

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

DECLARATION OF ROBERT HORST, Ph.D.

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. PROFESSIONAL BACKGROUND	1
III. COMPENSATION	5
IV. LEGAL CONSIDERATION	6
A. Obviousness	6
B. Claim Interpretation	8
V. OVERVIEW OF TASK & BASIS FOR OPINIONS	8
VI. PERSON OF ORDINARY SKILL IN THE ART	10
VII. A POSITA’S UNDERSTANDING OF THE FIELD OF TECHNOLOGY	11
A. Multi-core Processors	11
B. Clock Generation	13
C. Dynamic Voltage Frequency Scaling (DVFS)	19
D. Processor Cache Memory and Cache Coherency	22
VIII. OVERVIEW OF THE ’339 PATENT	25
IX. IEEE PUBLICATION PRACTICES RELEVANT TO THE WILLIAMS AND LEE PAPERS	29
X. GROUNDS CHALLENGING THE ’339 PATENT	31
A. Ground 1: White, Talwar and Lee Render Obvious Claims 1-3, 5, 8, 10-11, 14, and 21	31
1. White	31
2. Talwar	35
3. Lee	38
4. Motivation to Combine	42
5. Modification of White in View of Talwar and Lee	46
6. Element-by-Element Analysis	50
a) Claim 1	51

b)	Claim 2	63
c)	Claim 3	66
d)	Claim 5	66
e)	Claim 8	70
f)	Claim 10	70
g)	Claim 11	71
h)	Claim 14	71
i)	Claim 21	72
B.	Ground 2: Elgebaly and Lee Render Obvious Claims 1-3, 5, 8, 10-11, 14, and 21	75
1.	Elgebaly	75
2.	Motivation to Combine Elgebaly and Lee	77
3.	Element-by-Element Analysis	82
a)	Claim 1	82
b)	Claim 2	94
c)	Claim 3	96
d)	Claim 5	96
e)	Claim 8	100
f)	Claim 10	101
g)	Claim 11	101
h)	Claim 14	102
i)	Claim 21	103
C.	Ground 3: Elgebaly, Lee and Vajda Render Obvious Claim 9	105
1.	Vajda	105
2.	Motivation to Combine	106
3.	Element-by-Element Analysis	106
a)	Claim 9	106

Exhibit List

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.	1014

Exhibit Description	Exhibit #
https://ieeexplonorthbridre.ieee.org/document/4137045 (last accessed on 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 (last accessed on 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 (last accessed on 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

Exhibit Description	Exhibit #
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
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

I, Dr. Robert Horst, Ph.D., make this declaration in connection with the proceeding identified above.

I. INTRODUCTION

1. I have been retained by counsel for Apple Inc. (“Petitioner” or “Apple”) as a technical expert in connection with the proceeding identified above. I submit this declaration in support of Apple’s Petition for *Inter Partes* Review (“IPR”) of United States Patent No. 8,549,339 (“the ’339 Patent”).

II. PROFESSIONAL BACKGROUND

2. All of my opinions stated in this Declaration are based on my own personal knowledge and professional judgment. In forming my opinions, I have relied on my knowledge and experience in designing, developing, researching, and teaching the technology referenced in this Declaration.

3. I am over 18 years of age and, if I am called upon to do so, I would be competent to testify as to the matters set forth herein. I understand that a copy of my current curriculum vitae, which details my education and professional and academic experience, is being submitted as EX1003. The following provides a brief overview of some of my experience that is relevant to the matters set forth in this Declaration.

4. I am an expert in the field of computer architecture, including power efficient multi-core computer architecture and design, electronics systems, power

systems and control electronics, processors and interconnect, and computer hardware.

5. I am currently an Adjunct Research Professor in the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign and am also an independent consultant at HT Consulting. I have more than 30 years of expertise in designing and architecting computer systems.

6. I am also currently Principal at Horst Tech LLC, where I perform research, design, and manufacturing of electronic systems and devices.

7. I earned my M.S. (1978) in electrical engineering and Ph.D. (1991) in computer science from the University of Illinois at Urbana-Champaign after earning my B.S. (1975) in electrical engineering from Bradley University. During my master's program, I designed, constructed, and debugged a shared memory parallel microprocessor system. During my doctoral program, I designed and simulated a massively parallel, multi-threaded task flow computer.

8. After receiving my bachelor's degree and while pursuing my master's degree, I worked for Hewlett-Packard Co. While at Hewlett-Packard, I designed the micro-sequencer and cache of the HP3000 Series 64 processor.

9. From 1980 to 1999, I worked at Tandem Computers, which was acquired by Compaq Computers in 1997. While at Tandem, I was a designer and architect of several generations of fault-tolerant computer systems and was the

principal architect of the NonStop Cyclone superscalar processor. The system development work at Tandem also included development of the ServerNet System Area Network and applications of this network to fault tolerant systems and clusters of database servers.

10. Since leaving Compaq in 1999, I have worked with several technology companies, including 3Ware, Network Appliance, Tibion, and AlterG in the areas of network-attached storage and biomedical devices, including the development of microprocessor and FPGA electronics as well as interconnect options for network-attached storage subsystems. I also represented Network Appliance in the PCI Express Advanced Switching working group.

11. At HP, Tandem, Compaq, 3Ware and Network Appliance my computer design work was done using computer aided design (CAD) tools, with most designs specified in a hardware description language. In fact, the designs used integrated circuits and some involved the creation of new ASICs (application specific integrated circuits) that were manufactured using masks based on the netlists produced by the CAD software.

12. From 2012 to 2015, I was Chief Technology Officer of Robotics at AlterG, Inc., where I worked on the design of anti-gravity treadmills and battery-powered orthotic devices to assist those with impaired mobility.

13. In 1998, the University of Illinois department of Electrical and Computer Engineering awarded me the Distinguished Alumni Award for “Pioneering Contributions to Fault-tolerant Computer Architecture.”

14. I have served on various IEEE committees and in 2001, I was elected an IEEE Fellow “for contributions to the architecture and design of fault tolerant systems and networks.”

15. I have authored over 30 publications, including, for example, “A Linear-Array WSI Architecture for Improved Yield and Performance,” in Proc. Int. Conf. WSI, San Francisco, CA, pp. 85-91, Jan. 1990; “Task Flow Computer Architecture,” in Proc. Int. Conf. Parallel Processing, Vol. I, pp. 533-540, Aug. 1990; “Task Flow: A Novel Approach to Fine-grain Wafer-scale Parallel Computing,” Coordinated Science Lab. Report CRHC-91-15, University of Illinois, April 1991; “Multiple Instruction Issue in the NonStop Cyclone Processor,” in Proc. 17th Int. Symp. Computer Architecture, May 1990; “Task-Flow Architecture for WSI Parallel Processing,” Computer, vol. 25, no. 4, pp. 10-18, April 1992.

16. I have also been published in the areas of Networks, Storage, CPU Architecture, Fault Tolerance, and Bionics.

17. I have also worked with patent attorneys on numerous patent applications, and I am a named inventor on 87 issued U.S. patents.

18. I have my own patents directed to clocking, memory system design, and aspects of circuit design, including U.S. Pat. No. 5,930,275, U.S. Pat. No. 5,146,589, U.S. Pat. No. 5,287,472, U.S. Pat. No. 5,329,629, U.S. Pat. No. 5,034,964, and U.S. Pat. No. 9,893,604.

19. In addition, some of my patents include those related to multi-processor systems including Patent No. 6,496,940 (Multiple processor system with standby sparing), U.S. Patent No. 5,751,932 (fail-fast, fail-functional, fault-tolerant multiprocessor system) and U.S. Patent No. 5,390,355 (computer architecture capable of concurrent issuance and execution of general-purpose multiple instructions), and storage including U.S. Patent No. 6,549,977 (the use of deferred write completion interrupts to increase the performance of disk operations).

20. I have testified as an expert witness on over 50 patent matters related to systems, processors, networks, and storage. I have also served as an expert witness in patent and intellectual property litigation as well as on over 30 *inter partes* review matters.

III. COMPENSATION

21. My work on this matter is being billed at my customary rate of \$750 an hour. Also, I am being reimbursed for reasonable expenses I incur in relation to my services. I have no financial interest in the outcome of this proceeding. I

understand I will be paid regardless of the outcome of any proceeding in which my work is used.

IV. LEGAL CONSIDERATION

22. I am not an attorney. My understanding of the law is based on information provided by counsel for Petitioners.

A. Obviousness

23. I understand that a claimed invention is obvious and, therefore, not patentable if the subject matter claimed would have been obvious to a person of ordinary skill in the art at the time that the invention of the '339 Patent (a "POSITA"), which I have been asked to treat as February 26, 2010.

24. I understand that a claim can be obvious in view of a single prior art reference (e.g., via modification of that prior art reference) or multiple prior art references (e.g., via a combination of two or more prior art references), if such a modification or combination was within the skill of a POSITA. I understand that there must be some articulated reasoning with some rational underpinning to support a conclusion of obviousness. I also understand that to establish a finding of obviousness one must show that one of ordinary skill in the art would have had a motivation to combine the prior art references to produce the claimed invention and a reasonable expectation that the combination would be successful.

25. I further understand that exemplary rationales that may support a conclusion of obviousness include: (1) simply arranging old elements in a way in which each element performs the same function it was known to perform, and the arrangement yields expected results, (2) merely substituting one element for another known element in the field, if the substitution yields no more than a predictable result, (3) combining elements in a way that was “obvious to try” because of a design need or market pressure, where there was a finite number of identified, predictable solutions, (4) that design incentives or other market forces in a field would have prompted variations in a work that were predictable to a person of ordinary skill in the art, and (5) that there was some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify or combine prior art references to arrive at the claimed invention. I further understand that certain factors may support or rebut the obviousness of a claim. I understand that such secondary considerations include, among other things, commercial success of the patented invention, skepticism of those having ordinary skill in the art at the time of the invention, unexpected results of the invention, any long-felt but unsolved need in the art that was satisfied by the alleged invention, the failure of others to make the alleged invention, praise of the alleged invention by those having ordinary skill in the art, and copying of the alleged invention by others in the field. I understand that there must be a nexus—that is, a connection between

any such secondary considerations and the alleged invention. I also understand that contemporaneous and independent invention by others is a secondary consideration tending to show obviousness. I am not aware of any allegations by the named inventor of the '339 Patent or any assignee of the '339 Patent that any secondary considerations tend to rebut the obviousness of the '339 Patent.

