is independent from the second clock signal. EX1002, ¶114.

[1.3] an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,

White/Talwar/Lee renders obvious Limitation [1.3]. EX1002, ¶115. White teaches a common bridge logic that “may be configured to translate the communication, such as the signals on one or more control lines, between the

different logic levels and/or between the different frequencies.” EX1006, 8:13-40, 14:55-15:17; EX1002, ¶115. As shown in Figure 3 below, White discloses an *interface block*, translation logic 300A, directly coupled to core 120A (*the first set of processor cores based on Talwar/White Combination*) and indirect coupling to logic core 120B (*the second set of processor cores based on Talwar/White Combination*) through communication protocol interfaces 320A and 320B and translation logic 300B in common bridge logic 110. EX1002, ¶¶115-116. The indirect coupling to 120B can also include the path to and from main memory 1104. EX1002, ¶115. White discloses that both translation logic 300A and 300B perform the same voltage and frequency translations as described in the ’339 Patent specification on interface block. EX1002, ¶116.




FIG. 3

EX1006, FIG. 3; EX1002, ¶115.

Further, Lee discloses a Domain Converter for communicating between domains via the FIFO. EX1002, ¶117. In the White/Talwar/Lee combination, Lee's Domain Converter (*interface block*) couples Domain1 and Domain2 which contain the White/Talwar clusters of cores (*first set of processor cores and second set of processor cores*). *Id.*




Fig. 2 Triple Domains Power Management

EX1010, FIG. 2; EX1002, ¶117.

[1.3.1] wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.

White/Talwar/Lee renders obvious Limitation [1.3.1]. EX1002, ¶118. As discussed in Section VI.A.1, White’s common bridge logic block is “an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores.” *Id.* As shown in modified Figure 2A and explained in White, the common bridge logic is “configured to communicate with the logic cores as well as with other elements or peripherals . . . such as . . . [a] memory subsystem.” EX1006, 7:1-7; EX1002, ¶118. Such communication is passed through the “Bus Bridge” that “provide[s] an interface between microprocessor 1100, main memory 1104, graphics controller 1108, and devices attached to PCI bus 1114” and “translates an operation from the protocol used by the source device or bus to the protocol used by the target device or bus”. EX1006, 25:21-31; EX1002, ¶118.

Based on White's explicit reference to a "memory subsystem" and the knowledge of a POSITA in 2010 that cores may read and write to a common memory space through shared memory and cache coherency (§ III.D), a POSITA would understand that each of cores of the first and second sets of processor cores may read and write to common space in the memory subsystem, assisted by cache coherency, via the common bridge logic. EX1002, ¶119. Accordingly, the translation logic 300, via common bridge logic provides a communication path between the cores via the memory subsystem. *Id.*

b. Claim 2

[2.1] The multi-core processor of claim 1, wherein the interface block further comprises a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set of processor cores for a first signal traveling from the first set of processor cores to the second set of processor cores.

White/Talwar/Lee renders obvious Limitation [2.1] for reasons discussed above in Limitation [1.3.1]. EX1002, ¶120. As discussed in Section III.B, it was well known in the art that multi-core processors implementing DVFS required mechanisms for translating logic levels between chip components operating at different voltages. *Id.* Specifically, White discloses a "voltage translation circuitry configured to adjust the voltage level used for one or more communication signals or lines between the common bridge and a logic core of the processor", i.e., a level translation (*a first level shifter*) to allow the cores to communicate with each other

through the common bridge logic which may operate at different voltage and frequency. EX1006, 4:36-43; EX1002, ¶122. Similarly, Lee discloses a mechanism called “LS: Level Shifter” in Figure 3a. EX1002, ¶123.




Fig. 3 (a) Level shifter (b) Synchronizer

EX1010, FIG. 3a (annotated); EX1002, ¶123.

A POSITA would have understood White’s teaching that the common bridge logic “may be configured to translate the communication ... between the different logic levels” to mean that the common bridge logic includes one or more level shifters that enable the bridge to translate between the logic levels used by the first set of processor cores and the second set of processor cores. EX1002, ¶124; *see* EX1006, 8:13-40, 14:55-15:17, FIG. 3. Further, there were well-known prior art techniques for doing so. *See e.g.*, EX1037, FIGs. 4-5, 5:45-57; EX1034, FIGs. 2(a), 2(b), 3; EX1002, ¶124. Additionally, a POSITA would have understood that the level shifter that translates between the logic level of the first set of processor cores and the logic level of the second set of processor cores must be referenced to the

second supply voltage to know at what logic level the second set of processor cores is operating. *Id.*

c. Claim 3

[3.1] The multi-core processor claim 1, wherein the interface block further comprises a second level shifter that is referenced to the first supply voltage and adapted to translate second logic levels associated with the second set of processor cores to first logic levels associated with the first set of processor cores for a second signal traveling from the second set of processor cores to the first set of processor cores.

White/Talwar/Lee renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1]. EX1002, ¶125. A POSITA would understand that processor communications is bidirectional. A first and second level shifter are needed as level shifters are required when cores at different voltage levels communicate at both directions. *Id.*

d. Claim 5

[5.1] The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.

White/Talwar/Lee renders obvious Limitation [5.1]. EX1002, ¶129. As best understood for the purposes of this petition, “periphery” refers to a physical location of the power control block 108 and clock control block 110 in Figure 1 of the ’339 Patent (arranged at two different sides of the multi-core processor 100). EX1002, ¶128. Lee meets the limitation at least by its disclosure of PMU1, PMU2, and PMU3 located on the outer edges (e.g., a periphery) of the chip. EX1010, FIGs. 2, 8;

EX1002, ¶128.





Fig. 8 Chip micro-photograph

EX1010, FIG. 8 (annotated); EX1002, ¶129.

