

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NXP USA, INC. and QUALCOMM INCORPORATED,

Petitioners,

v.

REDSTONE LOGICS LLC,

Patent Owner.

Case: IPR2025-00485
U.S. Patent No. 8,549,339

PETITION FOR *INTER PARTES* REVIEW UNDER 35 U.S.C. § 312

TABLE OF CONTENTS

TABLE OF CONTENTS.....	ii
TABLE OF EXHIBITS.....	v
CLAIM LISTING	viii
I. INTRODUCTION	1
II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8.....	2
A. REAL PARTY IN INTEREST (37 C.F.R. §42.8(B)(1)).....	2
B. RELATED MATTERS (37 C.F.R. §42.8(B)(2)).....	3
1. Litigation.....	3
2. Administrative Proceedings	3
C. DESIGNATION OF COUNSEL (37 C.F.R. §42.8(B)(3)).....	4
D. SERVICE INFORMATION (37 C.F.R. §42.8(B)(4)).....	4
III. GROUNDS FOR STANDING.....	5
IV. THE GROUNDS IN THIS PETITION ARE NOT CUMULATIVE.....	5
A. ORIGINAL PROSECUTION	5
B. IPR2025-00085	6
V. STATEMENT OF PRECISE RELIEF REQUESTED FOR EACH CLAIM CHALLENGED	7
A. CLAIMS FOR WHICH REVIEW IS REQUESTED (37 C.F.R. § 42.104(B)(1)).....	7
B. STATUTORY GROUNDS OF CHALLENGE (§42.104(B)(2)).....	7
VI. LEVEL OF ORDINARY SKILL IN THE ART	9
VII. FIELD OF TECHNOLOGY	9
A. DYNAMIC VOLTAGE AND FREQUENCY SCALING (DVFS)	9
B. MULTI-CORE PROCESSORS AND DVFS	10
C. CLOCK GENERATION.....	11
D. PROCESSOR CACHE MEMORY AND CACHE COHERENCY.....	13

VIII. THE '339 PATENT	15
A. OVERVIEW	15
B. PROSECUTION HISTORY	17
IX. CLAIM CONSTRUCTION	18
A. “THE FIRST CLOCK SIGNAL IS INDEPENDENT FROM THE SECOND CLOCK SIGNAL”	19
B. INDEFINITE TERMS	19
X. REASONS FOR THE RELIEF REQUESTED (§§42.22(a)(2), 42.104(b)(4)) .	21
A. GROUND 1: WHITE AND TALWAR RENDER OBVIOUS CLAIMS 1-3, 5, 8, 10-11, 14, 21	21
1. <i>White</i>	21
2. <i>Talwar</i>	24
3. <i>Motivation to Combine</i>	25
4. <i>Modification of White in View of Talwar</i>	26
5. <i>Element-by-Element Analysis</i>	28
B. GROUND 2: WHITE, TALWAR, AND CHO RENDER OBVIOUS CLAIMS 5, 14	42
1. <i>Cho</i>	42
2. <i>Motivation to Combine White/Talwar with Cho</i>	43
3. <i>Modification of the White/Talwar Combination in View of Cho</i>	44
4. <i>Element-by-Element Analysis</i>	46
C. GROUND 3: WHITE, TALWAR, AND VAJDA RENDER OBVIOUS CLAIMS 8-9.....	47
1. <i>Vajda</i>	47
2. <i>Motivation to Combine White/Talwar with Vajda</i>	48
3. <i>Modification of the White/Talwar Combination in View of Kato</i>	48
4. <i>Element-by-Element Analysis</i>	49

D.	GROUND 4: ELGEBALY AND OZER RENDER OBVIOUS CLAIMS 1-3, 5, 8, 10-11, 14, AND 21	50
1.	<i>Elgebaly</i>	50
2.	<i>Ozer</i>	57
3.	<i>Motivation to Combine Elgebaly and Ozer</i>	59
4.	<i>Modification of Elgebaly in View of Ozer</i>	61
5.	<i>Element-by-Element Analysis</i>	63
E.	GROUND 5: ELGEBALY, OZER, AND VAJDA RENDER OBVIOUS CLAIMS 8-9	76
1.	<i>Motivation to Combine Vajda with Elgebaly and Ozer</i>	76
2.	<i>Modification of Elgebaly/Ozer in View of Vajda</i>	77
3.	<i>Element-by-Element Analysis</i>	78
F.	FOUNDATIONS 6-7: FOUNDATIONS 4-5 AND WILLIAMS RENDER OBVIOUS CLAIMS 1-3, 5, 8-11, 14, AND 21	78
XI.	CONCLUSION	80
	CERTIFICATE OF SERVICE	82

TABLE OF EXHIBITS

Exhibit	Short Name	Description
EX1001	'339 Patent	U.S. Patent No. 8,549,339
EX1002	Annavaram Declaration	Declaration of Dr. Murali Annavaram
EX1003	Annavaram CV	<i>Curriculum Vitae</i> of Dr. Murali Annavaram
EX1004	Elgebaly	U.S. Patent App. Pub. No. 2007/0096775
EX1005	'339 Prosecution History	File History of U.S. Patent No. 8,549,339
EX1006	White	U.S. Patent No. 7,263,457
EX1007	Talwar	U.S. Patent App. Pub. No. 2009/0271646
EX1008	Williams	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.
EX1009	RESERVED	RESERVED
EX1010	Cho	U.S. Patent App. Pub. No. 2010/0073068
EX1011	Vajda	U.S. Patent App. Pub. No. 2012/0060170
EX1012	RESERVED	RESERVED
EX1013	Ozer	U.S. Patent App. Pub. No. 2008/0209133
EX1014	NXP Complaint	Doc. No. 1, Complaint, <i>Redstone Logics LLC v. NXP Semiconductors N.V., et al.</i> , No. 7:24-cv-00028-DC-DTG (W.D. Tex.)
EX1015	Qualcomm Complaint	Doc. No. 1, Complaint, <i>Redstone Logics LLC v. Qualcomm Inc., et al.</i> , No. 7:24-cv-00231-ADA (W.D. Tex.)

Exhibit	Short Name	Description
EX1016	MediaTek Complaint	Doc. No. 1, Complaint, <i>Redstone Logics LLC v. MediaTek, Inc., et al.</i> , No. 7:24-cv-00029-DC-DTG (W.D. Tex.)
EX1017	NXP Opening Markman Brief	Doc. No. 39, Defendant NXP USA Inc.'s Opening Claim Construction Brief, <i>Redstone Logics LLC v. NXP USA, Inc.</i> , No. 7:24-cv-00028-DC-DTG (W.D. Tex.)
EX1018	Redstone Responsive Markman Brief	Doc. No. 41, Plaintiff's Responsive Claim Construction Brief, <i>Redstone Logics LLC v. NXP USA, Inc.</i> , No. 7:24-cv-00028-DC-DTG (W.D. Tex.)
EX1019	Kaxiras Textbook	S. Kaxiras, Computer Architecture Techniques for Power-Efficiency, 2008.
EX1020	Fryman	U.S. Patent No. 8,312,225
EX1021	Capps	U.S. Patent No. 7,584,369
EX1022	Pelley	U.S. Patent No. 7,777,330
EX1023	Intel Multi-Core Processor White Paper	R. Ramanathan, Intel multi-core processors: Making the move to quad-core and beyond, Technology@Intel Magazine, December 2006.
EX1024	Panov Paper	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.
EX1025	Flautner	U.S. Patent App. Pub. No. 2005/0034002
EX1026	Liptay 1968	J.S. Liptay, Structural aspects of the System/360 Model 85, II. The cache, IBM Systems Journal, Vol. 7, No. 1, 1968.
EX1027	Gold	U.S. Patent App. Pub. No. 2009/0063888
EX1028	Cesky	U.S. Patent No. 7,538,625
EX1029	Kim Paper	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.

Exhibit	Short Name	Description
EX1030	Stenstrom Paper	P. Senstrom, A Survey of Cache Coherence Schemes for Multiprocessors, Computer, Volume 23, Issue 6, June 1990.
EX1031	Campbell	U.S. Patent No. 10,482,952
EX1032	Crandall	U.S. Patent No. 4,873,703
EX1033	Kato	U.S. Patent App. Pub. No. 2004/0210738
EX1034	Redstone / NXP Infringement Chart	Doc. No. 1-2, Exhibit 2 to Redstone’s Complaint, <i>Redstone Logics LLC v. NXP Semiconductors N.V., et al.</i> , No. 7:24-cv-00028-DC-DTG (W.D. Tex.)
EX1035	Redstone / Qualcomm Infringement Chart	Doc. No. 14-2, Exhibit 2 to Redstone’s Second Amended Complaint, <i>Redstone Logics LLC v. Qualcomm Inc., et al.</i> , No. 7:24-cv-00231-ADA (W.D. Tex.)
EX1036	Ishimi	U.S. Patent No. 8,200,878
EX1037	Vajda PCT	WO 2010/138031
EX1038	Williams Paper IEEE Webpage	IEEE Webpage for Williams et al., “Low Latency Clock Domain Transfer for Simultaneously Mesochronous, Plesiochronous and Heterochronous Interfaces,” 13 th IEEE International Symposium on Asynchronous Circuits and Systems, 2007. https://ieeexploreonorthbridge.ieee.org/document/4137045 (last accessed on January 17, 2025).
EX1039	Tran Level Shifter Paper	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.
EX1040	Williams Paper PDF Metadata	Screenshot of EX1008’s pdf metadata, generated by Dr. Annavaram on January 20, 2025.

CLAIM LISTING

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

Claim	Limitation #	Limitation
5	[5.1]	The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.
8	[8.1]	The multi-core processor of claim 1, wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor.
9	[9.1]	The multi-core processor of claim 8, wherein the first region and the second region are overlapping regions of the multi-core processor.
10	[10.1]	The multi-core processor of claim 8, wherein the first region and the second region are non-overlapping regions of the multi-core processor.
11	[11.1]	The multi-core processor of claim 8, wherein the first region corresponds to a first row of the multi-core processor, and wherein the second region corresponds to a second row of the multi-core processor.
14	[14.1]	The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores.
21	[21.pre]	A multi-core processor, comprising:
	[21.1]	a first set of processor cores of the multi-core processor,
	[21.1.1]	wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block;
	[21.2]	a second set of processor cores of the multi-core processor,
	[21.2.1]	wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage from the power control block and a second output clock signal from a second PLL having a second clock signal as input in the clock control block,

Claim	Limitation #	Limitation
	[21.2.2]	wherein the first supply voltage is independent from the second supply voltage, and
	[21.2.3]	the first clock signal is independent from the second clock signal; and
	[21.3]	an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,
	[21.3.1]	wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.