B. Claim Interpretation

26. I understand that a claim term is interpreted according to its ordinary and customary meaning as a POSITA would have understood the term in light of the surrounding claim language, other claims, the specification, and the patent's prosecution history, which are referred to as intrinsic evidence. I also understand that prior art references cited in the patent's prosecution history are considered intrinsic evidence. I further understand that evidence outside the patent and its prosecution history (e.g., dictionaries and technical articles), may inform the context in which a POSITA would have understood the claims of a patent. I understand this ordinary and customary meaning applies absent unique circumstances, such as where a patent clearly expresses an intent to set forth a special meaning for a term.

V. OVERVIEW OF TASK & BASIS FOR OPINIONS

27. I have been asked to review the '339 Patent. I have been asked to provide opinions related to certain issues from the perspective of a person of

ordinary skill, having knowledge of the relevant art, as of February 26, 2010, and—except where otherwise noted—the opinions stated in this declaration are from that perspective. The qualifications and abilities of such a person are described in paragraphs 31-32 below.

28. I was further asked to review the documents filed in IPR 2025-00485 filed by NXP USA, INC. and QUALCOMM INCORPORATED, including the Declaration of Murali Annavaram, Ph.D, (EX1013), and the exhibits cited therein. Parts of this declaration are based on this material.

29. In preparing this declaration, I have considered this patent and its prosecution history as well as the general knowledge of those familiar with the field of technology relating to multi-core processors, which includes techniques for generating and managing distinct clock signals, power management techniques such as dynamic voltage frequency scaling, intra-chip communication techniques, efficient memory hierarchy designs such as caches and cache coherency protocols, and common semiconductor die layouts.

30. My opinions are based on my education, training, and experience as well as items that I reviewed to prepare my opinions, including the documents in the Exhibit List included with the Petition, and any other items that I reference in my below analysis.

31. I understand that other issues may arise that require further explanation, and I will provide that explanation if appropriate. As a result, I respectfully reserve the right to update and supplement this declaration and the information and opinions provided herein.

VI. PERSON OF ORDINARY SKILL IN THE ART

32. Based on the disclosure of the '339 Patent and the priority date, a person of ordinary skill in the art ("POSITA") pertinent to the '339 Patent would have had at least (1) a bachelor's degree in electrical engineering, computer engineering, computer science, or a similar field, and (2) at least two years of academic or industry experience designing or analyzing electronic circuits, semiconductors, processors, or power management, and related firmware and software, or the equivalent. The prior art also evidences this level of ordinary skill. Individuals with different education and additional experience could still be of ordinary skill in the art if that experience compensates for a deficit in their education stated above. Additional education can substitute for some experience, and vice versa.

33. Here, the background technology described in Section VII and the prior art described in Section X 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 frequency scaling, intra-chip communication techniques, memory hierarchy designs, and common semiconductor die layouts.

VII. A POSITA’S UNDERSTANDING OF THE FIELD OF TECHNOLOGY

A. Multi-core Processors

34. Multi-core processors were well known prior to the priority date of the ’339 Patent. 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.

35. 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].

36. 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]. 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. 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].

37. Multi-core processors feature a number of known design constraints that must be balanced to create an effective design, including power consumption, high temperatures that create overheating risks, high system complexity that limits performance, and limited chip space that restricts the number of components that may be cost-effectively included on any given chip. *See e.g.*, EX1006, 1:14-26 (limited chip space); EX1029, 4 (same); EX1006, 1:27-52 (overheating limiting performance); EX1011, [0063] (same); EX1021, 1:32-40 (same); EX1007, [0002]-[0004] (desire to minimize power consumption); EX1004, [0006] (same); EX1009, [0010] (same).

38. Of particular relevance to this Petition, it was known in the art that circuitry to support “the regulation of operating voltage for integrated circuits, such as microprocessors” requires “large values of capacitance which may consume inordinate amount of real estate on an IC die.” EX1006, 1:14-26; *see also* EX1024,

1 (describing how “conventional” voltage regulator designs would require too much space on the motherboard to be “practical” for 2000s-era microprocessors requiring high currents and well-regulated voltages). In part because of the amount of space required for voltage regulator circuitry, multi-core processors often “don’t have the same number of individual voltage sources [as] processors.” EX1007, [0003]. However, restricting the number of voltage sources to a single source is also sub-optimal since that would force all the components on a chip to operate at a single voltage. *See also* EX1007, [0004]; EX1019, 52-53. It is well known in art that a small number of heterogeneous mix of voltages (and correspondingly a few frequencies) would yield the best design. *See e.g.*, EX1007, [0011], FIG. 1; EX1019, 53 (describing the benefits of independently scaling the supply voltage and frequency of each core).

B. Clock Generation

39. The basic design principles for processor clock signals, including their generation, distribution, and synchronization, were widely known by 2010. The basic design to generate a single clock signal uses an external reference oscillator (e.g., crystal oscillator) paired with a phase locked loop (PLL) to multiply or divide the reference oscillator’s signal to reach the desired clock frequency. EX1028, 2:61-3:12. Other designs are also common. For example, a clock signal may be provided via one or more of: (a) an external reference oscillator (e.g., crystal

oscillator), (b) a PLL, and (c) “clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals.” *Id.*; EX1006, 7:23-45.




FIG. 1

EX1028, FIG. 1 (annotated). As shown in the figure above, the key components of a PLL are a phase/frequency detector, charge pump, low-pass filter (LPF), voltage-controlled oscillator (VCO), and N divider that work together to align the output signal’s phase with the phase of the reference crystal oscillator. *Id.* The basic PLL clock design generates a clock signal at a single frequency.

40. In addition to the basic PLL clock design, it was well known in the art that multiple PLLs, multiple crystal oscillators, or a combination of both may be used to create multiple clock signals at different frequencies. *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]. Such a design is useful, for example, in a processor where different functional units run at different speeds using different clock frequencies, also known as operating in different clock domains. *See* EX1008, 1-2. It was well known in the art that it was beneficial for some components of the processor to operate in one clock domain, while other components operate in a different clock domain. *Id.* The use of multiple clock domains within a processor allows certain components, such as memory, to run at lower frequencies and consume less power, while other components, such as adders or other computational units to run at higher frequencies for improved performance.

41. An example of a known design disclosed in Williams to generate two independent clock domains is shown below:




Figure 1: Generic clock domain diagram

Id., FIG. 1. As shown in the figure, the two clock domains are “clocked by different PLLS . . . [that] run off reference clocks produced by different clock generators and different crystal oscillators.” *Id.*, 1. Specifically, the clock signal in “Clock Domain 1” is derived from the [green] “Reference Clock Distribution 1” circuit that includes “Crystal 1” and the independent clock signal in “Clock Domain 2” is derived from the [red] “Reference Clock Distribution 2” circuit that includes “Crystal 2.” *Id.*

42. In addition to the multiple crystal oscillator, multiple PLL design above, it was also known in the art that clock signals at different frequencies may be generated using a clock tree or other circuitry configured to split or divide a single clock source into multiple signals at different frequencies. *See e.g.*, EX1006, 7:23-45; EX1027, [0027], FIG. 3. An example of this kind of configuration is depicted below:




Fig. 3

EX1027, FIG. 3.

43. It was also known to generate clocks from multiple PLLs, with each PLL receiving an independent clock at a frequency determined by a divider. An example of this technique was shown by Lee Paper Two, et. al:




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

EX1031, FIG. 2.

44. In this design, a domain specifies a desired frequency with an 11-bit control signal that can specify a frequency that is a factor of 5 to 283 times the Reference Clock frequency. In other words, the clock signal fed to the Phase Frequency Detector is the target frequency divided by a number from 5 to 283. In a separate paper presented by Lee in the same year (2007), “Lee”, Lee shows a design with multiple domains, each with a PLL that has its frequency set in such a way by independent clock signals.

45. As explained further in the following section on DVFS, it was also known that each frequency has an ideal supply voltage to sustain that frequency and hence multiple clock domains may also benefit from having multiple such voltage sources.

C. Dynamic Voltage Frequency Scaling (DVFS)

46. Dynamic voltage and frequency scaling, or DVFS, describes a broad range of well-known prior art methods to reduce microprocessor power consumption using “dynamic adjustments to supply voltage, clock frequency, or both.” EX1019, 35. 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. But a higher clock rate comes with costs; voltage must be increased to support the clock rate, increasing both power drain and heat. 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. By 2010, DVFS was “[t]he most popular technique for CPU power management.” EX1007, [0002].

47. Fundamentally, DVFS techniques are derived from the basic power equation for processors: $P = CV^2Af$, where f is processor frequency, V is voltage, C is capacitance, and A is an activity factor related to how often transistors switch on a particular processor. EX1019, 35. This equation demonstrates that, “[i]f we can reduce voltage by some small factor, we can reduce power by the square of that factor,” while keeping activity and capacitance unchanged. *See id.* However, “reducing supply voltage often slows transistors such that reducing the clock

frequency is also required.” *Id.*, *see also* EX1007, [0002] (“the frequency of the CPU is roughly linear in voltage”). Reducing a processor’s clock frequency “may also linearly degrade performance,” which can cause programs to “run[] at lower power dissipation levels, but for longer durations.” EX1019, 35. As a result of this trade-off, DVFS techniques focus on ways “to reduce clock frequency without having the workload experience a proportional reduction in performance.” *Id.*, 36.

48. In multi-core processors using DVFS techniques, it was known to divide cores and other components into multiple voltage and/or clock domains that are adjusted independently. EX1019, 52-53; EX1006, 2:15-44; EX1004, [0032]. By 2010, “[i]ndependent DVFS for each core,” “per-chip DVFS,” and “a mixture chip-wide DVFS and core allocation,” were all known configurations of DVFS for multi-core processors. EX1019, 52-53. Additionally, it was known that “synchronizer circuits are required to avoid metastability between domains” operating at different clock frequencies.” EX1019, 52-53; *see e.g.*, EX1032, FIGs. 1-3D, 1:55-2:4.