To the extent “periphery” is understood to refer to “peripherals,” e.g., controllers and devices connected to a processor through an I/O (input/output) bus, the limitation is met at least by White’s disclosure of controllers located in a periphery of the multi-core processor; for example, White’s Figure 8 shows PCI devices (the most common IO bus in the 2010 timeframe) 1112A and 1112B that are connected to the bus bridges. EX1002, ¶127. A POSITA would understand that such controllers send one or more control signals that are received by the processor cores.

Id.

A POSITA would understand that control signals, as recited in the limitation, would serve to adjust the frequency and voltage of the cores. EX1002, ¶129. As annotated below, Lee's Figure 2 shows *VDD and CLK signals* communicated from each *PMU* to its corresponding domain (*sets of processor cores*). *Id.* The VDD and CLK signals are control signals, or alternatively, Lee discloses control signals from the PMUs in the periphery to the cores because clock and voltages would typically be accompanied by other control signals (such as to indicate the signals are valid). *Id.* In the White/Talwar/Lee combination, each domain contains multiple cores. Hence, White/Talwar/Lee discloses sets of *processor cores [that] are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor. Id.*



Combination of EX1010, FIG. 2 and EX1007, FIG. 1 (annotated); EX1002, ¶129.

e. Claim 8

[8.1] The multi-core processor of claim 1, wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor.

White/Talwar/Lee renders obvious Limitation [8.1]. EX1002, ¶130. As illustrated in Figure 8, Lee explicitly discloses that Domain 1, Domain2, and Domain3 are located in separate *regions* of the chip. Lee in combination with White and Talwar show that the *first, second and third regions* could include a *set of processor cores*. *Id.* Further, as discussed in Section VI.A.5, White/Talwar includes a first cluster of logic cores 120A and 120C (the first set of logic cores) that share a voltage signal, clock signal, and cluster-level clock control circuitry and a second cluster of logic cores 120B and 120D (the second set of logic cores) that share a different voltage signal, clock signal, and cluster-level clock control circuitry. *Id.*

f. Claim 10

[10.1] The multi-core processor of claim 8, wherein the first region and the second region are non-overlapping regions of the multi-core processor.

White/Talwar/Lee renders obvious Limitation [10.1] for reasons discussed in Limitation [8.1]. EX1002, ¶131. The die photograph in Figure 8 and block diagram in Figure 2 of Lee clearly show that the three domains (*regions*) are non-overlapping. *Id.* The combined Lee/Talwar drawing shows the *multi-core processor* distributed to the Lee domains, and the die photograph in Lee's Figure 8 shows domains located in *non-overlapping regions of the multi-core processor*. *Id.*

g. Claim 11

[11.1] The multi-core processor of claim 8, wherein the first region corresponds to a first row of the multi-core processor, and wherein the second region corresponds to a second row of the multi-core processor.

White/Talwar/Lee renders obvious Limitation [11.1] for reasons discussed in Limitation [8.1]. EX1002, ¶132. Further, the die photograph in Figure 8 of Lee explicitly shows that the three domains (*regions*) are in separate rows. *Id.*

h. Claim 14

[14.1] The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores

White/Talwar/Lee renders obvious Limitation [14.1]. Limitation [5.1] includes all elements of Limitation [14.1] except that in Limitation [14.1], the control blocks are *located in a common region that is substantially central*. EX1002, ¶133. A POSITA would have understood that circuitry design involves layout decisions that balance performance, area, and thermal considerations, and that placement of control logic circuitry along the perimeter of the multi-core processor or substantially central to processor cores is an obvious design choice. *Id. See also* Kato EX1033, FIG. 8, [0056], [0103]-[0107] (depicting chip “floor plan” with “a controller for signal transmission from/to storage control units SC[]” that are “located at the side nearer to the SCs [storage control units],” which is substantially central to the processor units.); EX1002, ¶133. Further, a POSITA would understand the benefits of placing control blocks around the periphery of core clusters to shorten

the length of wires between the cores and the control logic to shorten delay. *Id.* In addition, this way, core clusters can be placed close to each other and have a shorter inter-cluster communication wires. *Id.*

i. Claim 21

[21.pre] A multi-core processor, comprising:

White/Talwar/Lee renders obvious Limitation [21.pre] for reasons discussed in Limitation [1.pre]. EX1002, ¶134.

[21.1] a first set of processor cores of the multi-core processor,

White/Talwar/Lee renders obvious Limitation [21.1] for reasons discussed in Limitation [1.1]. EX1002, ¶135.

[21.1.1] wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block;

White/Talwar/Lee renders obvious Limitation [21.1.1] for reasons discussed in Limitation [1.1.1]. EX1002, ¶136. This claim is the same as Limitation [1.1.1] except that here, the voltage is received from a power control block, and the clock signal is received as input in a clock control block. *Id.* As explained in Section VI.A.5 and shown in modified Figure 2A, in White/Talwar/Lee each processor in the first / second cluster receives a first / second voltage signal (the first / second supply voltage) from voltage regulator 130 (the power control block). *Id.* A POSITA would have understood White's "voltage regulator 130" and Lee's PMU to be power

control blocks, and White's "frequency regulator 140" and Lee's PMU to be clock control blocks. *Id.*

[21.2] a second set of processor cores of the multi-core processor,

White/Talwar/Lee renders obvious Limitation [21.2] for reasons discussed in Limitation [1.2]. EX1002, ¶137.

[21.2.1] wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage from the power control block and a second output clock signal from a second PLL having a second clock signal as input in the clock control block,

White/Talwar/Lee renders obvious Limitation [21.2.1] for reasons discussed in Limitations [1.2.1] and [21.1.1]. EX1002, ¶138. As shown in the composite diagram in Limitation [1.1.1], the power control block is White's frequency regular and voltage regulator as implemented by Lee's PMUs. *Id.*

[21.2.2] wherein the first supply voltage is independent from the second supply voltage, and

White/Talwar/Lee renders obvious Limitation [21.2.2] for reasons discussed in Limitation [1.2.2]. EX1002, ¶139.