I. INTRODUCTION

The Challenged Claims¹ relate to a “multi-core processor” having a “first” and “second” “set of processor cores,” each “configured to dynamically receive a [first] / [second] supply voltage and a [first] / [second] output clock signal of a first phase lock loop (PLL) having a [first] / [second] clock signal as input.” EX1001, Claim 1. They further recite that “the first supply voltage is independent from the second supply voltage, and the first clock signal is independent from the second clock signal.” *Id.* Both multi-core processors and applying Dynamic Voltage and Frequency Scaling (DVFS) to sets of cores in a multi-core processor were well known prior to the Challenged Claims’ earliest priority date.

White discloses a multi-core processor where each core “may be configured to operate at frequencies and voltages *independently of* the other logic cores.” EX1006, 2:15-44. Talwar discloses “divid[ing]” the cores of a multi-core processor “into clusters” where each “cluster receive[s] power from a single voltage source” and “[e]ach cluster may have a different cluster frequency.” EX1007, [0011], [0018]-[0019]. Applying Talwar’s teaching to White therefore renders obvious the challenged independent claims as shown below.

¹ Claims 1-3, 5, 8-11, 14, and 21.

Likewise, Elgebaly discloses “the clock for each processing core 130 may be set *independently*” “and the supply voltage for each processing core 130 may be set *independently*.” EX1004, [0033]. Elgebaly further discloses that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” *Id.* Accordingly, Elgebaly (along with known memory sharing techniques disclosed by Ozer) renders obvious the challenged independent claims as shown below.




FIG. 2A

EX1006, FIG. 2A (modified)




FIG. 2
(MODIFIED)

EX1004, FIG. 2 (modified)

II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8

A. Real Party in Interest (37 C.F.R. §42.8(b)(1))

The real parties-in-interest are Petitioners NXP USA, Inc. and Qualcomm Incorporated, as well as Qualcomm Technologies, Inc. NXP USA, Inc. is directly or indirectly owned by Freescale Semiconductor Holdings V, Inc., NXP B.V., and NXP Semiconductors N.V. (NASDAQ: NXPI). NXP Semiconductors N.V. and

NXP B.V. were defendants in *Redstone Logics LLC v. NXP USA, Inc.*, 7-24-cv-00028 (W.D. Tex 2024), but were dismissed.

B. Related Matters (37 C.F.R. §42.8(b)(2))

1. Litigation

Petitioner is aware of the following judicial matters involving the '339 Patent:

Case Heading	Number	Court	Filed
<i>Redstone Logics LLC v. Samsung Electronics Co., Ltd. et al</i>	2-23-cv-00485	E.D. Tex.	October 17, 2023
<i>Redstone Logics LLC v. NXP USA, Inc.</i> (“the NXP Litigation”)	7-24-cv-00028	W.D. Tex.	January 26, 2024
<i>Redstone Logics LLC v. MediaTek Inc. et al</i>	7-24-cv-00029	W.D. Tex.	January 26, 2024
<i>Redstone Logics LLC v. Qualcomm Inc. et al</i> (“the Qualcomm Litigation”)	7-24-cv-00231	W.D. Tex.	September 20, 2024

The NXP Litigation and Qualcomm Litigation are collectively referred to as “The Litigation.”

2. Administrative Proceedings

The '339 Patent is subject to a pending IPR Petition IPR2025-00085 filed by MediaTek, Inc. and MediaTek USA, Inc. (collectively, the “IPR2025-00085 Petitioners”). Petitioners are neither a privy nor a real party in interest (RPI) of MediaTek, Inc. or MediaTek USA, Inc.

C. Designation of Counsel (37 C.F.R. §42.8(b)(3))

Petitioners provide the following designation of counsel.

Lead Counsel	Backup Counsel
Eagle H. Robinson Reg. No. 61,361	Daniel S. Leventhal Reg. No. 59,576 Eric C. Green Reg. No. 59,817 Richard S. Zembek Reg. No. 43,306 G. Duncan Waldrop Reg. No. 83,028

D. Service Information (37 C.F.R. §42.8(b)(4))

Lead Counsel	Backup Counsel
Eagle H. Robinson Norton Rose Fulbright US LLP 98 San Jacinto Blvd., Ste. 1100 Austin, TX 78746 512.536.3083 (telephone) 512.536.4598 (facsimile) eagle.robinson@nortonrosefulbright.com	Daniel S. Leventhal Norton Rose Fulbright US LLP 1550 Lamar Street, Ste. 2000 Houston, TX 77010-4106 713.651.8360 (telephone) 713.651.5246 (facsimile) daniel.leventhal@nortonrosefulbright.com Eric C. Green Norton Rose Fulbright US LLP 98 San Jacinto Blvd., Ste. 1100 Austin, TX 78746 512.536.3009 (telephone) eric.green@nortonrosefulbright.com

Lead Counsel	Backup Counsel
	Richard S. Zembek Norton Rose Fulbright US LLP 1550 Lamar Street, Ste. 2000 Houston, TX 77010-4106 713.651.5283 (telephone) richard.zembek@nortonrosefulbright.com G. Duncan Waldrop Norton Rose Fulbright US LLP 98 San Jacinto Blvd., Ste. 1100 Austin, TX 78746 512.536.3131 (telephone) duncan.waldrop@nortonrosefulbright.com

Petitioner consents to electronic service.

III. GROUNDS FOR STANDING

Petitioner certifies that the '339 Patent is available for IPR, and that Petitioner is not barred or estopped from requesting IPR.

IV. THE GROUNDS IN THIS PETITION ARE NOT CUMULATIVE

A. Original Prosecution

Other than White, the relied-upon prior art was not before the Examiner. EX1001. White was cited in an IDS, but was not addressed by the Examiner. EX1005, 112/194, 102/194; *see also Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte Ges.m.b.H et al.*, No. IPR2019-01469, Paper 6 at 10-11 (P.T.A.B. Feb. 13, 2020). Further, the Grounds herein are not cumulative of art considered by the Examiner. For example, the applicant distinguished Jacobowitz

and Kim during prosecution by arguing that neither disclosed independent supply voltages for each set of cores. EX1005, 82/194. Here, however, each of the Grounds includes independent supply voltages for each set of cores.

B. IPR2025-00085

Discretionary denial, under either *General Plastics* or *Advanced Bionics* is not appropriate. First, Petitioners are not related, are not co-defendants or RPIs with, and are not accused of infringing based on the same products as the petitioners in IPR2025-00085. *Compare* EX1014, 1/46 *and* EX1015, 1/65 *with* EX1016, 1/6. “[W]here, as here, the first and second petitioners are neither the same party, nor possess a significant relationship under *Valve, General Plastic* factor one necessarily outweighs the other *General Plastic* factors.” *Videndum Prod. Solutions, Inc. v. Rotolight Ltd.*, IPR2023-01218, Paper 12 (Director Review) at 5-6. Second, addressing this Petition would not be an inefficient use of resources or interfere with the Board’s decision in IPR2025-00085. This Petition relies on different art than that in IPR2025-00085. Therefore, *General Plastics* factors 2, 6, and 7 also weigh against discretionary denial. Finally, Petitioners are timely filing this Petition; PO sued co-Petitioner NXP in January of 2024 and sued co-Petitioner Qualcomm in September of 2024. *See* EX1014, 6/46; EX1015, 6/65.

V. STATEMENT OF PRECISE RELIEF REQUESTED FOR EACH CLAIM CHALLENGED

A. Claims for Which Review Is Requested (37 C.F.R. § 42.104(b)(1))

Petitioner requests cancellation of Claims 1-3, 5, 8-11, 14, and 21.

B. Statutory Grounds of Challenge (§42.104(b)(2))²

Ground 1: Claims 1-3, 5, 8, 10-11, 14, 21 are unpatentable under §103 over US 7,263,457 (“White”) in view of US 2009/0271646 (“Talwar”). White published August 28, 2007, and is thus prior art under at least pre-AIA §102(b). EX1006. 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.

Ground 2: Claims 5, 14 are unpatentable under §103 over White and Talwar, further in view of US 2010/0073068 (“Cho”). Cho was filed September 22, 2008, and is thus prior art under at least pre-AIA §102(e). EX1010.

Ground 3: Claims 8-9 are unpatentable under §103 over White and Talwar, further in view of US 2012/0060170 (“Vajda”). Vajda claims priority to a May 26, 2009 PCT filing, and is thus prior art under at least pre-AIA §102(e).³ EX1011.

² All Grounds are supported by a POSITA’s general knowledge. *Koninklijke Philips N.V. v. Google LLC*, 948 F.3d 1330, 1337-38 (Fed. Cir. 2020).

³ The same content is also included in WO 2010/138031. EX1037.

Ground 4: Claims 1-3, 5, 8, 10-11, 14, 21 are unpatentable under §103 over US 2007/0096775 (“Elgebaly”) in view of US 2008/0209133 (“Ozer”). Elgebaly was published May 3, 2007, and is thus prior art under at least pre-AIA §102(b). EX1004. Ozer published August 28, 2008, and is thus prior art under at least pre-AIA §102(b). EX1013.

Ground 5: Claims 8-9 are unpatentable under §103 over Elgebaly and Ozer, further in view of Vajda.

Ground 6: Claims 1-3, 5, 8, 10-11, 14, and 21 are unpatentable under §103 over Elgebaly and Ozer, further in view of *Low Latency Clock Domain Transfer for Simultaneously Mesochronous, Plesiochronous and Heterchronous Interfaces*, Wade L. Williams et al., 13th IEEE International Symposium on Asynchronous Circuits and Systems, March 12-14, 2007 (“Williams”). Williams was presented publicly at the 13th IEEE International Symposium on Asynchronous Circuits and Systems, which took place from March 12-14, 2007, and made publicly available in print form by IEEE on or shortly after March 26, 2007. EX1008; EX1002, ¶¶55-58, *see also* EX1038 1/2, EX1040, 1/1. Williams is thus prior art under at least pre-AIA §§102(a), (b).

Ground 7: Claims 8-9 are unpatentable under §103 over Elgebaly, Ozer, and Vajda, further in view of Williams.

VI. LEVEL OF ORDINARY SKILL IN THE ART

A person of ordinary skill in the art (“POSITA”) would have had at least (1) a bachelor’s degree in electrical engineering, computer engineering, computer science, or a similar field, and (2) at least two years of academic or industry experience designing or analyzing electronic circuits, semiconductors, processors, or power management, and related firmware and software, or the equivalent. EX1002, ¶31. The prior art evidences this level of ordinary skill. *Id.* Additional education might substitute for some experience, and vice versa. *Id.*

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, efficient memory hierarchy designs such as caches and cache coherency protocols, and common semiconductor die layouts. *Id.*, ¶32.