49. 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, FIGs. 2(a), 2(b), 3. “A logic-swing level shifter is a key circuit component” in

DVFS implementations both because, often, “I/O blocks are operated under fixed V_{DD} ” and each DVFS block on a chip could be operating under different voltages at any given point in time. EX1034, 1. Thus, it is possible that a high voltage level signal of a first voltage domain may be recognized as a low voltage level by a second voltage domain, and vice versa, resulting in communication errors. Level shifter circuits address this problem. An example of a conventional level shifter circuit that was known in 2005 is shown below:




Figure 2 (a) Conventional level shifter.

Id., 2. The conventional design shown in Figure 2(a) operates as follows:

“[a]ssuming that initially input signal, IN, is ‘[Low, or a digital 0]’ MN1 is turned on and provides a conducting path to ground while MN2 is cut off. Therefore, node A is pulled down to ground, which makes MP2 on and node B is pulled up to ‘[High, or a digital 1].’ Thus, the output signal, OUT, becomes ‘[Low, or a digital 0].’” *Id.*, 1. If IN is a digital 1 in the first voltage domain (i.e., V_{DDL}), then node B is pulled down to ground, before being inverted to V_{DDH} so that OUT reflects a

digital 1 in the second voltage domain (i.e., V_{DDH}). *See id.* Thus, voltage signals used by the circuits in the first voltage domain can be understood by circuits in the second voltage domain. Similar designs were well known in the art and an essential feature of chip designs that use multiple voltage domains.

D. Processor Cache Memory and Cache Coherency

50. 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 have to access the data in slower memory. *See* EX1026, 1; EX1030, 1. By 2010, it was well known that multiple levels of caches was an optimal design, balancing the trade-off between memory speed and memory size. *See e.g.*, EX1006, FIGs. 1, 8. For example, to keep up with ever increasing processor clock speeds and “[t]o further improve speed of access to data within such a multi-processing system, it is known to provide each of the processing units with its own local cache.” EX1009, [0005]. An example of this cache design, also known private caching, is shown in the figure below:




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

EX1030, FIG. 2. As shown in the figure, each core may have a local cache that is unique to each processor core (P) and not directly accessible by other cores, hence the local cache is also termed as a private cache. EX1030, 1; *see also* EX1009, FIG. 1A. In addition, 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], *see e.g.*, EX1023 5; EX1006, FIGs. 1, 8. Shared memory has the added benefit that it allows multiple, “low-cost microprocessors [to be] economically interconnected with shared memory modules,” making this design particularly “cost-effective.” EX1030, 1.

51. Because prior art multi-core processors included both shared main memory and dedicated local caches, it became “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. The process of ensuring all processors operate using up-to-date data is known as cache coherency and is an essential feature of processor designs that use a shared memory. *See* EX1009, [0006]-[0009]; EX1030, 1.

52. Cache coherency schemes for multi-core processors have been well known since at least 1990. *See* EX1030, 1-2. Hardware-based cache coherency protocols were also well known since at least 1990, including “write-invalidate” and “write-update” schemes. EX1030, 3. In a write-invalidate scheme, when a first processor writes to a memory address, any copies of that memory address stored in the other processors’ local caches are invalidated. *Id.* In a write-update scheme, by contrast, instead of invalidating the copies held in the other processors’ local caches, the system updates them to reflect the modified value written to the memory address by the first processor. *Id.* For example, the cache of the processor that writes to the shared data may send the updated copy of the data to any other cache that holds a local copy of that data along a bus.

53. Both write-invalidate and write-update protocols “require that cache invalidation and update commands . . . be sent to at least those caches having copies of the [data] block.” *Id.*, 4. One such approach to transmitting these invalidation and update commands is to use protocols called “snoopy cache protocols because every cache ‘snoops’ on the network for every incoming

consistency command.” *Id.*, *see also* EX1009, [0009]. A bus on the microprocessor is one example of a common “network” used to implement snooping operations. EX1030, 4; *see e.g.*, EX1009, FIG. 1A. Alternate “network” configurations are also well known, such as mesh, ring, and hyper-cube topologies.

VIII. OVERVIEW OF THE ’339 PATENT

54. The ’339 Patent relates to “power management for a multi-core processor” and related “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. For example, “[b]ased on the measured operational information (e.g., supply voltage and local temperature) and the performance data, the power control block 108 may then be configured to select a supply voltage for each strip.” *Id.*, 2:50-54.




FIG. 1

Id., FIG. 1 (annotated). Figure 1 depicts a multi-core processor divided into rows of processor cores, or “stripes 112 [green], 114 [red], 116, and 118” that “may be associated with an independent power profile.” *Id.*, 2:25-27. By contrast, the ’339 Patent asserted that “[e]ach processor core in a conventional multi-core processor generally shares the same supply voltage and clock signal to simplify the interfaces between the processor cores.” *Id.*, 1:7-10.

55. The ’339 Patent associates each stripe with an “independent power profile” to address this perceived problem in the prior art. *Id.*, 2:26-27, 1:10-14. A “power profile may include . . . power-supply voltages . . . clock rates . . . clock

multipliers . . . power throttling . . . and/or sleep state cycles of the core processor.” *Id.*, 1:62-2:3. For example, each “stripe,” or set of processors “may be powered by a supply voltage received from a power control block 108 and/or may be associated with an independent clock domain defined by a clock signal received from a clock control block 110.” *Id.*, 2:28-31. Each stripe’s power profile “may be determined based on the computational requirements of the tasks assigned to the processor cores in the stripe.” *Id.*, 2:41-43. Thus, the ’339 Patent’s goal of “operat[ing] at high power and high clock frequency when needed and at low power when the computing requirements are reduced” can be achieved on a stripe-by-stripe basis. *Id.*, 1:10-14.

56. The ’339 Patent also states that “The multi-core processor 100 may be further divided into regions. In some implementations, the regions of multi-core processor 100 may correspond to rows of the two-dimensional array, and the regions may or may not be overlapping.” *Id.*, 2:20-23.

57. The ’339 Patent also discloses an “interface block” to facilitate communications between processor cores operating at different voltage levels and using different clocks. *Id.*, FIGs. 2-3. By passing inter-processor communications through an “interface block,” signals from different processors operating at different supply voltages and driven by independent clock signals can be

“translate[d] . . . such that each of the processor cores operates correctly.” *Id.*, 3:47-52, Abstract. This process is depicted in the figures below:




FIG. 2

FIG. 3

EX1001, FIGs. 2, 3. In these figures, “the processor core 152 belongs to the [green] stripe 112; the processor core 154 belongs to the [red] stripe 114; and the processor core 156 belongs to the stripe 116.” *Id.*, 3:21-26.

58. FIG. 2 (above) depicts a process for translating communications from cores operating a different voltage levels. As shown in the figure, “the level shifter 202 may be adapted to translate first logic levels associated with the [green] stripe 112 to second logic levels associated with the [red] stripe 114.” *Id.*, 3:52-54. This translation allows each processor core to “properly interpret the voltages as valid

logic levels even though [the] processor cores are powered by different supply voltages.” *Id.*, 3:50-52.

59. FIG. 3 (above) depicts a process for translating communications from cores operating using different, independent clock signals. As described above, each stripe is “associated with an independent clock domain.” *Id.*, 2:29-30. Being in different stripes, each processor core is clocked by the output of a PLL whose input is a different, independent clock signal: “clock signal 1,” “clock signal 2,” and “clock signal 3.” *Id.*, 4:4-10. To “facilitate communication between two processor cores,” the “[yellow] interface blocks” include “synchronizer 302” and “synchronizer 306.” *Id.*, 3:21-24, 4:1-10, Abstract.

IX. IEEE PUBLICATION PRACTICES RELEVANT TO THE WILLIAMS AND LEE PAPERS

60. In addition to my review of the ’339 Patent and the Petition, I have been asked to attest to my experience with the Institute of Electrical and Electronics Engineers (IEEE) and their publication practices. I have also been asked to review the Williams, Lee paper, and Lee II as downloaded from IEEE’s website and their related metadata.

61. I have served on various IEEE committees and I was elected an IEEE Fellow “for contributions to the architecture and design of fault tolerant systems and networks”. As such, I am generally familiar with IEEE’s publishing practices. My papers have been published through IEEE more than a dozen times, ranging

from my first publication in 1989 to today. Many of these papers were presented publicly at a conference, before being published either on IEEE's website, in print, or both. In most cases, conference papers are publicly available on IEEE's website shortly (e.g., within a month) after the conference occurs.

62. I have reviewed the IEEE website page for the Williams paper (EX1008) and can confirm that it lists the paper as having been "[a]dded to IEEE *Xplore*" on March 26, 2007. *See also* EX1014, 1. Based on my experience with IEEE and my familiarity with IEEE's publishing practices, this means that the paper was publicly available on IEEE's website as of that date.

63. I have also reviewed the IEEE website page for the Lee paper (EX1010) and can confirm that it lists the paper as having been "[a]dded to IEEE *Xplore*" on January 7, 2008. *See also* EX1015, 1. Based on my experience with IEEE and my familiarity with IEEE's publishing practices, this means that the paper was publicly available on IEEE's website as of that date.

64. Further, I have also reviewed the IEEE website page for the Lee II paper (EX1031) and can confirm that it lists the paper as having been "[a]dded to *IEEE Xplore*" on June 25, 2007. *See also* EX1016, 1. Based on my experience with IEEE and my familiarity with IEEE's publishing practices, this means that the paper was publicly available on IEEE's website as of that date.

X. GROUNDS CHALLENGING THE '339 PATENT

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

1. *White*

65. White published on August 28, 2007, and is thus prior art under at least pre-AIA §102(b). EX1006. 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” based on various considerations including power and temperature management. EX1006, 2:15-44.

66. A POSITA 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). *See supra* Section VII.C.

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




FIG. 2A

EX1006, FIG. 2A (annotated).