[21.2.3] the first clock signal is independent from the second clock signal; and

White/Talwar/Lee renders obvious Limitation [21.2.3] for reasons discussed in Limitation [1.2.3]. EX1002, ¶140.

[21.3] an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,

White/Talwar/Lee renders obvious Limitation [21.3] for reasons discussed in

Limitation [1.3]. EX1002, ¶141.


[21.3.1] wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.

White/Talwar/Lee renders obvious Limitation [21.3.1] for reasons discussed in Limitation [1.3.1]. EX1002, ¶142.

B. Ground 2: Elgebaly and Lee Render Obvious Claims 1-3, 5, 8, 10-11, 14, 21

1. Elgebaly

Elgebaly was published on May 3, 2007, and is thus prior art under at least pre-AIA §102(b). EX1004; EX1002, ¶143. Elgebaly discloses “[t]echniques for adaptively scaling supply voltage” and generating “[a] target clock frequency for the processing core” “based on computational requirements for the core.” EX1004, Abstract, [0009], FIGs. 1-2; EX1002, ¶143. While Elgebaly’s Figures 1 and 2 each show an ASIC (120) with two processor cores, “ASIC 120 may include ... **any number of processor cores.**” EX1004, [0030]; EX1002, ¶145 (emphasis added). Consistent with Elgebaly’s teaching and as described in Section III.B, multi-core processors containing four cores were also well known in the art and available commercially well before 2010. EX1002, ¶143; *see e.g.*, EX1023, 2, EX1004, [0030], [0033]. Elgebaly therefore teaches an ASIC 120 with four cores, as in Modified Figure 1 below, with additional [blue] Processor Core 0 and an additional [red] Processor Core 3. EX1002, ¶144.



EX1004, FIG. 1 (modified to show additional processor cores); EX1002, ¶144.

a. PLLs to generate clock signals:

Elgebaly also discloses that “Clock generator 140 generates clocks used by the processing units within ASIC 120 and may couple to external circuitry 142, which may include crystals, inductors, capacitors, and so on. Clock generator 140 may include one or more phase locked loops (PLLs) that control the oscillators used to generate the clocks.” EX1004, [0028] (emphasis added); EX1002, ¶146.

Thus, Elgebaly describes the use of PLLs to generate clocks. *Id.*

b. Grouping of multiple cores sharing common clocks and voltages

Additionally, Elgebaly teaches that “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” and

that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1004, [0033]; EX1002, ¶147. Thus, Elgebaly teaches grouping of multiple processing cores to share a common clock and for a common Supply Voltage to be jointly set for these multiple cores. *Id.*

Elgebaly is analogous art to the claimed invention because it is directed to the same field of endeavor—power management of computer processors. (EX1001, 1:6-14; EX1004, Abstract, [0005]-[0010].) By disclosing techniques for scaling processor voltage for a target clock frequency to “conserve battery power” in electronic devices, Lee is reasonably pertinent to the problem faced by the ’339 patent inventors—dynamic control of core voltage and clock speed to optimize “power consumption management.” (EX1001, 1:6-14; EX1010, Abstract, [0005]-[0010], [0025].)

2. *Motivation to Combine Elgebaly with Lee*

Elgebaly teaches “highly efficient adaptive voltage scaling” and that such “[a]daptive voltage scaling thus exploits the variations in computational requirements to scale the supply voltage and reduce the average power consumption while maintaining the required performance.” EX1004, [0034], [0035]; EX1002, ¶148. A POSITA would have been motivated to apply Lee’s DVFS clocking, domain converters, and chip layout to the cores in the processing units of Elgebaly’s ASIC

in order to achieve the common goal of “highly efficient adaptive voltage scaling,” via Lee’s “dynamical scaling the clock frequency and the supply voltage [to achieve] extremely efficient energy saving.” EX1004, [0035]; EX1010, 1; EX1002, ¶148.

Elgebaly presents an embodiment in which “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” (EX1004, [0033]), but has few details of how those independent clocks are produced. EX1002, ¶149. A POSITA implementing this embodiment of Elgebaly would naturally turn to a reference such as Lee to fill in the details on how the independent clocks may be generated using PLLs. *Id.*; EX1010, FIG. 4.

Further, Elgebaly shows in Figure 1, that multiple processor cores communicating to memory over a common bus, yet, it does not show the circuitry to allow them to communicate when the core frequency or voltage does not match that of the target memory or core. EX1004, FIG. 1; EX1002, ¶150. Lee discloses in Figures 2, 3a, and 3b the use of Domain Converter to perform level translation and clock synchronization. EX1010, FIGs. 2, 3a, and 3b; EX1002, ¶150. Specifically, Lee discloses an *interface block* comprising a Domain Converter and FIFO for communicating between domains. EX1002, ¶151. Thus, a POSITA would turn to a reference such as Lee to show the circuits to perform the level translation and clock synchronization. EX1002, ¶150.

In the Elgebaly/Lee combination, Lee's Domain Converter and FIFO (interface block) couples Domain1 and Domain2 which contain the Elgebaly cores first set of processor cores and second set of processor cores. EX1002, ¶151.




Fig. 2 Triple Domains Power Management

EX1010, FIG. 2; EX1002, ¶151.