VII. FIELD OF TECHNOLOGY

A. Dynamic Voltage and Frequency Scaling (DVFS)

The amount of work a processor core can accomplish in a given period of time depends on clock rate, with higher clock rates resulting in more work. EX1002, ¶42. But a higher clock rate comes with costs; voltage must be increased to support the clock rate, increasing both power drain and heat. *Id.* DVFS, “dynamic adjustments

to supply voltage, clock frequency, or both,” was well known to increase clock rate and voltage when needed for higher workloads, but also decrease them when possible to reduce power consumption and heat. EX1019, 35/219; EX1002, ¶¶42-43. By 2010, DVFS had already become “[t]he most popular technique for CPU power management.” EX1007, [0002].

B. Multi-core Processors and DVFS

Multi-core processors were also well known before the priority date of the ’339 Patent. EX1002 ¶33; EX1019, 52-53/219. By 2008, multi-core processors were already “becom[ing] the predominant general-purpose, high-performance microprocessor platform.” EX1019, 52/219; EX1023, 2/8; EX1007, FIG. 1, [0014]; EX1020, 1:16-20; EX1021, 2:53-57, FIG. 1; EX1022, FIG. 1. Multi-core processors “are typically used to gain higher performance by arranging the different processor cores to execute respective data processing operations in parallel” by “shar[ing] access to shared memory.” EX1013, [0004]; EX1002 ¶33.


By 2010, “[i]ndependent DVFS for each core,” cluster or “domain”-based DVFS, and “per-chip DVFS” had already become well known. EX1019, 52-53/219; EX1006, 2:15-44; EX1004, [0032]; EX1002, ¶34. 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, 52-53/219; EX1002,

¶34. Balancing costs of component and limited die space against the desire to set operating conditions granularly, it was well known to group cores into a “cluster” or “domain” that “share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1007, [0011]; EX1004, [0033]; EX1002, ¶34.

C. Clock Generation

The basic design principles for processor clock signals were also widely known by 2010. EX1002, ¶37. A clock signal may be provided via one or more of: (a) an external reference oscillator (e.g., crystal oscillator), (b) a phase locked loop (PLL), and (c) “clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals.” EX1028, 2:61-3:11; EX1006, 7:23-45; EX1002, ¶¶37-40.



EX1028, Figure 1 (annotated).

There were known alternatives to create multiple clock signals at different frequencies. EX1002, ¶38; *see e.g.*, EX1004, [0028]; EX1006, 7:23-45; EX1008, 1-2; EX1027, [0005]-[0010]. For example, it was known that “[t]he simplest way to implement multiple clock signals is to provide multiple oscillators.” EX1027, [0005]; EX1002, ¶38. Williams illustrates such a configuration with two external crystal oscillators each fed into a separate PLL:




Figure 1: Generic clock domain diagram

EX1008, Figure 1.

As another example, “clock trees or other circuitry” such as “a plurality of dividers” may be “configured to split or divide a clock signal into multiple individual signals” at different frequencies. EX1006, 7:23-45; EX1027, [0027]; EX1002, ¶40.




Fig. 3

EX1027, FIG. 3.

D. Processor Cache Memory and Cache Coherency

Processor designers have long used memory storage hierarchies, otherwise known as caches, to compensate for slow memory access speeds which may limit the performance of applications that depend on the rate at which data can be

accessed. EX1002, ¶46; *see* EX1026, 1/7; EX1030, 1/14. By 2010, it was well known that multiple levels of cache was a desirable configuration for multi-core processors, balancing the trade-off between memory size and speed. EX1002, ¶46; *see e.g.*, EX1006, FIGs. 1, 8. As shown in the figure below, each core may have a local cache that is unique to each core and not directly accessible by other cores. EX1030, 1/14; EX1002, ¶46; *see also* EX1013, FIG. 1A.




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


EX1030, Figure 2. But, to allow the multiple processor cores to execute data processing operations in parallel, it was known for the local caches to “share access to shared memory.” EX1013, [0004]; EX1002, ¶46; *see e.g.*, EX1023 5/8; EX1006, FIGs. 1, 8. Additionally, shared memory allows multiple “low-cost microprocessors [to be] economically interconnected with shared memory modules,” making this design particularly “cost-effective.” EX1030, 1/14.

It was “important to ensure that those processors will access the up-to-date data when seeking to access the associated address in shared memory.” EX1013, [0006]; EX1030, 1/14; EX1002, ¶47. The process of ensuring all processors operate using up-to-date data, cache coherency, was a known design requirement for processor designs that use a shared memory. EX1002, ¶47; *see* EX1013, [0006]-[0009]; EX1030, 1/14. Cache coherency schemes for multi-core processors have been well known since at least 1990. EX1002, ¶¶48-49; *see* EX1030, 1-2/14.

VIII. THE '339 PATENT

A. Overview

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



Id., FIG. 1 (annotated).

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




FIG. 2




FIG. 3

EX1001, FIGs. 2, 3 (annotated).

B. Prosecution History

Patent Owner filed Application No. 12/713,220, which resulted in the '339 Patent, on February 26, 2010, and was granted an allowance after Patent Owner amended the claims in response to the single Office Action issued by the Examiner. EX1005, 30-33/194 (Notice of Allowance), 90-100/194 (Non-Final Office Action rejecting claims on 101, 112, 102, and 103 grounds).

To overcome the 102 and 103 rejections in the Office Action, Patent Owner argued in an interview with the Examiner that the prior art relied on by the Examiner “disclosed using *a single reference clock*,” which was different than FIG. 3 of the

'339 Patent where “clock signals 1 through 3 were *different/independent clock signals* input to the PLLs.” EX1005, 72/194 (Applicant-Initiated Examiner Interview following the Office Action). In the same interview, “[t]he examiner agreed that Jacobowitz discloses using *a single reference clock*,” but disagreed the current claims excluded such an arrangement. *Id.* “Applicant’s representative stated that he would discuss amends [sic] to the claims.” *Id.* Two days later, Patent Owner filed a response amending the independent claims to add the phrase “the first clock signal *is independent* from the second clock signal.” *Id.*, 76, 78-80/194. The Examiner allowed the claims following this claim amendment without further comment. *Id.*, 30-33/194.

IX. CLAIM CONSTRUCTION

Petitioners interpret the Challenged Claims according to the *Phillips* claim construction standard. 83 Fed. Reg. 51340, 51340-44 (Oct. 11, 2018); *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005). For the reasons explained below, construction of these terms, or any others is not necessary to address the mappings presented in this Petition. The terms of the Challenged Claims can be given their plain and ordinary meaning for purposes of this proceeding. *See, e.g., Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017).

A. “the first clock signal is independent from the second clock signal”

Co-Petitioner NXP has requested that the district court construe this term to be given its plain and ordinary meaning, which “requires that the first and second clock signals are provided by or processed (i.e., divided or multiplied) from different reference oscillator clocks” (“Petitioner Construction”). EX1017, 4/25. This construction comes from prosecution history statements discussed in Section VIII.B. Patent Owner asserts that “the use of ‘independent’ is best understood as simply meaning ‘different’” (“PO Construction”). EX1018, 5/29. All Grounds include at least one mapping applicable to both constructions. The White Grounds include First and Second Configurations applying both constructions and a Third Configuration applying the PO Construction. The Elgebaly Grounds apply both constructions.

B. Indefinite Terms

Co-Petitioner NXP has challenged three terms as indefinite in the NXP Litigation because the bounds of each term cannot be determined by a POSITA with reasonable certainty. EX1017, 15/25, 17/25, 20/25. Without adopting PO’s positions (*Western Digital Corp. v. SPEX Technologies*, IPR2018-00084, Paper 14 at 11 (PTAB Apr. 25, 2018), the Petition applies Patent Owner’s urged scope:

- “each processor core from the first/second set of processor cores is configured to dynamically receive a first/second supply voltage and a

first/second output clock signal” “merely refers to the set of processor cores being configured to receive changing voltage and clock signals, such as for jobs with differing computational demands.” EX1018, 11/29

- “located in a periphery of the multi-core processor:” “[T]he periphery would be along the rectangle one could draw around all three components,” referring to the recited first and second sets of cores and “interface block.” *Id.*, 15-16/22.
- “located in a common region that is substantially central to the first set of cores and second set of processor cores:” “Common region” “references a region common to the first and second sets of processor cores,” and “substantially central” “refers to not along the outside of the multi-core processor but the inside.” *Id.*, 17-19/22.

X. REASONS FOR THE RELIEF REQUESTED (§§42.22(a)(2), 42.104(b)(4))

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

1. *White*

White discloses a multi-core processor in which each core “may be configured to operate at frequencies and voltages independently of the other logic cores.”

EX1006, 2:15-44. FIG. 2A illustrates the disclosed structure.




FIG. 2A

Id., FIG. 2A (annotated). “[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. “As with power management logic 200,” “temperature control logic 210” may also be configured to “adjust the voltages and frequencies supplied to logic cores 120A and 120B.” *Id.*, 13:27-43. “[Blue] voltage regulator 130 . . . may be configured to supply independent voltages

to each logic core,” each of which can be “different from the voltage levels at which other logic cores are operated.” *Id.*, 7:23-45, FIG. 2A. Similarly, “[purple] frequency regulator 140” may “supply multiple, independent clock signals” to the [green, red] logic cores 120A and 120B. *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 a particular voltage to [red] logic core 120B. *Id.*, 10:11-24. “[purple] frequency regulator 140” “instruct[s] frequency regulator 140 to supply a particular frequency or clock signal to [green] logic core 120A” and a particular clock signal to [red] logic core 102B. *Id.*

White discloses multiple options for generating and changing clock signals using a combination of elements at 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 configured to split or divide a clock signal into multiple individual signals.” *Id.*, 7:40-43. Additionally, each “logic core may generate its internal clock using a phase-locked loop (PLL) locked to an externally provided clock signal” from “frequency regulator 140.” *Id.*, 9:62-65. White thus

discloses or renders obvious a number of configurations relevant to mapping the Challenged Claims:

- **First Configuration:** Frequency regulator 140 includes “multiple clock sources” in the form of multiple crystal “oscillators” “in order to supply multiple, independent clock signals of different frequencies” (EX1006, 7:40-45), each feeding “a phase-locked loop (PLL)” at one of logic cores 120A or 120B (*id.*, 9:62-65). EX1002, ¶62.
- **Second Configuration:** Frequency regulator 140 includes “multiple clock sources” in the form of multiple PLL’s, each fed by a corresponding crystal “oscillator” “in order to supply multiple, independent clock signals of different frequencies” (EX1006, 7:40-45), each feeding “a phase-locked loop (PLL)” at one of logic cores 120A or 120B, or directly feeding the cores themselves (*id.*, 9:62-65). EX1002, ¶63; EX1008, FIG. 1, Section VII.C.
- **Third Configuration:** Frequency regulator 140 includes “clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals” (EX1006, 7:40-43), each feeding “a phase-locked loop (PLL)” at one of logic cores 120A or 120B, or directly feeding the cores themselves (*id.*, 9:62-65). EX1002, ¶64.