68. “[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.” *Id.*, 8:59-65. “[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,” as illustrated by the two separate lines extending from [blue] voltage regulator 130 to [green, red] logic cores 120A, 120B, respectively. *Id.*, 7:23-45, 10:31-32, FIG. 2A. Similarly, “[purple] frequency regulator 140” may “supply multiple, independent clock signals” to the [green, red] logic cores 120A and 120B to allow “each to operate at frequencies and voltages independent of the frequencies or voltages that the other logic cores, or other components, may be operating.” *Id.*, 6:41-49, 7:23-

45. “[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 “with [purple] frequency regulator 140 to instruct frequency regulator 140 to supply a particular frequency or clock signal to [green] logic core 120A.” *Id.*, 10:11-24. Further, “as with power management logic 200, described above, temperature control logic 210 may, in some embodiments be configured to communicate or interface with voltage regulator 130 and frequency regulator 140 in order to adjust the voltages and frequencies supplied to logic cores 120A and 120B.” *Id.*, 13:38-43. Accordingly, each logic core is arranged to receive an independent voltage and clock signal.

69. White discloses multiple options for generating and changing clock signals using a combination of elements at both frequency regulator 140 and the cores. 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.*, 7:38-45. Frequency regulator 140 may also include “clock trees or other circuitry configure to split or divide a clock signal into multiple individual signals.” *Id.*, 7:40-43. Additionally, each logic core may also “generate its internal clock using a phase-locked loop (PLL) locked to an

externally provided clock signal, such as may be provided by frequency regulator 140.” *Id.*, 9:62-65.

70. White also discloses the use of “common bridge logic 110, such as a Northbridge,”—a well-known means for communicating between the cores of a multi-core processor and its memory—“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. Because the logic cores may be operating at different voltage levels and frequencies, White discloses that the “[c]ommon bridge logic 110” is “supplied with multiple voltages and/or frequencies” so that it can “interface with a logic core using the same logic levels and/or clock timing as the logic core.” *Id.*, 7:64-8:10. In addition, because different voltage levels and frequencies are generally incompatible with one another, common bridge logic further includes an interface “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.” *Id.*, 8:21-26, FIG. 3. A POSITA would understand this to mean that the common bridge logic includes one of the 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* Section VII.C, *see also* EX1010, FIG. 3, EX1034, 1-2, Figures 2(a), 2(b), 3.

71. 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...[where] a processor may include multiple logic cores, or processor cores, configured to operate at frequencies and voltages independently of each other.” EX1006, 6:56-60; *see also* 6:63-67, claims 1, 25 (“wherein at least one of the logic cores is configured to operate at frequency independent of a frequency at which another of the plurality of logic cores operates”).


2. *Talwar*

72. 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. 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 [have] the same frequency.” EX1007, [0011], [0018]-[0019]. By providing a “single voltage source” to “multiple cores in a cluster,” “the frequency of all cores in a cluster” and the single voltage source can be jointly adjusted “[d]uring power management.” *Id.*, [0011]-[0012], [0018]-[0019], [0021]-[0025], [0030] (“The cores in a cluster, which receive power in one power domain, can be independently controlled from other clusters.”).

73. Similar to both the '339 Patent and White, Talwar discloses the use of “[d]ynamic voltage frequency scaling (DVFS)” to reduce CPU power consumption, which is “[t]he most popular technique used for CPU power management.” *Id.*, [0002]. DVFS “reduces the frequency [of the CPU] when the CPU utilization is below a certain threshold.” *Id.*, [0003]. Additionally, as was well known in the art, “if the CPU frequency is reduced, the required voltage is reduced” because “the frequency of the CPU is roughly linear in voltage.” *Id.*, [0002]. Talwar recognized, however, that this approach is “not as efficient in multi-core systems” because “they don’t have the same number of individual voltage sources” as processor cores. *Id.*, [0004]. For example, if “one processor’s frequency is scaled down, the voltage to the processor doesn’t change because the other processor [on the same chip] is still running at a higher frequency and needs the higher voltage.” *Id.*, [0004].

74. Talwar’s design solved the problem that “in current multi-core systems, all the processors use a single voltage source which renders frequency scaling [power management] technique[s] often inefficient.” *Id.*, [0003]-[0004]. By grouping the cores into “clusters,” Talwar’s design gains the power efficiency of multiple DVFS groups of cores per chip without requiring a unique voltage source for every single core. Additionally, Talwar’s design scales to “any number of voltage sources and cores,” while still requiring “less voltage sources than cores on

the chip.” *Id.*, [0014]. Providing one voltage (and clock) source per cluster enables efficiently performing DVFS on each cluster in “stark contrast to conventional multi-core chips” with only a single voltage source. *Id.*, [0012].




EX1007, FIG. 1 (annotated).

“maintains all the cores in a cluster at the same frequency” and “[e]ach cluster may have a different cluster frequency.” EX1007, [0018]-[0019]. By using the same supply voltage and frequency for each cluster, “maximum power savings” can be achieved (i.e., some clusters can be run at higher frequencies and voltages, while others can be dynamically set to lower frequencies and voltages). EX1007, [0021].

3. *Lee*

77. “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. Lee discloses a “DVFS scheme for Multi-power domains.” It 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.



EX1010, Fig. 2.

78. Lee also discloses techniques for communicating between the domains:

In this case, **each domain has different clock and supply voltage.** System has to convert the voltage level and adjust the phase of data which goes from domain 1 to domain 2 and from domain 2 to domain 3. The level shifter (LS) and synchronizer (Sync.) in Fig 3 change the voltage level and convert the clock frequency of transferred data, respectively. Therefore, the domain converters composed with LS and Sync. are inserted between each domain to adjust voltage level and to avoid metastability of data.

EX1010, 2 (emphasis added).

79. Lee discloses a system where the performance of the total system is determined by a RISC (known to a POSITA as a type of microprocessor). That performance 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.




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

EX1010, Fig. 3.

80. Lee also discloses a power management unit (PMU) for each domain that generates that domain's clocks and voltages:

the PMU generates the clock and the supply voltage to adjust the speed of target domain

EX1010, 2.

81. The clocks in the PMU are generated by the PLL shown in Figure 4:

The DVFS scheme in Fig. 4 uses the **phase locked loop (PLL)** structure to find supply voltage level and clock frequency level on the MEP. It merges PLL and linear regulator (LR) as shown in Fig. 5

EX1010, 3 (emphasis added), Fig. 4.




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

EX1010, Fig. 4 (modified with a signal label).

82. Lee also discloses a chip layout for the domains. The chip layouts are arranged with domains in rows (stripes):




Fig. 8 Chip micro-photograph

EX1010, Fig. 8.

83. As shown in the logical and physical illustrations in Lee’s Figs. 2 and 8, the domains are independent. Each of Lee’s PLLs receives both a reference clock and a clock signal derived from a division of the clock for that domain. The divider output signal is labeled as “Div_Clk” in modified Figure 4. The PLL compares Div_Clk with the Reference clock, and adjusts the voltage-controlled oscillator until Div_Clk oscillates at the same frequency and phase as the Reference Clock. The division is based on a divisor value set by the Power Management Unit based on the “Performance Index”. Each domain has an independent PMU and Performance Index with a PLL that receives a clock signal from a divider unique to that domain.

84. Each domain has its own reference clock input, meaning that this design also discloses separate independent reference clocks to each of the domains. The independence is also indicated because the PMUs are at different corners of the die and a POSITA would understand that the reference clocks may be supplied from independent off-chip sources.

4. *Motivation to Combine*

85. A POSITA would be motivated to apply Talwar's DVFS "clustering" of processor cores to White's DVFS implementation in order to achieve the goal common to White and Talwar—DVFS in a multicore processor internal to a chip—while meeting practical chip space constraints. Specifically, a POSITA would be motivated to combine White with a reference that clearly discloses the use of more than two cores, like Talwar, because (1) White contemplates the use of "any number of cores," (*see* EX1006, 6:50-60), and (2) when designing a chip with "any number of cores," such as the quad-core processors already commercially available before 2010, a POSITA would be motivated to add Talwar's DVFS "clusters," or groupings of processor cores to White's DVFS implementation to save valuable chip space and achieve improved power performance. *See also* EX1023, 2.

86. A POSITA would further be motivated to combine White/Talwar with Lee because it provides further DVFS implementation details. Lee presents

clocking, synchronization, level shifting and chip layout, all of which are required for the detailed design of a multi-domain DVFS system.

87. First, a POSITA would be motivated to combine White and Talwar because they address similar design considerations and Talwar provides additional disclosure on the optimization of additional cores. White teaches that “one or more of the cores of a processor may be adjusted to operate at a lower frequency and/or voltage in order to save power in certain situations.” EX1006, 2:26-28. Similarly, Talwar teaches that “[d]ynamic voltage frequency scaling (DVFS)” can be used to reduce CPU power consumption in multi-core processors. EX1007, [0002]. White recognized that cores in prior art multi-core processors generally “operate at a common voltage level.” EX1006, 1:14-26. White proposed a multi-core processor in which “each logic core . . . operate[s] at a particular voltage and frequency.” *Id.*, 2:19-22. And Talwar recognized that “multiple clusters . . . for performing power management . . . is a stark contrast to conventional multi-core chips and conventional DVFS.” EX1007, [0012]. In other words, while White disclosed an example of individually controlling each core, Talwar explicitly recognized benefits—in terms of both efficiency and chip area—of controlling cores in clusters. Additionally, Talwar explicitly discloses an embodiment that includes up to 12 cores organized into independently controllable clusters. *See* EX1007,

FIG. 1. For these reasons, a POSITA would have naturally looked to Talwar to supplement White in a multi-core processor with more than two cores.