Elgebaly presents that the techniques may be applied to a number of different types of devices, including “an application specific integrated circuit (ASIC), a digital signal processor (DSP), a digital signal processing device (DSPD), a programmable logic device (PLD), a field programmable gate array (FPGA), a processor, a controller, a microprocessor.” EX1004, [0025]; EX1002, ¶152. However, Elgebaly does not give guidance on how the power domains on ASICs or other devices should be arranged on the silicon chip. *Id.*, ¶153 A POSITA would look to a reference such as Lee to see that one favorable way to arrange the power domains in stripes. *Id.*

Elgebaly's Background discusses that "[a]n effective method for reducing power consumption ... is to scale or adjust the supply voltage." EX1004, [0006]; EX1002, ¶154. Reducing power consumption can be important, for example, to extend battery life. *Id.*; EX1004, [0006]. To this end, and as in Figure 2 (modified to show multiple processor cores and annotated below), Elgebaly discloses an embodiment of ASIC 120 with "adaptive voltage scaling (AVS)," in which "one [green] AVS [adaptive voltage scaling] unit 230 is provided for each [blue, red] processing core 130 ... to adaptively scale the supply voltage for that core." *Id.*, [0031]; EX1002, ¶154. Elgebaly may be extended to "multiple processing cores" in which the clock signal and supply voltage are "jointly set for these cores." EX1004, [0033]; EX1002, ¶155. Further, "ASIC 120 may include . . . any number of processor cores." EX1004, [0030]; EX1002, ¶¶154-155.





FIG. 2

EX1004, FIG. 2 (modified with multiple cores sharing clocks and voltage); EX1002, ¶155.

As shown in the figure above, “[yellow] clock generator 140 generates the clock at the target frequency and provides this clock to the [processor] core.” EX1004, [0032]; EX1002, ¶156. The “target frequency may be selected based on the processing load and computational requirements for the core.” *Id.*; EX1004, [0032].

As discussed above, Elgebaly discloses that the clock generator 140 uses “one or more phase locked loops (PLLs) that control the oscillators” to generate corresponding clock signals. EX1004, [0028]; EX1002, ¶156. A POSITA would

understand how to configure the PLLs to generate the clocks as shown in the modified Figure 2 diagram below, and would naturally turn to a reference such as Lee for the details on implementing the PLLs. EX1010, FIG. 4; EX1002, ¶157.



EX1004, FIG. 2 (modified to show PLLs internal to Clock Generator 140 and multiple processor cores sharing the same clocks and voltage); EX1002, ¶157.

The POSITA would have had a reasonable expectation of success in making the combination, at least because of the specific implementation detail provided by Elgebaly and Lee, as described above; the use of common components (e.g., clock and voltage supply circuitry); and the advanced state of the art as of the priority date.

3. *Element-by-Element Analysis*

a. **Claim 1**

[1.pre] A multi-core processor, comprising:

To the extent the preamble may be limiting, Elgebaly/Lee renders obvious Limitation [1.pre]. EX1002, ¶158. As discussed in Section VI.B.1, Elgebaly discloses an “ASIC 120 [that] includes various processing units that support monitoring and/or communication with one or more communication systems.” EX1004, [0027], FIG. 1 EX1002, ¶158. In the embodiment of Elgebaly shown in Figure 1, the ASIC includes two processor cores: “processor cores 130c and 130d.” *Id.* In addition, “ASIC 120 may include ... any number of processor cores.” EX1004, [0030]; EX1002, ¶158. Thus, Elgebaly’s ASIC is a *multi-core processor*. *Id.*

[1.1] a first set of processor cores of the multi-core processor,


Elgebaly/Lee renders obvious Limitation [1.1]. EX1002, ¶159. Elgebaly teaches that its multi-core processor “may include ... any number of processor cores.” EX1004, [0030]; EX1002, ¶159. 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.* For the reasons explained in Sections VI.B.1-VI.B.2 above, the Elgebaly/Lee Processor depicted in modified Figure 2 teaches a first set of [blue] processor cores (“Processor Core 0” and “Processor Core 1”). Thus, the

Elgebaly/Lee combination renders obvious a second set of processor cores. *Id.*, ¶160.


And thus, Elgebaly renders obvious a *first set of processor cores*. *Id.*, ¶161.

[1.1.1] 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;

Elgebaly/Lee renders obvious Limitation [1.1.1]. EX1002, ¶162. As explained in Sections VI.B.1.b, VI.B.2, Elgebaly discloses that “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.” EX1004, [0033]; EX1002, ¶162. And, as explained in Sections VI.B.1.b, VI.B.2, Elgebaly also discloses that a set of “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1004, [0033]; EX1002, ¶162. The Elgebaly/Lee Processor depicted in modified Figures 1 and 2 below include a *first set of processor cores* (blue “Processor Core 0” and “Processor Core 1”), and a *second set of processor cores* (red “Processor Core 2” and “Processor Core 3”). *Id.*



EX1004, FIG. 1 (modified to show Processors 130c and 130d each comprising 2 cores); EX1002, ¶162.



EX1004, FIG. 2 (modified to show cores 0-1 and 2-3 receiving separate clocks from

PLLs in Clock Generator 140 and separate voltages from Power Management 180); EX1002, ¶162.

Lee's PLL, as shown in Lee's Figure 4 can be combined with Elgebaly's Figure 2. *Id*, ¶163. The combination of Elgebaly and Lee is shown in the following diagram:




Fig. 4 The DVFS scheme for multi-domains power management

first/second clock signal as input
Targ Freq 3 or Targ Freq 4

EX1010, FIG. 4 (modified with labels to show the combination with Elgebaly); EX1002, ¶163.

Consistent with Elgebaly's teaching and as described in Section III.B and throughout the Petition, multi-core processors containing four cores were well known in the art and available commercially well before the '339 Patent's priority date. *See e.g.*, EX1023, 2, EX1004, [0030], [0033]. EX1009, FIG. 1A; EX1002,

¶164. A POSITA would have understood Elgebaly to teach that its methods of voltage and frequency control could be implemented on an ASIC with four cores, such as the one depicted in modified Figure 1. *Id.*

Further, the second modified Elgebaly Figure 2 shows the clocks being driven by PLLs as taught by Elgebaly: “Clock generator 140 may include one or more phase locked loops (PLLs) that control the oscillators used to generate the clocks.” EX1004, [0028]; EX1002, ¶165. In addition, Elgebaly’s teachings that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores” is shown in the modified Figure 2 on sharing of common clocks and voltages to multiple cores. EX1004, [0033]; EX1002, ¶166.