White also discloses the use of “common bridge logic, such as a Northbridge” “configured to communicate with the logic cores as well as with other elements or peripherals . . . such as a memory subsystem.” EX1006, 7:1-7, 17:29-32, FIGS. 2A, 2B. “Common bridge logic 110” can “interface with a logic core using the same logic levels and/or clock timing as the logic core.” *Id.*, 7:67-8:10. Common bridge logic further includes an interface “to translate the communication, such as the signals on one or more control lines . . . between the different frequencies.” *Id.*, 8:21-26, FIG. 3, EX1002, ¶65.

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

2. *Talwar*

Talwar discloses a multi-core processor design in which “[t]he cores and voltage sources are divided into clusters, whereby multiple cores in a cluster receive power from a single voltage source” and “[e]ach cluster may have a different cluster frequency,” but “all the cores in a cluster [operate] at the same frequency.” EX1007, [0011], [0018]-[0019]. The values of: (a) the “single voltage source” for the cluster, and (b) “the frequency of all cores in a cluster” can be adjusted “[d]uring power management.” *Id.*, [0011]-[0012], [0018-0019], [0021-0025], [0030] (“The cores in

a cluster ... can be independently controlled from other clusters.”). Providing one voltage and clock source per cluster enables “maximum power savings” by efficiently performing DVFS on each cluster in “stark contrast to conventional multi-core chips” with only a single voltage source. *Id.*, [0012], [0021]; EX1002, ¶69.

3. *Motivation to Combine*

A POSITA would be motivated to apply Talwar’s DVFS “clustering” of processor cores to White’s DVFS implementation in order to achieve the goal common to White and Talwar—DVFS in a multicore processor below the chip level—while meeting practical chip space constraints. EX1002, ¶71. Specifically, while White illustrates two logic cores, it explains that a device may “include ***any number of logic cores*** configured to operate at independent voltages and/or frequencies.” EX1006, 6:56-60. Following this disclosure, a device may include four—or eight or sixteen—cores, each with separate voltage and clock sources. *Id.*; EX1002, ¶74. However, a POSITA would have understood such a design would need to be adapted as the number of cores increases due to chip space constraints associated with adding an additional voltage and clock control and supply circuitry for each additional core. *Id.*, ¶¶73-74 (citing EX1029, 3-4/12, EX1019, 52-53/219).

Talwar discloses achieving the same goal as White of performing DVFS below the chip level by grouping cores into “clusters” and achieving “maximum

power savings”—in “stark contrast to conventional multi-core chips” with only a single voltage source (*id.*, [0021], [0012])—but simultaneously reducing the costs of the “more expensive” “per-core DVFS” (EX1019, 52-53/219) by having “less voltage sources than cores on the chip” (EX1007, [0014]). EX1002, ¶72. A POSITA would have been further motivated to apply Talwar’s DVFS “clustering” of processor cores to White’s DVFS implementation because White explicitly states that its teachings are not limited to per-core voltage and clock signals, but may more generally “be applied to any integrated circuit or component that includes multiple areas or sections of circuitry configured to operate at independent operating voltages and/or frequencies.” EX1006, 8:44-49; *id.*, 3:37-52, 6:63-67; EX1002, ¶75.

4. *Modification of White in View of Talwar*

As discussed above, White explicitly discloses applying its disclosures to a device that “include[s] *any number of logic cores* configured to operate at independent voltages and/or frequencies.” EX1006, 6:56-60. Following this disclosure, White’s FIG. 2 embodiment, including logic cores 120A-120B can be modified to include four logic cores 120A-120D. EX1002, ¶77. White’s common bridge logic would have a connection to each additional logic core. EX1002, ¶79. Then, applying Talwar’s clustering disclosure, [green] logic cores 120A and 120C constitute a first cluster and [red] logic cores 102B and 120D constitute a second cluster as shown below:




FIG. 2A

See EX1006, FIG. 2A (modified); EX1007, FIG. 1.

In accordance with Talwar’s disclosure providing a “single voltage source” for the cluster and the same frequency for “all cores in a cluster” (EX1007, [0011], [0019]), which a POSITA would understand refers to a single clock source for the cluster, each cluster (rather than each core) would receive a voltage signal from voltage regulator 130 and a single clock signal from frequency regulator 140. EX1002, ¶¶77-78. Further, clock control circuitry that White describes at the core level (e.g., a PLL at the core and accompanying “control registers” storing “PLL configuration parameters”) would operate at the cluster level. EX1002, ¶78. With

those modifications, the system would perform the same power management, or DVFS methods as White, but would do so on a cluster-by-cluster basis. *Id.*, ¶¶77, 80.

5. *Element-by-Element Analysis*

a. Claim 1

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

To the extent the preamble may be limiting, White/Talwar renders obvious Limitation [1.pre]. EX1002, ¶82. 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,

White/Talwar renders obvious Limitation [1.1]. EX1002, ¶83. As discussed in Section X.A.4 and shown in modified Figure 2A, White/Talwar implements Talwar’s teaching that “[t]he cores and voltage sources are divided into clusters.” EX1007, [0011], Claims 1, 25; EX1002, ¶83. Cores 120A and 120C are a first cluster (first set of processor cores) and cores 120B and 120D are a second cluster (second set of processor cores) of integrated circuit 100 (the multi-core processor). EX1002, ¶84.




FIG. 2A

EX1006, FIG. 2A (modified).

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

White/Talwar renders obvious Limitation [1.1.1]. EX1002, ¶85. 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 and a first / second clock signal from frequency regulator 140. EX1006, 10:11-11:23; EX1007, [0011], [0019]; EX1002, ¶¶85, 93.




FIG. 2A

EX1006, FIG. 2A (modified).

As discussed in Section X.A.1, White also teaches that the provided voltage and clock signals are dynamic as they are “adjusted” as needed, for example to support “power management” or “temperature management” requirements. EX1006, 2:15-44; EX1002, ¶85.

In the **First Configuration** described in Sections X.A.1, X.A.4 in which frequency regulator 140 includes “multiple clock sources” in the form of multiple crystal “oscillators” “in order to supply multiple, independent clock signals of different frequencies” (EX1006, 7:40-45), each feeding “a phase-locked loop (PLL)” at the first / second cluster (*id.*, 9:62-65), the output of the first / second

crystal oscillator is the first / second clock signal and an input to the PLL at the first / second cluster whose output is the first /second output clock signal. EX1002, ¶¶87, 90-91.

In the **Second Configuration** described in Section X.A.1, X.A.4 in which frequency regulator 140 includes “multiple clock sources” in the form of multiple PLL’s, each fed by a corresponding crystal “oscillator” “in order to supply multiple, independent clock signals of different frequencies” (EX1006, 7:40-45), each feeding “a phase-locked loop (PLL)” at the first / second cluster, or directly feeding the cores themselves (*id.*, 9:62-65), the output of the first / second crystal oscillator is the first / second clock signal and an input to the first / second PLL at frequency regulator 140 whose output is the first /second output clock signal. EX1002, ¶¶88, 90-91.

In the **Third Configuration** described in Section X.A.1, X.A.4 in which Frequency regulator 140 includes “clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals” (EX1006, 7:40-43), each feeding “a phase-locked loop (PLL)” at the first / second cluster, or directly feeding the cores themselves (*id.*, 9:62-65), the first / second individual signal output of the clock tree or other circuitry at frequency regulator 140 is the first / second clock signal and an input to the PLL at the first / second cluster whose output is the first /second output clock signal. EX1002, ¶¶89-91.

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

White/Talwar renders obvious Limitation [1.2] for the reasons discussed for Limitation [1.1]. EX1002, ¶94.

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

White/Talwar renders obvious Limitation [1.2.1] for the reasons discussed in for Limitation [1.1.1]. EX1002, ¶95.

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

White/Talwar renders obvious Limitation [1.2.2]. EX1002, ¶96. 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]; EX1002, ¶96. Therefore, in the White/Talwar combination, the first supply voltage to the first cluster is independent from the second supply voltage to the second cluster. EX1002, ¶96.

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

White/Talwar renders obvious Limitation [1.2.3]. EX1002, ¶97. In the **First and Second Configurations** described in Sections X.A.1, X.A.4 in which the output of the first / second crystal oscillator is the first / second clock signal, the first and second clock signals are independent under both Petitioner and PO's Constructions

because they are “multiple, independent clock signals of different frequencies” (EX1006, 7:38-45) originating from two different external crystal oscillators. EX1002, ¶90. In the **Third Configuration** described in Sections X.A.1, X.A.4 in which the first / second individual signal output of the clock tree or other circuitry at frequency regulator 140 is the first / second clock signal, the first and second clock signals are independent under PO’s Construction because they are different “multiple individual signals” created by “split[ting] or divid[ing]” a clock signal (EX1006, 7:38-45). EX1002, ¶90.

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

White/Talwar renders obvious Limitation [1.3]. EX1002, ¶98. As discussed in Section X.A.1, White’s common bridge logic block is “an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores.” *Id.* As shown in modified Figure 2A and explained in White, the common bridge logic is “configured to communicate with the logic cores as well as with other elements or peripherals . . . such as a memory subsystem.” EX1006, 7:1-7; EX1002, ¶98.




FIG. 2A

EX1006, FIG. 2A (modified). A POSITA would have understood that each processor core must be separately connected to the common bridge logic to allow each processor to interface with “other elements and peripherals” of the system, as each processor executes a separate thread or application. EX1002, ¶98; *see* EX1006, 7:1-7. 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. EX1002, ¶99. Accordingly, the common bridge logic, memory subsystem, and/or cache coherency logic are the interface block.

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

White/Talwar renders obvious Limitation [1.3.1] at least under PO's interpretation of this element. EX1002, ¶100. PO has asserted that "maintaining coherency between caches in the cores or complexes and/or with the L3 shared memory," including "a cache coherent interconnect," is sufficient to meet this element. EX1015, 58/65; EX1014, 40/26. White/Talwar renders this functionality obvious. As discussed for Limitation [1.3], a POSITA would understand that each of cores of the first and second sets of processor cores may read and write to common space in the memory subsystem, assisted by cache coherency, via the common bridge logic. EX1002, ¶101; Section VII.D. At least under PO's interpretation of this element, the memory subsystem enabling cores from the first and second set of processor blocks to write to a common memory, as well as the use of cache coherency, facilitate communication between the first and second set of processor cores. EX1002, ¶101.