88. Second, a POSITA would have understood that White’s in-depth exemplar design—if applied with a distinct voltage regulator for each individual core—would need to be adapted as the number of cores increases due to chip space constraints associated with adding an additional voltage source for each additional core. By 2010, it was well known in the field of multi-core processor design that it was increasingly important to minimize overhead and to devote as much of the chip as possible to additional cores. *See supra* Section VII.A. For example, in 2008, with respect to off-chip voltage regulators like the one disclosed in White, it was known that “due to cost and system board area constraints, it may not be practical to implement multiple power domains using off-chip voltage regulators” because the switching voltage regulators “better suited for loads employing DVFS” are “bulky and consume large area.” EX1029, 3-4. Further, independently-regulated voltages for each core would require separate “bulky” regulator circuits for each independently-regulated core. *See id.*, 3-4. A POSITA therefore would have been motivated to find a solution that would allow for more cores without requiring a separate voltage regulator for each core. *See also* EX1019, 52-53 (surveying 2000s-era research into varying granularities of DVFS from “per-chip” to “per-core”).

89. Additionally, by 2010, processors with sixteen or more cores were being contemplated, studied, and designed. *See e.g.*, EX1020, 1:16-21, 1:52-53, EX1021, 2:54-55, FIG. 1. Thus, a POSITA would have understood White’s disclosure of a processor with “any number of cores” to realistically apply to multi-core processors with as many as 16 cores, each with a separate voltage and clock source. A POSITA extending the White design to more cores would recognize the benefits of the Talwar approach to reducing the overhead of voltage and frequency sources.

90. Talwar recognized that multi-core processors often “don’t have the same number of individual voltage sources” as processors. EX1007, [0004]. Working within this extra design consideration (i.e., fewer voltage sources than processors), Talwar’s clusters allow for cluster-by-cluster DVFS, while still enabling “maximum power savings” by ensuring all cores associated with a particular voltage source “run at the same frequency.” *Id.*, [0021]. Talwar’s “clustering” also fits White’s explicitly teaching that its design is not limited to per-core voltage and clock signals, but may be applied more generally “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, 3:37-52, 6:63-67.

91. Talwar's clusters, each set to its own voltage and frequency, correspond to Lee's DVFS domains. The specifics on how to generate the clock signals for each cluster/domain are supplied by Lee.

92. Accordingly, for the reasons described above, a POSITA designing a multi-core processor for a power-sensitive and area-sensitive system based on White would look to Talwar's method of clustering cores to supplement White's design at higher number of cores per chip. The POSITA would also look to Lee for details on implementing the clocks with PLLs from independent clock inputs per cluster/domain, for communicating between the cluster with synchronizers and level shifters, and for laying out the clusters in rows on an integrated circuit. *See supra* Section VII.A.

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

93. As discussed above, a POSITA would have been motivated to add Talwar's clustering of logic cores to White's dynamic voltage and frequency power management system. The combination would implement a multi-core processor following the block diagram of White, but would replace the individual logic cores with groups, or clusters of cores. Thus, the system would perform the same power management, or DVFS methods as White, but would do so on a cluster-by-cluster basis. White's Figure 2 embodiment, including logic cores 120A-120B can be modified to include four logic cores 120A-120D. Then,

applying Talwar's clustering disclosure, [green] logic cores 120A and 120C constitute a first cluster and [red] logic cores 120B and 120D constitute a second cluster as shown below:




FIG. 2A

See EX1006, FIG. 2A (modified to show Talwar's multiple cores per cluster from EX1007, FIG. 1.) In the combination, voltage regulator 130 provides one voltage level to cores 120A and 120C (added) and an independent voltage level to cores 120B and 120D (added). Frequency regulator 140 provides one external clock signal to cores 120A and 120C and an independent external clock signal to cores 120B and 120D. The details on how to generate a clock for each cluster are supplied by Lee's PLL design.

94. Because each core in a cluster operates at the same voltage, there is a corresponding ideal frequency for the entire cluster due to the inherent relationship between a CPU's supply voltage and achievable clock frequency. *See* EX1007, [0003], [0019], [0021]; *supra* Section VII.C. Talwar teaches that “reducing the frequency of one core in a cluster provides insignificant power savings unless all the cores in the cluster have their frequency reduced. EX1007, [0012]. Therefore, a POSITA would have understood that the internal clock signal for each core, as disclosed by White, should be shared among the cluster of cores operating at the same voltage level to avoid the need for additional, redundant clock circuitry. *See also* EX1006, 9:62-65; EX1007, [0012]. Thus, in the combination depicted above, frequency regulator 140 provides one clock signal to a local PLL shared by cores 120A and 120C and an independent clock signal to a separate, local PLL shared by cores 120B and 120D. *See* EX1006, 9:62-65. Each PLL generates the internal clock signal, as taught by White and Lee, for a cluster of cores that all operate at the same frequency, as taught by Talwar. *See* EX1006, 9:62-65; EX1007, [0021]. In short, a POSITA would have understood that each cluster has one, shared PLL to derive the shared clock for the cores in the cluster.

95. In the combination, White's common bridge logic need only be expanded to add a connection to each additional logic core. White discloses both that the common bridge logic “may include multiple ports or interfaces, each of

which may communicate with a different one of multiple logic cores,” and that there may be “any number of logic cores” in the multi-core processor. EX1006, 4:11-30, 6:56-60. Further, Talwar’s clustering does not increase the complexity of the common bridge logic by requiring it to interface with cores operating at additional voltages and frequencies because the cores within each cluster share a common voltage and frequency (i.e., the combination has two voltage and clock domains, the same number as in White’s disclosed embodiments).

96. The assignment of frequencies and voltages to clusters/domains in the White/Talwar/Lee combination can be done by Lee’s scheme where one domain is set by a microprocessor and the others automatically adapt based on data rates. Alternatively, the performance assignments of all clusters could be determined by Talwar’s power management policies or White’s temperature limits. Any of these methods meets the claims of the ’339 Patent because the assignment of frequencies to clusters/domains is not specified in the claims. A POSITA would be capable of determining which method to use based on the requirements of the system being designed.

97. White’s power management logic remains the same in the combination. White provides multiple teachings that its design is compatible with power management across groups of processor components, such as clusters of cores. *See e.g.*, EX1006, 3:37-52 (when adjusting the voltage and frequency of one

logic core, “the operating voltage and/or frequency for others of the logic cores may also be reduced, either at the same time or subsequently”), 6:63-67 (“an integrated circuit may include two or more logic cores that each may be configured to operate at frequencies and voltages independent of the frequencies or voltages that other logic cores . . . may be operating”), 8:44-49 (“the systems, methods, and techniques described herein for independently controlling operating voltages and/or frequencies may 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”).

98. A POSITA would have been motivated to combine the White/Talwar combination with Lee which shows the detailed logic design for implementing the clock generation and synchronizers required for a multi-cluster (multi-domain) DVFS system. Lee shows a clocking scheme for implementing DVFS on a chip with three power domains. Each domain has an independent input reference clock input that feeds a PLL in each domain in order to generate clocks for the computing elements in that domain.

6. *Element-by-Element Analysis*

99. The analysis of each element below incorporates the above descriptions of each reference, the motivation to combine the references, and the resulting modification.

a) Claim 1

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

100. In my opinion, White/Talwar/Lee renders obvious Limitation [1.pre].

As discussed in Section X.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.

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

101. In my opinion, White/Talwar/Lee renders obvious Limitation [1.1].

As discussed in Section X.A.4 and shown in modified Figure 2A, below, the proposed White/Talwar combination implements Talwar’s teaching that “[t]he cores and voltage sources are divided into clusters, whereby multiple cores in a cluster receive power from a single voltage source.” EX1007, [0011], claim 1; EX1006, claims 1, 25.




FIG. 2A

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

102. As shown in the White/Talwar version of Figure 2A, cores 120A and 120C of integrated circuit 100 (i.e., the multi-core processor) are grouped into a first cluster that is a first set of processor cores of the multi-core processor and cores 120B and 120D are grouped into a second cluster that is a second 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;

103. In my opinion, White/Talwar/Lee renders obvious Limitation [1.1.1].

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:

As illustrated in FIG. 2A, power management logic 200 may be coupled to, or otherwise interfaced to, voltage regulator 130 and frequency regulator 140 in order to adjust the operating voltages and frequencies for the logic cores 120 and/or for common bridge logic 110. Thus, when adjusting the operating characteristics for logic core 120A, **power management logic 200 may be configured to communicate, such as via one or more control signals, with voltage regulator 130 to instruct voltage regulator 130 to supply a particular voltage to logic core 120A.** 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.

104. White also teaches that the supply voltages and clock signal frequencies should be “adjusted” as needed, for example to support “power management” or “temperature management” requirements. EX1006, 2:15-44.

105. Lee shows separate PMUs (power management units) for three different domains. A POSITA would understand that those domains could include *first set of processor cores* such as White’s “any number of processor cores” or Talwar’s clusters of cores. Lee’s PMU is shown in Fig. 4, and each PMU has a PLL (*first phase lock loop (PLL)*) that receives a (*first clock signal*):

The DVFS scheme in Fig. 4 uses the phase locked loop (PLL) structure to find supply voltage level and clock frequency level on the MEP. It merges PLL and linear regulator (LR) as shown in Fig. 5.

EX1010, 3.

106. The combination of White and Lee is illustrated by the annotated FIG. 2A of White and Fig. 4 of Lee below. Labels have been added to show which signals in White correspond to those in Lee. The PMU in Lee's Fig. 4 has multiple instances (not shown) which correspond to the number of different frequency/voltage domains. Each instance is unique and includes a frequency generator and a voltage regulator as shown by the purple and blue shaded areas of Lee's Fig.4 diagram. The output of Lee's divider, labeled Div_Clk, sets the frequency and is the *clock signal input* to the PLL. Each instance of the PLL receives an independent Div_Clk (*clock signal input*) and produces an independent *output clock signal* for that domain.




FIG. 2A

[illegible]

Div Clk AC or Div Clk BD

EX1010, Fig.. 4 (modified with labels added).

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

107. In my opinion, White/Talwar/Lee renders obvious Limitation [1.2] based on the same evidence as presented in Limitation [1.1]. The White/Talwar combination implements Talwar's teaching that "[t]he cores and voltage sources are divided into clusters, whereby multiple cores in a cluster receive power from a single voltage source." EX1007, [0011], claim 1; EX1006, claims 1, 25. As shown in modified FIG. 2A, cores 120B and 120D of integrated circuit 100 (i.e., the multi-core processor) are grouped into a second cluster that is *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,

108. In my opinion, White/Talwar/Lee renders obvious Limitation [1.2.1].

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 based on the same evidence as presented in Limitation [1.1.1].