In other words, the common cores Cores 0-1 dynamically receive a first supply voltage (Vdd3). Common cores 0-1 also receive a first output clock signal (Clock3) of a first phase lock loop (PLL) having a first clock signal as input (Target Freq 3). *Id.*, ¶167. A POSITA would have understood that the processor cores could be grouped into two sets, each group managed by one of the AVS units, and the clocks for each set share a clock derived from the same clock generator, both explicitly disclosed in Elgebaly. EX1004, FIG. 2; EX1002, ¶168.

Thus, the Elgebaly/Lee Processor depicted in modified Figure 2, the cores of the first set of [blue] processor cores (“Processor Core 0” and “Processor Core 1”) would be configured to receive a first supply voltage (“Vdd3”) and a first output

clock signal (“Clock 3”), and the cores of the second set of [red] processor cores (“Processor Core 2” and “Processor Core 3”) would be configured to receive a second supply voltage (“Vdd4”) and a second output clock signal (“Clock 4”). EX1002, ¶169.

Additionally, Elgebaly teaches that the supply voltages are “adaptive,” meaning the system “adjusts the supply voltage to a level that satisfies performance requirements,” as they change over time, e.g., when entering an “idle mode.” EX1004, [0034], [0006]; EX1002, ¶170. Thus, in the Elgebaly/Lee Processor, the cores in the first set are configured to dynamically receive a first supply voltage, and the cores in the second set are configured to dynamically receive a second supply voltage. *Id.*; see EX1004, [0033].

Elgebaly’s cores are also configured to dynamically receive respective clock signals. EX1002, ¶171. “Each core 130 provides to clock generator 140 and to its AVS unit 230 a target frequency for its clock.” EX1004, [0032]. “For each processing core 130, clock generator 140 generates the clock at the target frequency and provides this clock to the core as well as to the associated AVS unit 230.” *Id.*

Thus, Elgebaly renders obvious each processor core from the first set of processor cores being 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.* *Id.*, ¶172.

[1.2] a second set of processor cores of the multi-core processor,

Elgebaly/Lee renders obvious Limitation [1.2] for the reasons discussed for Limitation [1.1]. EX1002, ¶¶173-174. Elgebaly discloses that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1004, [0033]; EX1002, ¶174. The Elgebaly/Lee Processor depicted in modified Figures 1 and 2 include a second set of [red] processor cores (“Processor Core 2” and “Processor Core 3”). *Id.* Thus, the Elgebaly/Lee combination renders obvious *a second set of processor cores. Id.*

[1.2.1] 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,

Elgebaly/Lee renders obvious Limitation [1.2.1] for the reasons discussed in for Limitation [1.1.1]. EX1002, ¶175. In the modified Figure 2 above, the common cores Cores 2-3 *dynamically receive a second supply voltage (Vdd4). Id.* Common cores 2-3 also *receive a second output clock signal (Clock4) of a second phase lock loop (PLL) having a second clock signal as input (Target Freq 4). Id.*

[1.2.2] wherein the first supply voltage is independent from the second supply voltage,

Elgebaly/Lee renders obvious Limitation [1.2.2] for reasons discussed for Limitation [1.1.1] and [1.2.1]. EX1002, ¶176. Specifically, Elgebaly teaches that “the supply voltage for each processing core 130 may be set independently by power management unit 180.” EX1004, [0033]; EX1002, ¶176. Similarly, in the

Elgebaly/Lee combination in which one AVS unit is used for each set of cores, a POSITA would understand that the supply voltages (“Vdd3” and “Vdd4”) for each set of cores would be set independently. *Id.*

[1.2.3] and the first clock signal is independent from the second clock signal; and

Elgebaly/Lee renders obvious Limitation [1.2.3] for the reasons discussed for Limitations [1.1.1] and [1.2.1]. EX1002, ¶177. As also explained in Sections VI.B.1-VI.B.2, Elgebaly teaches a clock generator that includes multiple PLLs, one for each independent clock signal, and that “each core 130 provides to clock generator 140 and to its AVS unit 230 a target frequency for its clock.” EX1004, [0032]; EX1002, ¶177. In the modified Elgebaly’s Figure 2, the clock generator receives separate inputs for the target frequencies for groups of cores, e.g. Target Freq 3 (*the first clock signal*) and Target Freq 4 (*the second clock signal*). *Id.* The sources of these signals are from two different processors, hence *the first clock signal is independent from the second clock signal. Id.*


[1.3] an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,

Elgebaly/Lee renders obvious Limitation [1.3]. EX1002, ¶178. Modified Figure 1 of Elgebaly, above, depicts a communication bus between the components of Elgebaly’s multi-core processor that couples the first set of processor cores and to the second set of processor cores. *Id. See supra* Limitation [1.1]. A POSITA would

understand that when processor cores running at different voltages and/or frequencies need to communicate, level translation and synchronization may be required. *Id.* The translation and synchronization are disclosed by Lee's Domain Converter that includes a Level Shifter and Synchronizer as shown in Lee's Figures 2, 3a and 3b. *Id.* The Domain Converter is the *interface block coupled to the first set of processor cores and also coupled to the second set of processor cores. Id.*

[1.3.1] wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.

Elgebaly/Lee renders obvious Limitation [1.3.1]. EX1002, ¶179. Based on Elgebaly's inclusion of internal memory 150 and the knowledge of a POSITA in 2010 that cores may read and write to a common memory space through shared memory and cache coherency (§ III.D), a POSITA would understand that each of the cores of the first and second sets of processor cores may read and write to common space in the memory subsystem, possibly assisted by cache coherency. *Id.* In the Elgebaly/Lee combination, the communication to and from memory 150 may require voltage translation and/or clock synchronization which is supplied by Lee's domain converters (*interface blocks*). *Id.*



EX1004 (modified FIG. 1 with Lee Domain converters); EX1002, ¶179.