Further, 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. EX1002, ¶102.

b. Claim 2

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

White/Talwar renders obvious Limitation [2.1] for reasons discussed above in Limitation [1.3.1]. EX1002, ¶103. As discussed in Section VII.B, it was well known in the art that multi-core processors implementing DVFS required mechanisms for translating logic levels between chip components operating at different voltages. *Id.* 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. *Id.*; see EX1006, 8:13-40, 14:55-15:17, FIG. 3. Further, there were well-known prior art techniques for doing

so. *See e.g.*, EX1031, FIGs. 4-5, 5:45-57, EX1039, 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. EX1002, ¶103.

c. Claim 3

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

White/Talwar renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1]. EX1002, ¶104.

d. Claim 5

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

White/Talwar renders obvious Limitation [5.1]. EX1002, ¶105. White discloses that the “logic circuitry or firmware configured to control the operating voltage and frequency of logic cores” may “reside in the processor itself.” EX1006, 5:17-24; *see also id.*, 13:7-26 (“logic circuitry...internal...to integrated circuit”), 14:32-38 (“temperature control logic [] and power management logic [] may...be

part of a single set of logic.”), FIGs. 2A, 2B. The logic circuitry within the processor may include the “power management logic” that sends control signals to the processor cores by “writ[ing]to one or more control registers” within each processor core and constitute the one or more control blocks. EX1006, 10:7-9. EX1002, ¶105. A POSITA would have understood that circuitry design involves layout decisions that balance performance, area, and thermal considerations, and that placement of control logic circuitry along the perimeter of the multi-core processor or substantially central to processor cores is an obvious design choice. EX1002, ¶105; *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 SC [storage control units],” which is substantially central to the processor units).

e. Claim 8

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

White/Talwar renders obvious Limitation [8.1]. EX1002, ¶106. 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. Further, it would have been obvious to a POSITA to arrange logic cores 120A and 120C in a first row constituting a first region and logic cores 120B and 120D in a second row constituting a second non-overlapping region. First, arranging logic cores in rows was a well-known approach to physical layout of a multicore processor. EX1002, ¶106 (citing EX1033, FIG. 8, EX1022, 2:8-10, FIGs. 1-2, EX1036, 4:53-57, FIG. 2, EX1011, FIG. 3) Second, arranging the first cluster in one row and second cluster in the second row enables efficient signal routing to route the shared clock signal and voltage signal to each core in a cluster. EX1002, ¶107. Third, placing the cores of each cluster together and, thus, in non-overlapping regions enables cluster-level clock control circuitry to be placed close to all cores in the cluster. EX1002, ¶108.

f. Claim 10

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

White/Talwar renders obvious Limitation [10.1] for reasons discussed in Limitation [8.1]. EX1002, ¶109.

g. Claim 11

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

White/Talwar renders obvious Limitation [11.1] for reasons discussed in Limitation [8.1]. EX1002, ¶110.

h. Claim 14

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

White/Talwar renders obvious Limitation [14.1] for reasons discussed in Limitation [5.1]. EX1002, ¶111.

i. Claim 21

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

White/Talwar renders obvious Limitation [21.pre] for reasons discussed in Limitation [1.pre]. EX1002, ¶112.

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

White/Talwar renders obvious Limitation [21.1] for reasons discussed in Limitation [1.1]. EX1002, ¶113.

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

White/Talwar renders obvious Limitation [21.1.1] for reasons discussed in Limitation [1.1.1]. EX1002, ¶114. 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 (the power control block). *Id.* In the **First and Third Configurations**

in which the PLL at the first / second cluster receives as an input the first / second clock signal, the region on the chip encompassing these local PLLs is a clock control block. EX1002, ¶115. In the **Second Configuration** in which the PLL at frequency regulator 140 receives as an input the first / second clock signal, frequency regulator 140 is a clock control block.

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

White/Talwar renders obvious Limitation [21.2] for reasons discussed in Limitation [1.2]. EX1002, ¶117.

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

White/Talwar renders obvious Limitation [21.2.1] for reasons discussed in Limitations [1.2.1] and [21.1.1]. EX1002, ¶118.

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

White/Talwar renders obvious Limitation [21.2.2] for reasons discussed in Limitation [1.2.2]. EX1002, ¶119.

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

White/Talwar renders obvious Limitation [21.2.3] for reasons discussed in Limitation [1.2.3]. EX1002, ¶120.

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

White/Talwar renders obvious Limitation [21.3] for reasons discussed in Limitation [1.3]. EX1002, ¶121.

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


White/Talwar renders obvious Limitation [21.3.1] for reasons discussed in Limitation [1.3.1]. EX1002, ¶122.

B. Ground 2: White, Talwar, and Cho Render Obvious Claims 5, 14

1. Cho

Cho discloses an “integrated circuit” that includes a plurality of “functional units...implemented on a die of the integrated circuit.” EX1010, [0007]. The functional units may each represent a “processor core.” *Id.*, [0038], Claim 17. In addition, a “temperature control unit (TCU)” is coupled to each of the functional units and “is configured to perform temperature control actions on each, independently of any temperature control actions performed on other ones of the functional units [e.g., processor cores].” *Id.*, [0038], [0047]-[0048]. Specifically, Cho explains that “in addition to the temperature signals, the clock signals, and the voltage control signals [green], processor cores [] may also be coupled to TCU [] by corresponding core control signals ctrl1 and ctrl2 [red], respectively.” *Id.*, [0042], Figs. 4. “These control signals may be bi-directional and/or may be transmitted via

multiple signal lines” and through the use of these signals, the “TCU [] may also perform temperature control actions that involve re-allocating the processing load between processor cores [] in an effort to control their respective temperatures.” *Id.*; *see also id.*, [0047]-[0048].



EX1010, Figures 1, 4 (annotated). As depicted in Figures 1 and 4, [blue] TCU may be located within the multi-core processor substantially central to the functional units (e.g., processor cores) as well as in the periphery of the multi-core processor. *Id.*; EX1002, ¶123.

2. Motivation to Combine White/Talwar with Cho

A POSITA would have been motivated to combine White/Talwar with Cho to achieve the goal of controlling processor core temperatures that is common to White and Cho. EX1002, ¶124. A POSITA would be motivated to include Cho’s temperature control functionality because it achieves temperature control in a

different way and is, thus, additive and adds extra flexibility. *Id.* Specifically, White discloses “[i]f temperature control logic 210 determines that the operating temperature of integrated circuit 100 is too high,” it may “automatically adjust the operating voltage and/or frequency of one or more logic cores.” EX1006, 13:27-31. Rather than adjusting operating voltage and/or frequency, Cho discloses when a processor core “reports a temperature that exceeds a predetermined threshold,” the “TCU [] may attempt to reduce the temperature of processor core 405 by re-allocating a portion of its workload.” EX1010, [0042]. Thus, Cho discloses a third approach—re-allocating a portion of a core’s workload—along with those already disclosed by White. A POSITA would be motivated to include this additional functionality both to increase flexibility in addressing temperature issues and because this method “may optimize both performance as well as thermal control.” EX1002, ¶¶125-126; EX1010, [0047].


3. *Modification of the White/Talwar Combination in View of Cho*

The functionality of Cho’s TCU 102 may be incorporated into White’s existing control logic circuitry (i.e., the “temperature control logic 210” and “power management logic 200”). EX1006, 5:17-20; EX1002, ¶127. White explains that these components “may reside in the processor itself” (EX1006, 5:17-20) and “may both be part of a single set of logic” (*id.*, 14:32-38). EX1002, ¶127. Cho provides additional guidance on the location of the control circuitry beyond White’s

disclosure of “in the processor” (i.e., integrated circuit 100), illustrating the control circuitry as alternatively substantially central to the functional units (*e.g.*, processor cores) (EX1010, Figure 1) or in the periphery of the multi-core processor (*id.*, Figure 4). EX1002, ¶128. A POSITA would have understood both of these locations to represent acceptable known design choices. EX1002, ¶128 (citing EX1033, FIG. 8, [0056], [0103]-[0107]).

Finally, Cho’s “core control signals cctrl1 and ccrtl2” that couple the “TCU” to “processor cores” follow the existing control path signals White describes to enable “power management logic” to “write to one or more control registers” in the processing cores. EX1006, 10:4-8; EX1002, ¶129.

Accordingly, as depicted below in the modified portions of Figure 2B, the “control signals cctrl1 and ccrtl2” [red] as taught by Cho could be integrated within the power/temperature management set of logic (i.e., logic circuitry [blue]) in the White/Talwar combination and may be located within the multi-core processor as taught by both White and Cho. EX1002, ¶130; EX1010, Fig. 2B (modified below).



EX1006, FIG. 2B (modified).

4. *Element-by-Element Analysis*

a. **Claim 5**

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

White/Talwar/Cho renders obvious Limitation [5.1]. EX1002, ¶131. White discloses that the “logic circuitry or firmware configured to control the operating voltage and frequency of logic cores” may “reside in the processor itself.” EX1006, 4:49-5:17-24, 13:7-26. Further, as explained in Section X.B.3, the logic circuitry [blue] provides control signals to the processor cores and is located along the perimeter of the multi-core processor as shown in Fig. 4. EX1002, ¶131; *see also* EX1010, [0038]-[0040], [0042], [0047]-[0048], Fig. 4.

b. Claim 14

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

White/Talwar/Cho renders obvious Limitation [14.1] for reasons discussed in Limitation [5.1]. EX1002, ¶132. Further, as explained in Sections X.B.3, in addition to locating logic circuitry [blue] along the periphery of a multi-core processor, a POSITA would have understood that an obvious design choice would be placement of logic circuitry that provides control signals from a common region that is substantially central to the functional units (e.g., processor cores). EX1002, ¶132; *see also* EX1010, [0018]-[0019], [0021]-[0022], Fig. 1.

C. Ground 3: White, Talwar, and Vajda Render Obvious Claims 8-9

1. Vajda

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]. Vajda provides guidance for “the geometry” (i.e., layout of the individual cores) “of the multi-core chip 230.” *Id.* 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 White/Talwar with Vajda*

A POSITA would have been motivated to apply Vajda’s core layout guidance in defining the physical layout of cores in White/Talwar to obtain Vajda’s stated benefit of a layout that may “avoid hot spots and over-heating.” EX1002, ¶134; EX1011, [0062]. As discussed in Section X.A.4, White/Talwar includes a first cluster of logic cores 120A and 120C and a second cluster of logic cores 120B and 120D, each having independent voltage and clock signals that enable operating one cluster at a different higher frequency—producing more heat—than the other. Thus, the problem Vajda seeks to address where a layout includes “processor cores 231 that can execute at higher frequencies” and “other processor cores 232, running at lower frequencies” is applicable to White/Talwar, and a POSITA would be motivated to apply Vajda’s layout guidance to “avoid hot spots and over-heating.” EX1002, ¶134; EX1011, [0062].