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

109. In my opinion, White/Talwar/Lee renders obvious Limitation [1.2.2].

As explained in Sections X.A.1 and X.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]. Lee discloses a system with three domains in which *the first supply voltage is independent from the second supply voltage*: “The voltage and frequency of each power domain are dynamically scaled according to the performance requirement of each domain.” EX1010, 1. Lee also indicates that “each domain has different clock and supply voltage.” EX1010, 2. Therefore, in the White/Talwar/Lee combination, the first supply voltage to the first cluster/domain is *independent from the second supply voltage* of the second cluster/domain.

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

110. In my opinion, White/Talwar/Lee renders obvious Limitation [1.2.3]. A POSITA would have understood that the Configurations disclosed by Lee disclose multiple, independent clock signals as inputs to the PLLs at each core. *See* Limitation [1.1.1]. The Div_Clk input to each PMU is a *clock signal* that is independent of other Div_Clk signals. 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.

111. Each PMU also receives a Reference Clock. The location of the PMUs in the Lee chip diagram suggest to a POSITA that the Reference Clocks come from an off-chip source. The reference clocks can also be considered *independent* input clock signals because they may be driven by different oscillators or clock dividers. *See* Section VII.B.




Fig. 2 Triple Domains Power Management




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

EX1010, Fig. 2 and Fig. 4 (emphasis and labels added).

112. A POSITA would have understood that Lee's three power management units (highlighted in yellow, blue and pink) are independent. Each of the three PMUs has one of the PLLs shown in Fig. 4. Hence, there are three

independent Div_Clks, one for each PMU. 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.

113. 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. A POSITA would understand these clocks to be independent, whether from separate oscillators, or from independent outputs from dividers of a single oscillator.

114. Thus, based both on White's explicit teachings that the frequency regulator "may include multiple clock sources in order to supply multiple, independent clock signals of different frequencies" and the disclosures of the independent reference clock in Lee, a POSITA would have understood that White/Talwar/Lee discloses *the first clock signal is independent from the second clock signal*. EX1006, 7:38-45; EX1010, Figs. 2-4.

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

115. In my opinion, White/Talwar/Lee renders obvious Limitation [1.3]. White teaches that the common bridge logic "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. White discloses *an interface block*, translation logic 300A, directly

coupled to core 120A. This interface block is also indirectly coupled to logic core 120B through communication protocol interfaces 320A and 320B and translation logic 300B in common bridge logic 110. The indirect coupling to 120B can also include the path to and from main memory 1104. Thus White's translation logic 300A is an *interface block coupled to 120A (the first set of processor cores) and also coupled to 120B (the second set of processor cores)*.




FIG. 3

EX1006, FIG. 3.

116. White's translation logic is performing the same voltage and frequency translations as described in the '339 Patent specification regarding its *interface block*¹:

117. Lee also discloses the Domain Converter which is an *interface block* for communicating between domains via the FIFO. 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*.




Fig. 2 Triple Domains Power Management

EX1010, FIG. 2.

¹ “FIG. 2 is a block diagram illustrating an example set of processor cores with example interface blocks having level shifters; FIG.3 is another block diagram illustrating an example set of processor cores with example interface blocks having synchronizers[.]” EX1001, 1:28-30.

[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.

118. In my opinion, White/Talwar/Lee renders obvious Limitation [1.3.1]. 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." As shown in modified FIG. 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 170." EX1006, 7:1-7. The communication is passed through the Bus Bridge as follows:

Bus Bridge 1102 may provide an interface between microprocessor 1100, main memory 1104, graphics controller 1108, and devices attached to PCI bus 1114. When an operation is received from one of the devices connected to Bus Bridge 1102, Bus Bridge 1102 identifies the target of the operation (e.g., a particular device or, in the case of PCI bus 1114, that the target is on PCI bus 1114). Bus Bridge 1102 routes the operation to the targeted device. Bus Bridge 1102 generally 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.

119. 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 (Section VII.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. Accordingly,

translation logic 300, via common bridge logic, provides a communication path between the cores via the memory subsystem. In other words, translation logic 300A is an *interface block, that is configured to facilitate communication between the first set of processor cores and the second set of processor cores.*

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.

120. In my opinion, White/Talwar/Lee renders obvious Limitation [2.1] for reasons discussed above in Limitation [1.3.1]. As discussed in Section VII.C, 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.

121. In addition to the disclosure described with respect to Limitation [1.3], White teaches that the common bridge logic “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. Because it is the processor cores that are operating at different logic levels and different frequencies, a POSITA would have understood this teaching to mean that the common bridge logic translates signals either directly between

processor cores, indirectly between processor cores via the use of a shared memory, or both, which also facilitates communication between the first and second set of processor cores.

122. White's translation logic may include level translation (*a first level shifter*) to allow the cores to communicate with each other through the common bridge logic which may operate at a different voltage and frequency:

For instance, the processor may include **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. Additionally, the processor may, in some embodiments, include circuitry configured to translate data communicated between the common bridge and a logic core operating at different frequencies.

EX1006, 4:36-43.

123. Lee discloses such a similar mechanism called "LS: Level Shifter" in Fig. 3a.




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

EX1010, Fig. 3a (annotated).

124. The *first logic level* and *second logic level* are shown as Level1 and Level2. 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. *See* EX1006, 8:13-40, 14:55-15:17, FIG. 3. As discussed, level shifters are an essential feature of chip designs that use multiple voltage domains. *See* Section VII.C. Further, level shifter designs were well known in the art. *See e.g.*, EX1037, FIGs. 4-5, 5:45-57; EX1034, 1-2, FIGs. 2(a), 2(b), 3. Additionally, a POSITA would have understood that the level shifter that translates between the logic level of the first set of processor cores to 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. *See also* EX1034, 1 (“each block on a chip could be operated under either high V_{DD} (V_{DDH}) or low V_{DD} (V_{DDL})”), Figure 2(a) (depicting a “conventional level shifter” referenced to both V_{DDH} and V_{DDL} to translate signals between the two logic levels).

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.

125. In my opinion, White/Talwar/Lee renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1]. A POSITA would understand that processor communications is bidirectional. When cores at different voltage levels communicate, level shifters are required both directions, hence the need for a *first* and *second level shifter*.

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.

126. The '339 Patent uses the term “periphery” only in claim 5 and in the specification where it only parrots the wording of the claims. There is no additional written description in the specification on what constitutes the “periphery.” In trying to understand claim 5, I can find two possible meanings.

127. First, “periphery” could be meant to describe “peripherals” which are controllers and devices connected to a processor through an I/O (input/output) bus. The most common IO bus in the 2010 timeframe was PCI which stands for Peripheral Component Interconnect. White’s FIG. 8 shows PCI devices 1112A and

1112B that are connected to the bus bridges. With this interpretation, White discloses controllers *located in a periphery of the multi-core processor* and a POSITA would understand that controllers inherently send *one or more control signals* that are received by the processor cores.


128. Alternatively, “periphery” could be referring to the physical location of the power control block 108 and clock control block 110 in ’339 Patent. EX1001, FIG. 1². FIG. 1 shows these control blocks on the left and right edge, which might be designated the *periphery*. With this interpretation, Lee discloses PMU1, PMU2 AND PMU3 that are shown in Figure 2 to be located on the outer edges (*periphery*) of the chip.

the PMU generates the clock and the supply voltage to adjust the speed of target domain


EX1010, 2.

² “the power control block 108 and the clock control block 110 may be arranged at two different sides of the multi-core processor 100 as shown in FIG. 1. In some other implementations, the power control block 108 and the clock control block 110 may be arranged at the same side of the multi-core processor 100. In yet some other implementations, the power control block 108 and the clock control block 110 may be arranged in a common area located near the center of the multi-core processor 100.” EX1001, 2:31-40.

129. The '339 Patent has no explanation of what constitutes a *control signal* from a *control block*, but a POSITA would understand that such control signals would serve to adjust the frequency and voltage of the cores. Lee's Figure 2 shows VDD and CLK signals communicated from each PMU to its corresponding domain. The VDD and CLK signals could be considered *control signals*, or alternatively a POSITA would understand that the Lee block diagram would not be expected to show every signal between the blocks. Clocks and voltages would typically be accompanied by other control signals, such as reset or error signals or signals to indicate to the cores that the voltage or clock signals are valid, hence the POSITA would understand Lee to teach control signals from the PMUs in the periphery to the cores. 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. See the annotated drawings below. Control blocks are highlighted green, control signals orange, and sets of processor cores blue.*



EX1010; EX1007 (combination of Lee's FIG. 2 and Talwar's FIG. 1 (annotated)).



EX1010, FIG. 8 (annotated).

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.

130. In my opinion, White/Talwar/Lee renders obvious Limitation [8.1].

The die photograph in Lee's Fig. 8 clearly shows that Domain1, 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*. As discussed in Section X.A.4, 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.

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.

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

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.

132. In my opinion, White/Talwar/Lee renders obvious Limitation [11.1] based on the evidence presented in claim 8. The die photograph in Lee's Fig. 8 clearly shows that the three domains (*regions*) are in separate rows.

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

133. In my opinion, White/Talwar/Lee renders obvious Limitation [14.1]. Limitation [5.1] includes all elements of Limitation [14.1] except that here, the control blocks are *located in a common region that is substantially central*. 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 choice³. *See also*

³ I also note that “substantially central” has no description in the specification of the '339 Patent and appears to lack any way to determine the bounds of this term.

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.). For example, the central placement of shared components, such as control logic, typically shortens the total length of wires between the cores and the control logic, shortening the delay between the cores and the control logic. On the other hand, a POSITA would also understand that it can be more efficient under some circumstances to place the control blocks, such as the power management controller, around the periphery of core clusters, rather than substantially central to the core clusters. That way the core clusters can be placed close to each other and may have shorter inter-cluster communication wires.

i) Claim 21

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

134. In my opinion, White/Talwar/Lee renders obvious Limitation [21.pre] for reasons discussed in Limitation [1.pre].

My analysis is based on the best I can do to understand what was meant by “substantially central.”