The modified Elgebaly's Figure 1 shows the path for Core 1, to write to Internal Memory which is then be read by Core 3. *Id.*, ¶180. A POSITA would understand that this is a typical operation in a multiprocessor, and in the case where Core 1 and Core 3 are in different voltage/frequency domains, would require the Level Shifting and Synchronization provided by Lee's Domain Converters. *Id.*

Accordingly, there is a communication path between the cores via the memory subsystem. In other words, the domain converter is an *interface block*, that is *configured to facilitate communication between the first set of processor cores and the second set of processor cores*. *Id.*, ¶181.

b. Claim 2

Elgebaly/Lee renders obvious Limitation [2.1] for reasons discussed above in Limitation [1.3.1]. EX1002, ¶182. As discussed in Section III.B, it was known that multi-core processors implementing multiple voltage domains required mechanisms for translating logic levels between chip components operating at different voltages. *Id.* A POSITA would have understood Figure 1 of Elgebaly, which depicts communication between processor cores operating using independent voltage levels, to require a voltage level shifter as part of the bus shown in the center of the ASIC. *Id.*; see EX1004, FIG. 1. Indeed, Lee discloses its “LS: Level Shifter” in Figure 3a. EX1002, ¶182. In annotated Lee’s Figure 3, the *first logic level* and *second logic level* are shown as Level1 and Level2. *Id.*, ¶183.




Fig. 3 (a) Level shifter (b) Synchronizer

EX1010, FIG. 3a (annotated); EX1002, ¶182.

Further, there were well-known prior art designs for such circuits. *See e.g.*, EX1037, FIGs. 4-5, 5:45-57, EX1034, FIGs. 2(a), 2(b), 3; EX1002, ¶184. Additionally, a POSITA would have understood that the level shifter that translates between the logic level of the first set to the logic level of the second set must be

referenced to the second supply voltage to know at what logic level the second set is operating. *Id.* For these reasons, it would have been obvious, based on Elgebaly's disclosure, to implement a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set for cache coherency operations between the first and second sets (i.e., "a first signal traveling from the first set of processor cores to the second set of processor cores"). *Id.*

c. Claim 3

Elgebaly/Lee renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1]. EX1002, ¶185. A POSITA would understand that communications in a multiprocessor system is bidirectional. *Id.* There is a need for a *first* and *second level shifter* when cores at different voltage levels communicate in both directions. *Id.* In the modified Elgebaly's Figure 1, the green arrow showing communication from core 1 to core 3 is accompanied by another arrow (not shown) indicating communication from core 3 to core one. *Id.*

d. Claim 5

Elgebaly/Lee renders obvious Limitation [5.1]. EX1002, ¶186. As best understood for the purposes of this petition, "periphery" refers to a physical location of the power control block 108 and clock control block 110 in Figure 1 of the '339 Patent (arranged at the left and right edge of the multi-core processor 100; this

limitation is met at least because Lee discloses PMU1, PMU2, and PMU3 located on the outer edges (e.g., a periphery) of the chip. EX1010, FIGs. 2 and 8; EX1002, ¶¶128, 186.




Fig. 8 Chip micro-photograph

EX1010, FIG. 8 (annotated); EX1002, ¶187.

A POSITA would understand that control signals as recited in the claim would serve to adjust the frequency and voltage of the cores. *Id.* As annotated below, Lee's Figure 2 shows *VDD and CLK signals* communicated from each *PMU* to its corresponding domain (*sets of processor cores*). *Id.* In the Elgebaly/Lee combination, each domain contains multiple cores. *Id.* Hence, Elgebaly/Lee discloses sets of *processor cores [that] are configured to receive one or more*




FIG. 2

EX1004, FIG. 2 (modified); EX1002, ¶188.

Thus, a POSITA would have understood the Elgebaly/Lee combination to render obvious processor cores configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor. EX1002, ¶189.

e. Claim 8

Elgebaly/Lee renders obvious Limitation [8.1]. EX1002, ¶190. As illustrated in Figure 8, Lee explicitly discloses that Domain1, Domain2, and Domain3 are located in separate *regions* of the chip. EX1010, FIG. 8; EX1002, ¶190. Lee in combination with Elgebaly shows that the *first*, *second* and *third regions* could include a *set of processor cores*. *Id.* As discussed in Section VI.B.2, Elgebaly/Lee

includes a first group of cores 130c (the first set of logic cores) that share a voltage signal, clock signal, and cluster-level clock control circuitry and a second cluster of logic cores 130d (the second set of logic cores) that share a different voltage signal, clock signal, and cluster-level clock control circuitry. *Id.*

f. Claim 10

Elgebaly/Lee renders obvious Limitation [10.1] for reasons discussed in Limitation [8.1]. EX1002, ¶191. The die photograph in Figure 8 and block diagram in Figure 2 of Lee clearly show that the three domains (*regions*) are non-overlapping. *Id.*

g. Claim 11

Elgebaly/Lee renders obvious Limitation [11.1] for reasons discussed in Limitation [8.1]. EX1002, ¶192. Further, the die photograph in Figure 8 of Lee explicitly shows that the three domains (*regions*) are in separate rows. *Id.*

h. Claim 14

Elgebaly/Lee renders obvious Limitation [14.1]. EX1002, ¶193. Limitation [5.1] includes all elements of Limitation [14.1] except that in Limitation [14.1], the control blocks are *located in a common region that is substantially central*. EX1002, ¶¶133, 193. A POSITA would have understood that circuitry design involves layout decisions that balance performance, area, and thermal considerations, and that placement of control logic circuitry along the perimeter of the multi-core processor or substantially central to processor cores is an obvious design choice. *See also* Kato

EX1033, FIG. 8, [0056], [0103]-[0107] (depicting chip “floor plan” with “a controller for signal transmission from/to storage control units SC[]” that are “located at the side nearer to the SCs [storage control units],” which is substantially central to the processor units.); EX1002, ¶¶133, 193. Further, a POSITA would understand the benefits of placing control blocks around the periphery of core clusters to shorten the length of wires between the cores and the control logic to shorten delay. *Id.* In addition, this way, core clusters can be placed close to each other and have a shorter inter-cluster communication wires. *Id.*

i. Claim 21

[21.pre]

Elgebaly/Lee renders obvious Limitation [21.pre] for reasons discussed in Limitation [1.pre]. EX1002, ¶194.