3. *Modification of the White/Talwar Combination in View of Kato*

As discussed in Section X.A.4, White/Talwar includes a first cluster of logic cores 120A and 120C and a second cluster of logic cores 120B and 120D, each having independent voltage and clock signals that enable operating one cluster at a different higher frequency than the other. In accordance with Vajda’s disclosure 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” (EX1011, [0062]), a POSITA would be motivated to arrange the cores of the first cluster and the second cluster in a checkerboard pattern (e.g., a first row of logic core 120A followed by 120B and a second row of logic core 120D followed by 120C) in order to “spread out” cores of one cluster from each other and avoid any hot spots created by grouping of cores of one cluster operating at a higher frequency. EX1002, ¶135.

4. Element-by-Element Analysis

a. Claim 8

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

White/Talwar/Vajda renders obvious Limitation [8.1]. EX1002, ¶136. As discussed in Section X.C.3, White/Talwar/Vajda includes logic cores 120A and 120C (the first set of processor cores) and logic cores 120B and 120D (the second set of processor cores) arranged in a checkerboard pattern (e.g., a first row of logic core 120A followed by 120B and a second row of logic core 120D followed by 120C). The region containing the first set of processor cores is the first region and the region containing the second set of processor cores is the second region. *Id.*

b. 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.

White/Talwar/Vajda renders obvious Limitation [9.1] for reasons discussed in Limitation [8.1]. EX1002, ¶137. The first and second regions are an overlapping region because the checkerboard pattern of cores results in overlapping regions containing each cluster. *Id.*

D. Ground 4: Elgebaly and Ozer Render Obvious Claims 1-3, 5, 8, 10-11, 14, and 21

1. Elgebaly

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, “ASIC 120 may include ... any number of processor cores.” *Id.*, [0030]. Elgebaly therefore discloses or at least renders obvious an ASIC 120 with four cores, as in Modified Figure 1 below, with additional [green] Processor Core 3 and [red] Processor Core 4. EX1002, ¶138.




FIG. 1
(MODIFIED)

EX1002, ¶138; *see* EX1004, FIG. 1, [0033]. Consistent with Elgebaly’s teaching and as described in Section VII.B, multi-core processors containing four cores were also well known and available commercially before 2010. EX1002, ¶138; *see e.g.*, EX1023, 2/8, EX1004, [0030], [0033].

a. Independent clock signals and voltages

Elgebaly’s Background that “[a]n effective method for reducing power consumption ... is to scale or adjust the supply voltage.” *Id.*, [0006]. Reducing power consumption can be important, for example, to extend battery life. *Id.* To this end, and as in Figure 2 (annotated below), Elgebaly discloses an embodiment of ASIC 120 with “adaptive voltage scaling (AVS),” in which “one [blue] AVS

[adaptive voltage scaling] unit 230 is provided for each [green, red] processing core

130 ... to adaptively scale the supply voltage for that core.” EX1004, [0031].




FIG. 2

Id., FIG. 2 (annotated).

“Each core 130 provides to clock generator 140 and to its AVS unit 230 a target frequency for its clock,” which “frequency may be selected based on the [core’s] processing load and computational requirements.” *Id.*, [0032]. “For each processing core 130, clock generator 140 generates the clock at the target frequency and provides this clock to the core [and] associated AVS unit 230.” *Id.* “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.” *Id.*, [0033].

Clock generator 140 uses “one or more phase locked loops (PLLs) that control the oscillators” to generate corresponding clock signals. *Id.*, [0028]. The clock generator also “may couple to external circuitry, which may include *crystals*.” *Id.* For the embodiment of Figure 2, in which clocks “may be set independently” for each core, a POSITA would have understood Elgebaly’s disclosure of one or more PLLs and external circuitry including plural “crystals” to include known clock distribution channels; specifically, multiple PLLs each receiving a distinct clock signal from a different external oscillator crystal and, in turn, outputting one of the desired clock signals. EX1002, ¶141. At the very least, such a configuration was an obvious design choice that would have allowed a POSITA to implement multiple, independently-settable clock signals and thereby achieve Elgebaly’s flexibility and potential power savings. EX1002, ¶142. Such a configuration would have been desirable to maximize independence and flexibility between the clock signals sent to AVS Unit 1 and AVS Unit 3 and thereby maximize Egebaly’s benefits of independently controlling clock signals and voltages. *Id.* Indeed, such a configuration using separate oscillators and PLLs was known, as explained in Section VII.C above. *Id.*

Once clock signals are generated, “[t]he supply voltage for each processing core 130 is set by the associated AVS unit 230 such that the core can operate at the target clock frequency.” EX1004, [0032]. Thus, for each processing core, “the clock ... may be set independently ..., and the supply voltage ... may be set independently.” *Id.*, [0033]. 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].

b. Sets of cores

As noted above, Elgebaly explains that “ASIC 120 may include ... any number of processor cores.” *Id.*, [0030]. Elgebaly therefore discloses or at least renders obvious an ASIC 120 with four cores. EX1002, ¶144. Elgebaly also discloses that “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1004, [0033]; EX1002, ¶144. Elgebaly therefore discloses or at least renders obvious a chip with four processor cores where: the first two of the processor cores share a first common clock, a first common supply voltage, and a first AVS Unit, and where the other two of the processor cores share a second common clock, a second common supply voltage, and a second AVS Unit, as shown in Modified Figure 2, inset below.




FIG. 2
(MODIFIED)

A POSITA would have also understood that this type of core grouping was a well-known approach to balancing competing design factors. EX1002, ¶145. For example, a POSITA would have understood that chip space is limited and, as the number of cores increases, adding individual power generation and control circuitry, AVS units, and clock generation circuitry for each core becomes more “expensive” both in terms of space and cost. *Id.* Additionally, each additional processor core provides a roughly linear increase in processing power, whereas the voltage and frequency flexibility of setting cores differently may give diminishing returns at the cost of valuable chip space for additional power generation and control circuitry, AVS units, and clock generation circuitry. *Id.* Accordingly, a POSITA would have

understood both from this knowledge and Elgebaly's teaching of grouping processor cores into sets, each with a common clock and voltage, that it would have been obvious in a chip with four or more processor cores to group the processor cores into two sets, such that each set is managed by one of Elgebaly's AVS units that receives a corresponding clock signal from clock generator 140, as disclosed in Elgebaly. EX1002, ¶145; *see* EX1004, FIG. 2; *see also* EX1007, [0011]-[0012] (disclosing that clusters of cores were known in the art).

c. Interface block configured to facilitate communication between sets of processor cores

Elgebaly also discloses an internal memory used by the ASIC's processors. EX1004, [0027]-[0028], FIG. 1. Figure 1 (inset below) shows a communication interface between the processors and a [yellow] internal memory 150 that “stores data and program codes used by the processing units within ASIC 120.” *Id.*, [0028].





FIG. 1

Id., FIG. 1 (annotated). A POSITA would have understood from this disclosure that this shared, internal memory facilitates communication via the storage and transfer of “data and program codes used by the processing units” (e.g., 130c, 130d). EX1002 ¶146; *see* EX1004, FIG. 1, [0028]; *see also* EX1004, FIG. 2, [0033]; Section VII.D. Elgebaly, however, provides little detail on the implementation of shared memory.

2. *Ozer*

Ozer discloses a cache coherency method for shared memory of a processor “compris[ing] a plurality of processor cores [], each provided with its own local cache circuitry.” EX1013, [0068]. Each core’s local cache is updated by a “cache coherency fabric” associated with a “snoop bus” that work in tandem to “ensure that when one processor updates data in its local cache, any copies of that data in the caches associated with the other processors will either be appropriately updated or invalidated in those caches to ensure coherency.” *Id.*, [0072]-[0073]. For example, a write-update scheme requires local caches to send the updated value of the data to each processor that holds an old copy of the data in its local cache. EX1002, ¶147, *supra* Section VII.D.

Ozer’s Figure 1A (inset below) depicts an example of a four-core processor cache coherency system:



EX1013, FIG. 1A. As shown, “processor cores 10, 20, 30, 40, [are] each provided with its own local cache circuitry 15, 25, 35, 45, respectively.” *Id.*, [0068]. That “cache circuitry may comprise a single cache or may comprise multiple caches.” *Id.* Each local cache stores data from “memory 70,” which includes at least one “shared memory region in which the processors 10, 20, 30, 40 can store data that is sharable amongst those processors.” *Id.*, [0071]. Since the cores share data, “it needs to be ensured that the copy maintained by the local cache circuitry is up-to-date.” *Id.* This coherency is maintained by the “cache coherency fabric 60” using “coherency operations.” *Id.*, [0072]-[0073]; EX1002, ¶148.

“When a processor core wishes to access data in memory 70, it issues an access request specifying an address of that data in memory.” EX1013, [0069].

Then, as described generally in Section VII.D, “a lookup operation is typically performed in the cache circuitry associated with that processor core in order to determine whether the data can be accessed locally in the cache.” *Id.* “[T]he cache coherency fabric 60 will initiate a coherency operation ... “each time a cache miss occurs on a read request to shared data, or whenever a write access request is issued in respect of shared data.” *Id.*, [0073]. When a coherency operation is required, “notifications,” or “snoop requests” are “issued to each of the other caches ... to perform a snoop operation in order to determine whether any of those caches store data at the address specified by the access request.” *Id.* Additionally, Ozer’s protocols are designed to meet cache coherency needs while minimizing chip power consumption. *Id.*, [0022]; EX1002, ¶150.

3. Motivation to Combine Elgebaly and Ozer

A POSITA would have understood that the use of a local cache at each processor core was an obvious design choice in a multi-core processor like that of Elgebaly. EX1002, ¶151; *see* EX1013, FIG. 1A. For example, it was known in art of processor design that the time required to access main memories can limit processor speeds. EX1002, ¶151; *see* EX1026, 1/7. Further, it was known by 2010 that processor idle time while waiting on memory accesses could be reduced with a hierarchy of caches, for example: a small, fast, local cache; a medium-sized, medium-speed shared cache; and a larger, slower bulk storage cache. EX1002, ¶151,

supra Section VII.D. A POSITA would have understood Elgebaly's disclosure of internal memory with off-chip volatile and non-volatile memory to contemplate such a memory hierarchy. EX1002, ¶151; *see also* EX1004, FIG. 1, [0028]-[0029]. A POSITA would have also expected each of Elgebaly's processor cores to include its own cache. EX1002, ¶151. But Elgebaly is silent on individual caches for each processor core. *Id.*

Ozer, however, explains that “[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 [for] a subset of the data held in the shared memory.” EX1013, [0005]. Ozer further explains that using local caches to improve access speed may also “complicate[] the issue of data coherency. *Id.*, [0005]. But Ozer also provides a method for managing cache coherency. *Id.*, Abstract.

A POSITA therefore would have been motivated to implement certain of Ozer's teachings with multi-core processors of Elgebaly to “improve speed of access to data” and manage cache coherency to realize the benefits of that improved speed. EX1002, ¶153; *see* EX1013, [0004]-[0006]; EX1004, [0058] (recognizing memory access delays exist and vary for different types of memory).