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

135. In my opinion, White/Talwar/Lee renders obvious Limitation [21.1] for reasons discussed in Limitation [1.1].

[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;

136. In my opinion, White/Talwar/Lee renders obvious Limitation [21.1.1] for reasons discussed in Limitation [1.1.1]. This claim is the same as 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*. As explained in Section X.A.4 and shown in modified Figure 2A, in White/Talwar each processor in the first / second cluster receives a first / second voltage signal (the first / second supply voltage) from voltage regulator 130. A POSITA would have understood White's "voltage regulator 130" and Lee's PMU to be *power control blocks*. A POSITA would have understood White's frequency regulator 140 and Lee's PMU to be *clock control blocks*.

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

137. In my opinion, White/Talwar/Lee renders obvious Limitation [21.2] for reasons discussed in Limitation [1.2].

[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,

138. In my opinion, White/Talwar/Lee renders obvious Limitation [21.2.1] for reasons discussed in Limitations [1.2.1] and [21.1.1]. The *power control block* is White's frequency regulator plus voltage regulator as implemented by Lee's PMUs. See the composite diagram in Limitation [1.1.1].

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

139. In my opinion, White/Talwar/Lee renders obvious Limitation [21.2.2] for reasons discussed in Limitation [1.2.2].

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

140. In my opinion, White/Talwar/Lee renders obvious Limitation [21.2.3] for reasons discussed in Limitation [1.2.3].

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

141. In my opinion, White/Talwar/Lee renders obvious Limitation [21.3] for reasons discussed in Limitation [1.3].

[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.


142. In my opinion, White/Talwar/Lee renders obvious Limitation [21.3.1] for reasons discussed in Limitation [1.3.1].

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

1. *Elgebaly*

143. Elgebaly was published on May 3, 2007, and is thus prior art under at least pre-AIA §102(b). EX1004. 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. While Elgebaly’s Figures 1 and 2 each show an ASIC (120) with two processor cores, Elgebaly expressly states that ASIC (120) may include “any number of processor cores” and “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” *Id.*, [0030], [0033]. Consistent with Elgebaly’s teaching and as described in Section VII.B, multi-core processors containing four cores were also well known in the art and available commercially well before 2010. *See e.g.*, EX1023, 2, EX1004, [0030], [0033].

144. Based on this disclosure, a POSITA would have understood Elgebaly to extend at least to an ASIC 120 with four cores, as shown in Modified Figure 1 below with an additional [blue] Processor Core 0 and an additional [red] Processor Core 3.



EX1004, FIG. 1 (modified to show additional processor cores).

145. Elgebaly explains the additional cores as follows:

[0030] In general, wireless device 100 may include fewer, more and/or different integrated circuits than those shown in FIG. 1. Furthermore, ASIC 120 may include fewer, more and/or different processing units than those shown in FIG. 1. In general, ASIC 120 may include any number of DSP cores and **any number of processor cores**. The number of processing units and the types of processing units included in ASIC 120 are typically dependent on various factors such as the communication systems, applications, and functions supported by wireless device 100.

EX1004, [0030] (emphasis added).

146. Elgebaly also describes the use of PLLs to generate clocks:

[0028] 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.** Internal memory 150 stores data and program codes used by the processing units within ASIC 120. External interface unit 160 interfaces with other units external to ASIC 120. Other circuits 170 may include a power control

unit that controls power to various processing units within ASIC 120, PLLs for receiver 114 and transmitter 116, and/or other circuitry.

EX1004, [0028] (emphasis added).

147. Elgebaly also discloses the grouping of multiple cores sharing common clocks and voltages:

[0033]For the embodiment shown in FIG. 2, 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. In other embodiments, **multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.**

EX1004, [0033] (emphasis added).

2. Motivation to Combine Elgebaly and Lee

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" (EX1004, [0035]), via Lee's "dynamical scaling the clock frequency and the supply voltage [to achieve] extremely efficient energy saving." EX1010, 1. Elgebaly also states that "Adaptive 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].

149. 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. 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. EX1010, Fig. 4.

150. Elgebaly shows multiple processor cores communicating to memory over a common bus (EX1004, FIG. 1.), but does not expressly show the circuitry to allow them to communicate when the core frequency or voltage does not match that of the target memory or core. A POSITA would turn to a reference such as Lee to show the circuits to perform the level translation and clock synchronization, namely Lee’s Domain Converter (EX1010, Figs. 2, 3a and 3b).

151. Lee discloses an *interface block* comprising a Domain Converter and FIFO for communicating between domains. 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*.




Fig. 2 Triple Domains Power Management

EX1010, Fig. 2.

152. Elgebaly presents techniques that can be applied to a number of different types of devices:

[0025] The adaptive voltage scaling techniques described herein may be used for various types of integrated circuits. For example, these techniques may be used for 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, and so on.

EX1004, [0025].

153. However, Elgebaly does not give guidance on how the power domains on ASICs or other devices should be arranged on the silicon chip. A POSITA would look to a reference such as Lee to see that one favorable way to arrange the power domains in stripes.

154. Elgebaly's Background explains that "[a]n effective method for reducing power consumption when operating in an active mode is to scale or adjust

the supply voltage.” *Id.*, [0006]. As was known in the art, this “may be achieved by identifying a critical signal path in an integrated circuit (IC), e.g., the signal path with the longest delay, and adjusting the supply voltage such that the critical signal path meets timing requirements.” *Id.*, [0007]. Elgebaly also recognized that “the critical signal path may vary from die to die based on IC process and temperature variations,” and also based on the current supply voltage and target clock frequency. *Id.*, [0007], [0009]. Elgebaly discloses the use of “adaptive voltage scaling (AVS)” units, where “one AVS unit 230 is provided for each processing core 130 and is used to adaptively scale the supply voltage for that core.” *Id.*, [0031].

155. Modified FIG. 2 of Elgebaly shows (green) AVS units to independently set the supply voltages to each (blue, red) processor core, and is modified to show multiple cores per processor. Elgebaly may be extended to “multiple processing cores” in which the clock signal and supply voltage are “jointly set for these cores.” *Id.* Further, “ASIC 120 may include . . . any number of processor cores.” *Id.*, [0030].




FIG. 2

EX1004, FIG. 2 (modified with multiple cores sharing clocks and voltage).

156. 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.” *Id.*, [0032]. The “target frequency may be selected based on the processing load and computational requirements for the core.” *Id.*

157. To generate clock signals, the clock generator 140 uses “one or more phase locked loops (PLLs) that control the oscillators” to generate corresponding clock signals. EX1004, [0028]. A POSITA would understand how to configure the PLLs to generate the clocks as shown in the modified FIG. 2 diagram below, and

would naturally turn to a reference such as Lee for the details on implementing the PLLs. EX1010, Fig. 4.




FIG. 2

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

3. Element-by-Element Analysis

a) Claim 1

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

158. In my opinion, Elgebaly/Lee renders obvious Limitation [1.pre]. As discussed in Section X.D.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. 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.” *Id.*, [0030]. Thus, Elgebaly’s ASIC is a *multi-core processor*.

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


159. In my opinion, Elgebaly/Lee renders obvious Limitation [1.1]. Elgebaly teaches that its multi-core processor “may include ...any number of processor cores.” *Id.*, [0030]. 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.

160. The Elgebaly/Lee Processor depicted in modified Figure 2 includes 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.


161. Thus, Elgebaly/Lee renders obvious a *first set of processor cores*.
[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;

162. In my opinion, Elgebaly/Lee renders obvious Limitation [1.1.1]. 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.” *Id.*, [0033]. The Elgebaly 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”).



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



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); *see id.*, FIG. 1, [0033].

163. The combination of Elgebaly and Lee is shown in the following diagram showing how the PLL of Lee, as shown in Lee's Fig. 4 can be combined with Elgebaly's FIG. 2:




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 (Lee modified with labels to show the combination with Elgebaly).

164. Consistent with Elgebaly's teaching and as described in Section VII.A and throughout the Declaration, multi-core processors containing four cores were well known in the art and available commercially well before 2010. *See e.g.*, EX1023, 2; EX1004, [0030], [0033]; EX1009, FIG. 1A. 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 FIG. 1.

165. The second modified Elgebaly FIG. 2 shows the clocks being driven by PLLs as taught by Elgebaly:

[0028]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).

166. The modified Figure 2 sharing of common clocks and voltages to multiple cores is also taught by Elgebaly:

[0033]For the embodiment shown in FIG. 2, 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. In other embodiments, **multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.**

Id., [0033] (emphasis added).

167. 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).

168. Accordingly, a POSITA would have understood that the processor cores could be grouped into two sets, such that each group of processor cores is managed by one of the AVS units explicitly disclosed in Elgebaly and the clocks for each set share a clock derived from the same clock generator explicitly disclosed in Elgebaly. *See* EX1004, FIG. 2; *see also* EX1007, [0011]-[0012] (disclosing that clusters of cores were known in the art). A POSITA would have added two or more additional processor cores to increase the processing power of

the ASIC without requiring additional AVS and clock generation circuitry, consistent with Elgebaly's teaching that cores may share a "common clock and/or a common supply voltage. See EX1004, [0033].

169. Thus, in 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").

170. Additionally, Elgebaly teaches that the supply voltages to the cores should be "adaptive," meaning that the processor "adjusts the supply voltage to a level that satisfies performance requirements," which may change over time, such as when the device enters an "idle mode." EX1004, [0034], [0006]. When the supply voltage is "jointly set for [each set of] cores," as described in Elgebaly and further in Limitation [1.1] above, the Elgebaly/Lee combination discloses the first set of processor cores being configured to *dynamically receive a first supply voltage*. See EX1004, [0033]. Thus, in the Elgebaly/Lee Processor, the cores in the first set of processor cores 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*.

171. Elgebaly's cores are also configured to *dynamically receive* respective clock signals. "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.*

172. 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*.