[21.1]

Elgebaly/Lee renders obvious Limitation [21.1] for reasons discussed in Limitation [1.1]. EX1002, ¶195.

[21.1.1] wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block;

Elgebaly/Lee renders obvious Limitation [21.1.1] for reasons discussed in Limitation [1.1.1]. EX1002, ¶196. This limitation is the same as Limitation [1.1.1] except that here, the voltage is received *from a power control block*, and the clock

signal is received *as input in a clock control block*. *Id.* In modified Figure 2A, in Elgebaly/Lee cores 0/1 receive a first voltage signal, Vdd3 and cores 2/3 receive a second voltage signal, Vdd4. *Id.* A POSITA would have understood Elgebaly's "other circuits 170" to include Lee's PMUs which include both *power control blocks* and *clock control blocks* to adjust the voltage and clocks of the domains. *Id.*

[21.2]

Elgebaly/Lee renders obvious Limitation [21.2] for reasons discussed in Limitation [1.2]. EX1002, ¶197.

[21.2.1]

Elgebaly/Lee renders obvious Limitation [21.2.1] for reasons discussed in Limitation [1.2.1]. EX1002, ¶198.

[21.2.2]

Elgebaly/Lee renders obvious Limitation [21.2.2] for reasons discussed in Limitation [1.2.2]. EX1002, ¶199.

[21.2.3]

Elgebaly/Lee renders obvious Limitation [21.2.3] for reasons discussed in Limitation [1.2.3]. EX1002, ¶200.

[21.3]

Elgebaly/Lee renders obvious Limitation [21.3] for reasons discussed in Limitation [1.3]. EX1002, ¶201.

[21.3.1]

Elgebaly/Lee renders obvious Limitation [21.3.1] for reasons discussed in Limitation [1.3.1]. EX1002, ¶202.

C. Ground 3: Elgebaly, Lee, and Vajda Render Obvious Claim 9

1. Vajda

Vajda claims priority to a May 26, 2009, PCT filing, and is thus prior art under at least pre-AIA §102(e).³ EX1011; EX1002, ¶203. Vajda discloses “a multi-core chip 230” that includes a number of “processor cores 231 that can execute at higher frequencies” and “other processor cores 232.” EX1011, [0062]-[0063]; EX1002, ¶203. Vajda provides guidance for “the geometry” (i.e., layout of the individual cores) “of the multi-core chip 230.” EX1011, [0063]; EX1002, ¶203. Specifically, Vajda proposes that “the processor cores 231 that can execute at higher frequencies may be spread out as evenly as possible and surrounded by other processor cores 232, running at lower frequencies” because such a layout “may be an advantage primarily in order to avoid hot spots and over-heating of the multi-core chip 230.” EX1011, [0063]; EX1002, ¶203.

Vajda is analogous art to the claimed invention because it is directed to a common field of endeavor—resource management (including power management) for computer processors. (EX1001, 1:6-14; EX1011, Abstract, [0002], [0030], [0037]-[0041], [0085].) By disclosing techniques for scalable and power-efficient

³ The same content is also included in the publication WO 2010/138031. EX1035.

scheduling of processes for multi-core processors, including to operate within defined power consumption limits, Vajda is reasonably pertinent to the problem faced by the '339 patent inventors—optimizing “power consumption management” in multi-core processors. (EX1001, 1:6-14; EX1011, Abstract, [0002], [0030], [0037]-[0041], [0085].)

2. *Motivation to Combine Vajda with Elgebaly and Lee*

A POSITA would be motivated to apply Vadja's chip layout as an alternative to the chip layout shown in the Lee die photo. EX1002, ¶204. A POSITA implementing the Elgebaly/Lee combination may find that in some applications, hotspots and overheating are a problem and, in those cases, opt for the Vadja type of layout to address the problem. EX1002, ¶204. It was well known that multi-core processors feature a number of design constraints such as high temperatures that create overheating risks. *See e.g.*, EX1006, 1:27-52 (overheating limiting performance); EX1011, [0063] (same); EX1021, 1:32-40 (same); EX1002, ¶204. In other cases, where overheating is not a problem, the designer may instead stay with the Lee stripe layout. *Id.* A POSITA would recognize that these references present two viable approaches to layout and may be applicable in particular situations involving various tradeoffs. *Id.* A POSITA familiar with the particular design being implemented would understand which layout is preferable, or if it is not clear, would realize that both are obvious to try. *Id.*

3. *Element-by-Element Analysis*

a. **Claim 9**

[9.1] The multi-core processor of claim 8, wherein the first region and the second region are overlapping regions of the multi-core processor.

Elgebaly/Lee/Vajda renders obvious Limitation [9.1]. EX1002, ¶205. Elgebaly/Lee/Vajda includes processor cores arranged in a checkerboard pattern (e.g., a first row with Processor Core 0 followed by Processor Core 2 and a second row with Processor Core 3 followed by Processor Core 1). *Id.* The region containing the first set of processor cores (i.e., Blue Processor Core 0 and Processor Core 1) is the first region and the region containing the second set of processor cores (i.e., Red Processor Core 2 and Processor Core 3) is the second region. *Id.* The first and second regions are an overlapping region because the checkerboard pattern of cores results in overlapping regions containing each cluster. *Id.*

VII. NOTICES AND STATEMENTS

Pursuant to 37 C.F.R. § 42.8(b)(1), Apple Inc., is the real party-in-interest for Petitioner.