In addition, Ozer's cache coherency technique “achieve[s] significant energy savings” by sending snoop requests selectively, rather than to every possible cache whenever a coherency operation is initiated (EX1013, [0018]-[0021]), which aligns

with Elgebaly's goal of "conserving battery power." EX1004, [0003], [0008]. A POSITA thus would have also been motivated by a desire to further improve energy savings in the combination of teachings from Elgebaly and Ozer. EX1002, ¶154.

Particularly given Elgebaly's and Ozer's shared focus of improving multi-core processors, a POSITA would have understood that such local caches and cache coherency methods would predictably improve the speed of memory access to Elgebaly's internal memory, which "stores data and program codes used by the processing units" and experiences "memory access delays." EX1004, [0028], [0058]; EX1002, ¶155.

4. Modification of Elgebaly in View of Ozer

As discussed above, a POSITA would have been motivated to modify Elgebaly to add a local cache and coherency circuitry to each processor and implement Ozer's techniques using the on-chip bus depicted in Figure 1 of Elgebaly. EX1002, ¶156. By 2010, local processor caches, such as those in Ozer were well known. *Id.*; *see supra* Section VII.D. Thus, Ozer tells a POSITA **how** to implement a known feature that is tangential to Elgebaly's focus, but was commonplace in the art of multi-core processors. EX1002, ¶156. The implementation of that memory hierarchy and cache coherency scheme would not change the operation of Elgebaly's invention, which is related to adaptive voltage and clock control and is fully compatible with Ozer's approach. *Id.*; *see* EX1004, FIGs. 1-3.

For the reasons explained in Sections X.D.1-3 above, at least one obvious combination of Elgebaly and Ozer (the “Elgebaly/Ozer Processor”) would have included the structure shown in the modified versions of Elgabaly’s Figures 1 and 2 (inset again below), in which (a) clock generator 140 includes separate PLLs, each receiving a signal from a respective one of two external oscillators, to generate distinct clock signals to each of AVS Unit 3 and AVS Unit 4, (b) each of Processor Cores 1-4 include local caches, (c) cache coherency fabric and protocols are implemented in ASIC 120 to facilitate sharing of data between and for use by Processor Cores 1-4 as detailed in Ozer. EX1002, ¶157.




FIG. 1
(MODIFIED)




FIG. 2
(MODIFIED)

5. Element-by-Element Analysis

a. Claim 1

[1.pre]

To the extent the preamble may be limiting, Elgebaly/Ozer renders obvious Limitation [1.pre]. EX1002, ¶158. As discussed in Section X.D.1, Elgebaly’s “ASIC 120 may include ... any number of processor cores.” EX1004, [0030].

[1.1]

Elgebaly/Ozer renders obvious Limitation [1.1]. EX1002, ¶159. Elgebaly teaches that its multi-core processor “may include ... any number of processor cores.” EX1004, [0030]. Elgebaly also 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]; EX1002, ¶159. For the reasons explained in Sections X.D.1-4 above, the Elgebaly/Ozer Processor depicted in modified Figure 2 would include a first set of [green] processor cores (“Processor Core 1” and “Processor Core 3”), and a second set of [red] processor cores (“Processor Core 2” and “Processor Core 4”).




FIG. 2
(MODIFIED)

EX1002, ¶¶159-161.

[1.1.1]

Elgebaly/Ozer renders obvious Limitation [1.1.1]. EX1002, ¶162. As explained in Section X.D.1.a, Elgebaly discloses that “the clock for each processing core 130 may be set independently by clock generator 140, and the supply voltage for each processing core 130 may be set independently by power management unit

180.” EX1004, [0033]. And, as explained in Section X.D.1.b, Elgebaly also discloses that a set of “multiple processing cores may share a common clock and/or a common supply voltage that may be jointly set for these cores.” EX1004, [0033]; EX1002, ¶162. Thus, for the reasons explained in Sections X.D.1-4 above, in the Elgebaly/Ozer Processor depicted in modified Figure 2, the cores of the first set of [green] processor cores (“Processor Core 1” and “Processor Core 3”) 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 4”) would be configured to receive a second supply voltage (“Vdd4”) and a first output clock signal (“Clock 4”). EX1002, ¶162.

Additionally, Elgebaly teaches that the supply voltages are “adaptive,” meaning the system “adjusts the supply voltage to a level that satisfies performance requirements,” as they change over time, e.g., when entering an “idle mode.” EX1004, [0034], [0006]; EX1002, ¶163. Thus, in the Elgebaly/Ozer Processor, the cores in the first set are configured to dynamically receive a first supply voltage, and the cores in the second set are configured to dynamically receive a second supply voltage. EX1002, ¶163; *see* EX1004, [0033].

Elgebaly’s cores are also configured to dynamically receive respective clock signals. EX1002, ¶164. “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.*

Finally, as explained in Section X.D.1.a, Elgebaly discloses or at least renders obvious that clock generator 140 includes a plurality of PLLs each receiving a respective clock signal from an external oscillator, with a first of such PLLs receiving a first clock signal from a first oscillator crystal and generating a first output clock signal (“Clock 3”) to AVS Unit 3 for the first set of cores (Processor Cores 1, 3), and with a second of such PLLs receiving a second clock signal from a second oscillator crystal and generating a second output clock signal (“Clock 4”) to AVS Unit 4 for the second set of cores (Processor Cores 2, 4). *See id.*, [0028]; EX1002, ¶165.

[1.2]

Elgebaly/Ozer renders obvious Limitation [1.2] for the reasons discussed for Limitation [1.1]. EX1002, ¶167.

[1.2.1]

Elgebaly/Ozer renders obvious Limitation [1.2.1] for the reasons discussed in for Limitation [1.1.1]. EX1002, ¶168.

[1.2.2]

Elgebaly/Ozer renders obvious Limitation [1.2.2] for reasons discussed for Limitation [1.1.1] and [1.2.1]. EX1002, ¶169. 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/Ozer combination in which one AVS unit is used for each set of cores, a POSITA would understand that the supply voltages (“Vdd3” and “Vdd4”) for each set of cores would be set independently. EX1002, ¶169.

[1.2.3]

Elgebaly/Ozer renders obvious Limitation [1.2.3] for the reasons discussed for Limitations [1.1.1] and [1.2.1]. EX1002, ¶170. As also explained in Sections X.D.1-4, Elgebaly discloses or at least renders obvious a clock generator that includes multiple PLLs receiving respective independent input clock signals from different external crystal oscillators, with each PLL generating an output clock signal to respective sets of the processor cores. *Id.*; EX1004, [0028].

[1.3]

Elgebaly/Ozer renders obvious Limitation [1.3] at least under PO’s interpretation of this element as discussed for Limitation [1.3.1] in Ground 1. EX1002, ¶171. As explained in Section X.D.1.c, a POSITA would understand from Elgebaly’s disclosure that Elgebaly’s ASIC 120 necessarily includes an interface block configured to facilitate communication between cores and sets of cores. EX1002, ¶171. As explained in Section X.D.2, Ozer discloses cache coherency protocols for implementing multi-core processors that also rely on such an interface block (e.g., cache coherency fabric, snoop bus). Finally, as explained in

Sections X.D.3-4, it was obvious to combine Ozer's teachings with those of Elgebaly, rendering obvious at least the Elgebaly/Ozer Processor with an interface block (e.g., cache coherency fabric, snoop bus) coupled to the sets of processor cores. EX1002, ¶171.

[1.3.1]

Elgebaly/Ozer renders obvious Limitation [1.3.1] at least under PO's interpretation of this element as discussed for Limitation [1.3.1] in Ground 1. As explained in Sections X.D.3-4, it was obvious to combine Ozer's teachings with those of Elgebaly, rendering obvious at least the Elgebaly/Ozer Processor that included an interface block (e.g., cache coherency fabric, snoop bus) configured to facilitate communication between the sets of processor cores. EX1002, ¶172.

For example, Ozer explains that per "a typical cache coherency protocol, certain accesses performed by a processor will require a coherency operation to be performed." EX1013, [0007]; EX1002, ¶173. "The coherency operation will cause a notification ([aka] ... snoop request) to be sent to the other processors identifying the type of access taking place and the address ..., " which "will cause those other processors to perform certain actions ..., and may also in certain instances result in certain information being fed back from one or more of those processors to the processor initiating the access" EX1013, [0007]; EX1002, ¶173.

As a further example, Ozer explains that when a core wishes to read or write to a data location in memory, its cache coherency fabric 60 selectively circulates a cache coherency notification to “other caches [of respective processor cores] to cause those caches ... to determine whether any of those caches store data at the address specified by the access request.” EX1013, [0073]. Each processor then takes action “dependent on the cache coherency protocol being used.” *Id.* Under known write-invalidate schemes, a POSITA would have understood that the processor receiving the request would mark that data in its cache as invalid. EX1002, ¶174, *supra* Section VII.D. Under known write-update schemes, a POSITA would have understood that the modified data is sent on the bus from the processor that issued the cache coherency notification to any processors holding old copies of the data in their local caches. EX1002, ¶174.

The cache coherency fabric and bus in the Elgebaly/Ozer Processor thereby facilitate communication between the sets of cores by passing cache coherency operations and data between the sets of cores at least under PO’s interpretation of this element. EX1002, ¶175.

b. Claim 2

Elgebaly/Ozer renders obvious Limitation [2.1] for reasons discussed above in Limitation [1.3.1]. EX1002, ¶176. As discussed in Section VII.B, it was known that multi-core processors implementing multiple voltage domains required

mechanisms for translating logic levels between chip components operating at different voltages. EX1002, ¶176. A POSITA would have understood Figure 1 of Elgebaly, which depicts communication between processor cores operating using independent voltage levels, to require a voltage level shifter as part of the bus shown in the center of the ASIC. *Id.*; *see* EX1004, FIG. 1. Further, there were well-known prior art designs for such circuits. *See e.g.*, EX1031, FIGs. 4-5, 5:45-57, EX1039, 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 to the logic level of the second set must be referenced to the second supply voltage to know at what logic level the second set is operating. EX1002, ¶176. For these reasons, it would have been obvious, based on Elgebaly's disclosure, to implement a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set for cache coherency operations between the first and second sets (i.e., "a first signal traveling from the first set of processor cores to the second set of processor cores"). *Id.*

c. Claim 3

Elgebaly/Ozer renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1]. EX1002, ¶177.

d. Claim 5

Elgebaly/Ozer renders obvious Limitation [5.1]. EX1002, ¶178. Elgebaly discloses “other circuits 170” onboard the ASIC, which “may include a power control unit that controls power to various processing units.” EX1004, [0028]. Elgebaly further discloses that the [orange] “other circuits 170,” which includes the power control unit, may be located in a periphery of the ASIC:




FIG. 2
(MODIFIED)

Id., FIG. 2 (modified). A POSITA would have understood Figure 2 to disclose power control unit—which controls power consumption of the processor cores by sending signals to each core—to be in a periphery of the multi-core processor. EX1002, ¶178. Further, as explained in Ground 1 for Limitation [5.1], it was obvious to a POSITA to place control logic circuitry along the perimeter of the multi-

core processor. *Id.*, ¶¶103, 178; *see also* EX1033, FIG. 8, [0056], [0103]-[0107]. Thus, a POSITA would have understood Elgebaly to disclose—or at least render obvious—a control block that supplies control signals to each processing core being located in a periphery of the multi-core processor. EX1002, ¶178.

e. Claim 8

Elgebaly/Ozer renders obvious Limitation [8.1]. EX1002, ¶179. As explained in Section X.D.1-4, Elgebaly/Ozer discloses or renders obvious a first set of processor cores 1 and 3 sharing a voltage signal, clock signal, and AVS Unit 3; and a second set of processor cores 2 and 4 sharing a different voltage signal, clock signal, and AVS Unit 4. While Elgebaly and Ozer are silent on physical layout of cores, it was obvious to a POSITA to arrange processor cores 1 and 3 in a first row constituting a first region and processor cores 2 and 4 in a second row constituting a second non-overlapping region, for example, as shown in Modified Figure 2 of Elgebaly (inset below). EX1002, ¶179.