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

173. Elgebaly/Lee renders obvious Limitation [1.2] for the reasons discussed in Limitation [1.1].

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]. 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"). Thus, the Elgebaly/Lee combination renders obvious *a second set of processor cores*.

[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,

175. In my opinion, Elgebaly/Lee renders obvious Limitation [1.2.1] for reasons discussed in Limitation [1.1.1]. In the modified FIG. 2 above, the common cores Cores 2-3 *dynamically receive a second supply voltage (Vdd4)*. 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)*.

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

176. In my opinion, Elgebaly/Lee renders obvious Limitation [1.2.2] for reasons discussed in Limitation [1.1.1] and [1.2.1]. Specifically, Elgebaly teaches that “the supply voltage for each processing core 130 may be set independently by power management unit 180.” EX1004, [0033]. 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 voltage for each set of cores would be set *independently*.

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

177. In my opinion, Elgebaly/Lee renders obvious Limitation [1.2.3] for the reasons discussed in Limitations [1.1.1] and [1.2.1]. As explained in Limitation [1.1.1], Elgebaly teaches a clock generator that includes multiple PLLs, one for each independent clock signal. In the modified Elgebaly’s FIG. 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*). The sources of these signals are from two different processors, hence *the first clock signal is independent from the second clock signal*.

[0032] Each processing core 130 receives from clock generator 140 a clock used to trigger synchronous circuits within the core. Each core 130 **provides to clock generator 140 and to its AVS unit 230 a target frequency** for its clock. This target frequency may be selected based on the processing load and computational requirements for the core. 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.

EX1004, [0032] (emphasis added).

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

178. In my opinion, Elgebaly/Lee renders obvious Limitation [1.3].


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. *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. The translation and synchronization is disclosed by Lee's Domain Converter that includes a Level Shifter and Synchronizer as shown in EX1010, Figs. 2, 3a and 3b. 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.

[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.

179. In my opinion, Elgebaly/Lee renders obvious Limitation [1.3.1].

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 (Section VII.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. 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*).



EX1004 (modified FIG. 1 with Lee Domain converters).

180. The modified Elgebaly's FIG. 1 shows the path for Core 1, to write to Internal Memory which is then be read by Core 3. 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.

181. 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*.

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.

182. In my opinion, Elgebaly/Lee renders obvious Limitation [2.1] for reasons discussed above in Limitation [1.3.1]. As discussed in Section VII.C, it was well known in the art that multi-core processors implementing multiple voltage domains required mechanisms for translating logic levels between chip components operating at different voltages. A POSITA would have understood FIG. 1 of Elgebaly, which depicts communication between processor cores operating at independent voltage levels, to require a voltage level shifter as part of the bus shown in the center of the ASIC. Lee discloses its “LS: Level Shifter” in Fig. 3a.




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

EX1010, Fig. 3 (annotated).

183. In annotated Lee's Fig. 3, the *first logic level* and *second logic level* are shown as Level1 and Level2.

184. As discussed, level shifters are an essential feature of chip designs that use multiple voltage domains. *See* Section VII.C. Further, level shifter circuits were well known in the art. *See e.g.*, EX1037, Figs. 4-5, 5:45-57; EX1034, 1-2, Figures 2(a), 2(b), 3. Additionally, a POSITA would have understood that the level shifter that translates between the logic level of the first set of processor cores to 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. *See* EX1034, Figure 2(a). For these reasons, it would have been obvious to a POSITA, 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 of processor cores for cache coherency operations between the first and second set of processor cores (i.e., "a first signal traveling from the first set of processor cores to the second set of processor cores").

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.

185. In my opinion, Elgebaly/Lee renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1]. A POSITA would understand that communications in a multiprocessor system is bidirectional. When cores at different voltage levels communicate, level shifters are required both directions, hence the need for a *first* and *second level shifter*. In the modified Elgebaly's FIG. 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.

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.

186. Assuming that "periphery" is referring to the physical location of the power control block 108 and clock control block 110 in '339 Patent. EX1001,

FIG. 1⁴ shows these control blocks on the left and right edge, which might be designated the *periphery*. With this interpretation, Lee discloses PMU1, PMU2 AND PMU3 that are shown in Figure 3 to be located on the outer edges (*periphery*) of the chip.


the PMU generates the clock and the supply voltage to adjust the speed of target domain

EX1010, 2.

187. The '339 Patent has no explanation of what constitutes a *control signal* from a *control block*, but a POSITA would understand that such control signals would serve to adjust the frequency and voltage of the cores. Lee's Fig. 2 shows VDD and CLK signals communicated from each PMU to its corresponding domain. In the Elgebaly/Lee combination, each domain contains multiple cores. Hence, Elgebaly/Lee discloses sets of *processor cores [that] are configured to*

⁴ "the power control block 108 and the clock control block 110 may be arranged at two different sides of the multi-core processor 100 as shown in FIG. 1. In some other implementations, the power control block 108 and the clock control block 110 may be arranged at the same side of the multi-core processor 100. In yet some other implementations, the power control block 108 and the clock control block 110 may be arranged in a common area located near the center of the multi-core processor 100." EX1001, 2:31-40.

receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor. See the annotated drawings below. Control blocks are highlighted green, control signals orange, and sets of processor cores blue.



EX1010; EX1004 (combination of Lee's Fig. 2 and Elgebaly's modified FIG. 2 (annotated)).




Fig. 8 Chip micro-photograph

EX1010, FIG. 8 (annotated).

188. Elgebaly discloses “external interface 160” and “other circuits 170” onboard the ASIC, which “may include a power control unit that controls power to various processing units.” EX1004, [0028]. In the modified Elgebaly’s FIG. 2 below, the [orange] “other circuits 170,” including the power control unit, may be located in a periphery of the ASIC. *Id.*, FIG. 1. Hence, in the Elgebaly/Lee combination, Elgebaly’s “other logic” includes Lee’s PMUs which are located on the *periphery*.




FIG. 2

EX1004, FIG. 2 (modified).

189. Thus, a POSITA would have understood the Elgebaly/Lee combination to disclose—or at least 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.

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.

190. In my opinion, Elgebaly/Lee renders obvious Limitation [8.1]. The die photograph in Lee's Fig. 8 clearly shows that Domain1, Domain2 and Domain3 are located in separate *regions* of the chip. Lee in combination with Elgebaly

shows that the *first, second and third regions* could include a *set of processor cores*. As discussed in Section X.A.4, 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.

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.

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

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.

192. In my opinion, Elgebaly/Lee renders obvious Limitation [11.1] based on the evidence presented in claim 8. The die photograph in Lee's Fig. 8 clearly shows that the three domains (*regions*) are in separate rows.

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.

193. In my opinion, Elgebaly /Lee renders obvious Limitation [14.1]. Limitation [5.1] includes all elements of Limitation [14.1] except that here, the control blocks are *located in a common region that is substantially central*. 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 choice⁵. *See also* 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.). For example, the central placement of

⁵ I also note that “substantially central” has no description in the specification of the ’339 Patent and appears to lack any way to determine the bounds of this term. My analysis is based on the best I can do to understand what was meant by “substantially central.”

shared components, such as control logic, typically shortens the total length of wires between the cores and the control logic, shortening the delay between the cores and the control logic. On the other hand, a POSITA would also understand that it can be more efficient under some circumstances to place the control blocks, such as the power management controller, around the periphery of core clusters, rather than substantially central to the core clusters. That way the core clusters can be placed close to each other and may have shorter inter-cluster communication wires.

i) Claim 21

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

194. In my opinion, Elgebaly/Lee renders obvious Limitation [21.pre] for reasons discussed in Limitation [1.pre].

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

195. In my opinion, Elgebaly/Lee renders obvious Limitation [21.1] for reasons discussed in Limitation [1.1].

[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;

196. In my opinion, Elgebaly/Lee renders obvious Limitation [21.1.1] for reasons discussed in Limitation [1.1.1]. This claim is the same as 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*. 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. 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.

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

197. In my opinion, Elgebaly/ Lee renders obvious Limitation [21.2] for reasons discussed in Limitation [1.2].

[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,

198. In my opinion, Elgebaly/ Lee renders obvious Limitation [21.2.1] for reasons discussed in Limitation [1.2.1].

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

199. In my opinion, Elgebaly/ Lee renders obvious Limitation [21.2.2] for reasons discussed in Limitation [1.2.2].

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

200. In my opinion, Elgebaly/ Lee renders obvious Limitation [21.2.3] for reasons discussed in Limitation [1.2.3].

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

201. In my opinion, Elgebaly/ Lee renders obvious Limitation [21.3] for reasons discussed in Limitation [1.3].

[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.

202. In my opinion, Elgebaly/Lee renders obvious Limitation [21.3.1] for reasons discussed in Limitation [1.3.1].

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

1. Vajda

203. Vajda claims priority to a May 26, 2009 PCT filing, and is thus prior art under at least pre-AIA §102(e). EX1011. 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]. Vajda provides guidance for “the geometry” (i.e., layout of the individual cores) “of the multi-core chip 230.” *Id.*, [0063]. 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.” *Id.*

2. *Motivation to Combine*

204. A POSITA would be motivated to apply Vajda's chip layout as an alternative to the chip layout shown in the Lee die photo. 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 Vajda type of layout. Indeed, as discussed in Section VII.A, 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). In other cases, where overheating is not a problem, the designer may instead stay with the Lee stripe layout. While these two references might seem to be in conflict with each other, a POSITA would recognize that they are two viable approaches to layout and are applicable in different situations. 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.

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.

205. Elgebaly/Lee/Vajda renders obvious Limitation [9.1].

Elgebaly/Lee/Vajda includes processor cores arranged in a checkerboard pattern.

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. The first and second regions are an overlapping region because the checkerboard pattern of cores (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) results in overlapping regions containing each set of cores.

All statements made herein of my own knowledge are true, all statements made herein on information and belief are believed to be true, and these statements were made with the knowledge that willful false statement and the like are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

Dated: 9/11/25

Respectfully submitted,

A handwritten signature in blue ink, appearing to read "Robert W. Horst".

Robert Horst, Ph.D.