Pursuant to 37 C.F.R. § 42.8(b)(2), Petitioner identifies the following related matters involving the '339 Patent:

Case Heading	Number	Court	Filed
<i>Redstone Logics LLC v. Samsung Electronics Co., Ltd. et al</i>	2-23-cv-00485	E.D. Tex.	October 17, 2023

<i>Redstone Logics LLC v. NXP USA, Inc.</i>	<i>7-24-cv-00028</i>	W.D. Tex.	January 26, 2024
<i>Redstone Logics LLC v. MediaTek Inc. et al</i>	<i>7-24-cv-00029</i>	W.D. Tex.	January 26, 2024
<i>Redstone Logics LLC v. Qualcomm Inc. et al</i>	<i>7-24-cv-00231</i>	W.D. Tex.	September 20, 2024
<i>MediaTek Inc. and MediaTek USA, Inc. v. Redstone Logics LLC</i>	<i>IPR2025-00085</i>	P.T.A.B.	October 22, 2024
<i>NXP USA, Inc. and Qualcomm Inc. v. Redstone Logics LLC</i>	<i>IPR2025-00485</i>	P.T.A.B.	January 22, 2025
<i>Redstone Logics LLC v. NVIDIA Corporation</i>	<i>7-25-cv-00184</i>	W.D. Tex.	April 18, 2025
<i>Redstone Logics LLC v. Apple Inc.</i>	<i>7-25-cv-00183</i>	W.D. Tex.	April 18, 2025
<i>Redstone Logics LLC v. Advanced Micro Devices, Inc.</i>	<i>7-25-cv-00182</i>	W.D. Tex.	April 18, 2025

Pursuant to 37 C.F.R. § 42.8(b)(3), Petitioner identifies the following counsel
(and a power of attorney accompanies this Petition).

Lead Counsel	First Backup Counsel
Alex S. Yap Registration No. 60,609 MORRISON & FOERSTER LLP 707 Wilshire Boulevard, Suite 6000 Los Angeles, CA 90017-3543 Email: AYap@mofo.com Tel: (213) 892-5200 Fax: (213) 892-5454	Mehran Arjomand Registration No. 48,231 MORRISON & FOERSTER LLP 707 Wilshire Boulevard, Suite 6000 Los Angeles, CA 90017-3543 Email: MArjomand@mofo.com Tel: (213) 892-5200 Fax: (213) 892-5454

Backup Counsel	Backup Counsel
Bryan Blumenkopf Registration No. 73,282 MORRISON & FOERSTER LLP 12531 High Bluff Drive, Suite 200 San Diego, CA 92130 Email: BBlumen@mofo.com Tel: (213) 892-5200 Fax: (213) 892-5454	Jean Nguyen Registration No. 71,051 MORRISON & FOERSTER LLP 707 Wilshire Boulevard, Suite 6000 Los Angeles, CA 90017-3543 Email: JNguyen@mofo.com Tel: (213) 892-5200 Fax: (213) 892-5454

Pursuant to 37 C.F.R. § 42.8(b)(4), service information for lead and back-up counsel is provided above. Petitioner consents to electronic service by email to Apple-Redstone-IPR@mofo.com.

Pursuant to 37 C.F.R. § 42.104(a), Petitioner certifies that the '339 Patent is available for *inter partes* review and that Petitioner is not barred or estopped from requesting an *inter partes* review challenging the patent claims on the grounds identified in this Petition.

The USPTO is authorized to charge any required fees, including the fee as set forth in 37 C.F.R. § 42.15(a) and any excess claim fees, to Deposit Account No. **03-1952** referencing Docket No. **10684-00009.02**.

VIII. CONCLUSION

Petitioner respectfully requests that the Board initiate *inter partes* review and cancel claims 1-3, 5, 8-11, 14, and 21 of the '399 patent.

Date: September 16, 2025

By: /Alex S. Yap/
Alex S. Yap
Registration No. 60,609
MORRISON & FOERSTER LLP
707 Wilshire Boulevard, Suite 6000
Los Angeles, CA 90017-3543
Email: AYap@mofo.com
Tel: (213) 892-5200

Counsel for Petitioner

Certification of Word Count (37 C.F.R. § 42.24)

I hereby certify that this Petition for *Inter Partes* Review has 13,310 words (as counted by the “Word Count” feature of the Microsoft Word™ word-processing system), exclusive of “a table of contents, a table of authorities, mandatory notices under § 42.8, a certificate of service or word count, or appendix of exhibits or claim listing.”

Date: September 16, 2025

By: /Bryan Blumenkopf/
Bryan Blumenkopf

Certificate of Service (37 C.F.R. § 42.6(e)(4) and 42.205(a))

I hereby certify that the attached Petition for *Inter Partes* Review and supporting materials were served as of the below date by UPS, which a means as fast and reliable as U.S. Express Mail on the Patent Owner at the correspondence address indicated for U.S. Patent No. 8,549,339.

GENE I. SU
REN-SHENG INTERNATIONAL IP MANAGEMENT LTD
7F, NO. 57, SEC. 2, DUN HUA S. ROAD
TAIPEI, TAIWAN
TELEPHONE: (408) 729-9496

Courtesy copies of the same documents are also being sent to litigation counsel for Patent Owners' exclusive licensee, Redstone Logic LLC via email at:

Reza Mirzaie
Email: rmirzaie@raklaw.com
Marc A. Fenster
Email: mfenster@raklaw.com
Neil A. Rubin
Email: nrubin@raklaw.com
Christian W. Conkle
Email: cconkle@raklaw.com
Jonathan Ma
Email: jma@raklaw.com
RUSS AUGUST & KABAT
12424 Wilshire Boulevard, 12th Floor
Los Angeles, California 90025
Telephone: (310) 826-7474
Facsimile: (310) 826-6991

Qi (Peter) Tong
Email: ptong@raklaw.com
RUSS AUGUST & KABAT
4925 Greenville Avenue, Suite 200
Dallas, TX 75206
Telephone: (310) 826-7474
Facsimile: (310) 826-6991

By: /Bryan Blumenkopf/
Bryan Blumenkopf