FIG. 2
(MODIFIED)

First, arranging processor cores in rows was a well-known approach to physical layout of a multicore processor. EX1002, ¶179 (citing EX1033, FIG. 8, EX1022, 2:8-10, FIGs. 1-2, EX1036, 4:53-57, FIG. 2, EX1011, FIG. 3) Second, arranging the first set in one row and the second set in a second row enables efficient signal routing to route the shared clock signal and voltage signal to each core in a set. EX1002, ¶179. Third, placing the cores of each set together and, thus, in non-overlapping regions enables the shared AVS Unit to be placed close to all cores in the set. *Id.*

f. Claim 10

Elgebaly/Ozer renders obvious Limitation [10.1] for reasons discussed in Limitation [8.1]. EX1002, ¶180.

g. Claim 11

Elgebaly/Ozer renders obvious Limitation [11.1] for reasons discussed in Limitation [8.1]. EX1002, ¶181.

h. Claim 14

Elgebaly/Ozer renders obvious Limitation [14.1]. EX1002, ¶182. Ozer discloses a “cache coherency fabric 60 [that] is located centrally” and “issue[s] snoop requests to any such identified caches identifying the segments to be subjected to the snoop operation.” EX1013, [0077], FIG. 1A. Ozer also discloses that “the cache coherency fabric 60 will ensure that when one processor updates data in its local cache, any copies of that data in the caches associated with the other processors will either be appropriately updated or invalidated in those caches to ensure coherency.” *Id.*, [0072]. As in Ozer’s Figure 1A, cache coherency fabric is located between, and interfaces with both of, the snoop bus and the shared memory. EX1013, FIG. 1A. A POSITA would have understood this to mean that Ozer discloses a control block located substantially central to the first set of processor cores and the second set of processor cores. EX1002, ¶182. As explained in Sections X.D.3-4, a POSITA would have been motivated to include Ozer’s cache coherency fabric as a control block within the multi-core processor to send control signals to each of the processor cores to ensure coherency between local caches at each core. *Id.*; *see also* EX1004, [0028]-[0028], [0048]; EX1013, [0004]-[0006],

[0072], FIG. 1A. Both Ozer and Elgebaly also teach that the interface between the on-chip bus and the internal memory is substantially central to the processor cores. EX1004, FIG. 1, [0028]; EX1013, [0077].

i. Claim 21

[21.pre]

Elgebaly/Ozer renders obvious Limitation [21.pre] for reasons discussed in Limitation [1.pre]. EX1002, ¶183.

[21.1]

Elgebaly/Ozer renders obvious Limitation [21.1] for reasons discussed in Limitation [1.1]. EX1002, ¶184.

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

Elgebaly/Ozer renders obvious Limitation [21.1.1] for reasons discussed in Limitation [1.1.1]. EX1002, ¶185.

[21.2]

Elgebaly/Ozer renders obvious Limitation [21.2] for reasons discussed in Limitation [1.2]. EX1002, ¶186.

[21.2.1]

Elgebaly/Ozer renders obvious Limitation [21.2.1] for reasons discussed in Limitation [1.2.1]. EX1002, ¶187.

[21.2.2]

Elgebaly/Ozer renders obvious Limitation [21.2.2] for reasons discussed in Limitation [1.2.2]. EX1002, ¶188.

[21.2.3]

Elgebaly/Ozer renders obvious Limitation [21.2.3] for reasons discussed in Limitation [1.2.3]. EX1002, ¶189.

[21.3]

Elgebaly/Ozer renders obvious Limitation [21.3] for reasons discussed in Limitation [1.3]. EX1002, ¶190.

[21.3.1]

Elgebaly/Ozer renders obvious Limitation [21.3.1] for reasons discussed in Limitation [1.3.1]. EX1002, ¶191.

E. Ground 5: Elgebaly, Ozer, and Vajda Render Obvious Claims 8-9

1. *Motivation to Combine Vajda with Elgebaly and Ozer*

Similar to the White/Talwar/Vajda combination in Ground 3, a POSITA would have been motivated to apply Vajda’s core layout guidance in defining the physical layout of cores in Elgabaly/Ozer to obtain Vajda’s stated benefit of a layout that may “avoid hot spots and over-heating.” EX1002, ¶192; EX1011, [0062]. As discussed in Section X.D.1, Elgebaly/Ozer includes a first set of processor cores 1 and 3 that share voltage and clock signals and a second set of processor cores 2 and 4 that share voltage and clock signals that are “set independently” from other voltage

and clock signals, thus enabling operating one set at a different higher frequency—producing more heat—than the other. Thus, the problem Vajda seeks to address where a layout includes “processor cores 231 that can execute at higher frequencies” and “other processor cores 232, running at lower frequencies” is applicable to Elgebaly/Ozer, and a POSITA would be motivated to apply Vajda’s layout guidance to “avoid hot spots and over-heating.” EX1002, ¶192; EX1011, [0062].

2. Modification of Elgebaly/Ozer in View of Vajda

As discussed in Section X.D.1, Elgebaly/Ozer includes a first set of processor cores 1 and 3 that share voltage and clock signals and a second set of processor cores 2 and 4 that share voltage and clock signals that are “set independently” from other voltage and clock signals. In accordance with Vajda’s disclosure 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” (EX1011, [0062]), a POSITA would be motivated to arrange the cores of the first set and the second set in a checkerboard pattern (e.g., a first row of processor core 1 followed by 2 and a second row of processor core 4 followed by 3) in order to “spread out” cores of one cluster from each other and avoid any hot spots created by grouping of cores of one cluster operating at a higher frequency. EX1002, ¶193.

3. Element-by-Element Analysis

a. Claim 8

Elgebaly/Ozer/Vajda renders obvious Limitation [8.1]. EX1002, ¶194. As discussed in Section X.E.2, Elgebaly/Ozer/Vajda includes processor cores 1 and 3 (the first set) and processor cores 2 and 4 (the second set) arranged in a checkerboard pattern (e.g., a first row of processor core 1 followed by 2 and a second row of processor core 4 followed by 3). The region containing the first set of processor cores is the first region and the region containing the second set of processor cores is the second region. EX1002, ¶194.

b. Claim 9

Elgebaly/Ozer/Vajda renders obvious Limitation [9.1] for reasons discussed in Limitation [8.1]. EX1002, ¶195. The first and second regions are an overlapping region because the checkerboard pattern of cores results in overlapping regions containing each set of cores. *Id.*

F. Grounds 6-7: Grounds 4-5 and Williams Render Obvious Claims 1-3, 5, 8-11, 14, and 21

To the extent Patent Owner contends that Elgebaly does not disclose using multiple external crystal oscillators, it would have been obvious to modify Elgebaly to include such functionality as disclosed in Williams. EX1002, ¶196.

Specifically, a POSITA would look to Williams to supplement Elgebaly's disclosure related to clock generation. EX1002, ¶197. Williams explains how to

design multi-core processor clock systems, such as Elgebaly's clock generator, a topic tangential to Elgebaly's disclosure. *Id.* It would have been obvious to a POSITA to combine Williams's disclosure of multiple crystal oscillators that generate independent clock signals for independent clock domains within a chip, with Elgebaly's chip architecture that uses multiple clock domains. *Id.*; *see* EX1004 [0028], FIG. 2. A POSITA would also have known that the use of multiple crystal oscillators is a simple, well-known method to generate multiple, independent clock signals to feed independent clock domains on a multi-core processor. EX1002, ¶197; *see* EX1027, [0005]. Such a design is in line with Elgebaly's teaching that the clock generator "may couple to external circuitry 142, which may include crystals, inductors, capacitors, and so on." *Id.*, [0028].

Relying on this teaching from Elgebaly and further in view of Williams, a POSITA would have implemented Elgebaly's clock generator with the multiple oscillator structure disclosed in Figure 1 of Williams. EX1002, ¶198; *see* EX1008, Figure 1. Such a modification would result in the same structure as discussed in Grounds 4-5, above. Accordingly, Grounds 6-7 render the Challenged Claims obvious for the same reasons as discussed in Grounds 4-5, respectively.

XI. CONCLUSION

Petitioner respectfully requests cancellation of the Challenged Claims.

Date: January 22, 2025

By: /Eagle H. Robinson/
Eagle H. Robinson
Reg. No. 61,361

Lead Counsel for Petitioners

CERTIFICATION UNDER 37 C.F.R. §42.24

Pursuant to 37 C.F.R. § 42.24(a)(i), the undersigned certifies that this Paper—exclusive of the table of contents, mandatory notices under § 42.8, certificate of service, and this certificate of word count—includes 13,838 words.

Date: January 22, 2025

/Eagle H. Robinson/
Eagle H. Robinson (Reg. No. 61,361)

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. § 42.6(e) and 37 C.F.R. § 42.105(a), the undersigned certifies that on January 22, 2025, complete copies of this Petition for Inter Partes Review, Petitioners' Powers-of-Attorney, and EX1001-EX1008, EX1010-EX1011, and EX1013-EX1040 were served on Patent Owner at the correspondence addresses of record listed below by USPS Priority Mail Express®:

GENIE I. SU
REN-SHENG INTERNATIONAL IP MANAGEMENT LTD
7F, NO. 57, SEC. 2, DUN HUA S. ROAD
TAIPEI, 106 TAIWAN
TELEPHONE: (408) 729-9496

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

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

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

/Eagle H. Robinson/
Eagle H. Robinson (Reg. No. 61,361)