

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

TCL ELECTRONICS HOLDINGS LTD. And TCL INDUSTRIES HOLDINGS
CO., LTD.,
Petitioner,

v.

INTELLECTUAL VENTURES I LLC,
Patent Owner.

Case IPR2024-01245

U.S. PATENT NO. 7,089,443

PETITION FOR *INTER PARTES* REVIEW

TABLE OF CONTENTS

	<u>Page</u>
1. INTRODUCTION	1
2. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED	2
A. 37 C.F.R. §42.104(b)(1): Claims for Which IPR Is Requested	2
B. 37 C.F.R. §42.104(b)(2): Grounds for Challenge	2
C. 37 C.F.R. § 42.104(b)(3): Claim Construction	5
D. 37 C.F.R. § 42.104(b)(4): How the Claims Are Unpatentable	6
E. 37 C.F.R. § 42.104(b)(5): Evidence Supporting Challenge.....	6
3. OVERVIEW OF THE '443 PATENT	6
A. Purported Invention.....	6
B. Relevant Prosecution History.....	9
C. Priority Date	11
4. PERSON OF ORDINARY SKILL IN THE ART.....	11
5. SUMMARY OF THE PRIOR ART	11
A. Iyer (EX1005).....	11
B. Soerensen (EX1006).....	12
C. Hemani (EX1007)	13
D. Moore (EX1008)	13
E. Njolstad (EX1009)	14
F. Liljeberg (EX1010)	14
6. SPECIFIC GROUNDS FOR PETITION	15

A.	Ground I: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Iyer (EX1005)	15
1.	Claim 1	15
2.	Claim 3	19
3.	Claim 4	19
4.	Claim 5	24
5.	Claim 8	25
6.	Claim 10	27
7.	Claim 12	28
B.	Ground II.A: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Njolstad (EX1009)	28
1.	Claim 1	28
2.	Claim 3	32
3.	Claim 4	33
4.	Claim 5	35
5.	Claim 8	37
6.	Claim 10	38
7.	Claim 12	38
C.	Ground II.B: Claims 5 and 10 Are Rendered Obvious By Njolstad (EX1009) in view of Liljeberg (EX1010)	39
1.	Motivation to Combine Njolstad and Liljeberg	39
2.	Claim 5	40
3.	Claim 10	42
D.	Ground III.A: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Soerensen	43

Petition for *Inter Partes* Review of U.S. Patent No. 7,089,443

1.	Claim 1	43
2.	Claim 3	48
3.	Claim 4	48
4.	Claim 5	50
5.	Claim 8	51
6.	Claim 10	54
7.	Claim 12	54
E.	Ground III.B: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Soerensen (EX1006) in view of Hemani (EX1007)	55
1.	Motivation to Combine Soerensen and Hemani	55
2.	Claim 1	56
3.	Claim 3	61
4.	Claim 4	62
5.	Claim 5	64
6.	Claim 8	66
7.	Claim 10	69
8.	Claim 12	70
F.	Ground III.C: Claims 5 and 10 Are Rendered Obvious By Soerensen (EX1006) in view of Moore (EX1008)	70
1.	Motivation to Combine Soerensen and Moore	70
2.	Claim 5	71
3.	Claim 10	73
7.	SECONDARY CONSIDERATIONS	73
8.	NO JUSTIFICATION FOR DISCRETIONARY DENIAL.....	73

Petition for *Inter Partes* Review of U.S. Patent No. 7,089,443

A.	The Board Should Not Decline Institution Pursuant to Fintiv.....	73
B.	The Board Should Not Decline Institution Pursuant to General Plastic	76
C.	The Board Should Not Decline Institution Pursuant To § 325(d)	76
9.	MANDATORY NOTICES UNDER 37 C.F.R. § 42.8(A)(1)	77
A.	37 C.F.R. § 42.8(b)(1): Real Party-in-Interest	77
B.	37 C.F.R. § 42.8(b)(2): Related Matters	77
C.	37 C.F.R. § 42.8(b)(3) and (4): Counsel and Service Information.....	78
10.	PAYMENT OF FEES UNDER 37 C.F.R. §42.103.....	78
11.	CERTIFICATION OF GROUNDS FOR STANDING	79
12.	CONCLUSION	79

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>Microsoft Corp. v. Parallel Networks Licensing, LLC</i> , IPR2015-00486, Paper 10 (P.T.A.B. July 15, 2015).....	76
<i>Sand Revolution II, LLC v. Continental Intermodal Group—Trucking LLC</i> , IPR2019-01393, Paper 24 (P.T.A.B. June 16, 2020)	74
Statutes	
35 U.S.C. §102.....	3
35 U.S.C. §103.....	3
35 U.S.C. § 314(a)	72
35 U.S.C. §315(e)(1).....	78
35 U.S.C. § 318(b)	2
Other Authorities	
37 C.F.R. § 1.68	6
37 C.F.R. § 42.8(A)(1).....	76
37 C.F.R. § 42.8(b)(1).....	76
37 C.F.R. § 42.8(b)(2).....	76
37 C.F.R. § 42.8(b)(3).....	77
37 C.F.R. § 42.8(b)(4).....	77
37 C.F.R. §42.10(b)	77
37 C.F.R. §42.103	77
37 C.F.R. §42.103	77

Petition for *Inter Partes* Review of U.S. Patent No. 7,089,443

37 C.F.R. §42.104(b)(1).....	2
37 C.F.R. §42.104(b)(2).....	2
37 C.F.R. § 42.104(b)(3).....	5
37 C.F.R. § 42.104(b)(4).....	6
37 C.F.R. § 42.104(b)(5).....	6

TABLE OF EXHIBITS

Exhibit No.	Description
1001	U.S. Patent No. 7,089,443 (the “’443 Patent”)
1002	Certified File History of the ’443 Patent
1003	Declaration of Dr. Nader Bagherzadeh
1004	Declaration of Dr. Sylvia Hall-Ellis
1005	Iyer, Anoop, and Diana Marculescu. “Power efficiency of voltage scaling in multiple clock, multiple voltage cores.” <i>Proceedings of the 2002 IEEE/ACM international conference on Computer-aided design</i> (2002).
1006	U.S. Patent App. No. 20030071657 (“Soerensen”)
1007	A. Hemani, T. Meincke, S. Kumar, A. Postula, T. Olsson, P. Nilsson, J. Oberg, P. Ellervee, and D. Lundqvist, “Lowering power consumption in clock by using globally asynchronous, locally synchronous design style,” <i>Proc. ACM/IEEE Design Automation Conference</i> , 1999.
1008	Moore, Simon, George Taylor, Bob Mullins, and Peter Robinson. “Channel communication between independent clock domains.” <i>First ACiD-WG Workshop of the European Commission’s Fifth Framework Programme</i> . 2001.
1009	Tormod Njolstad, Odd Tjore, Kjetil Svarstad, Lars Lundheim, Tor O. Vedal, Jukka Typo, Tor Ramstad, Lars Wanhammar, E. J. Aar, and Helge Danielsen. “A socket interface for GALS using locally dynamic voltage scaling for rate-adaptive energy saving.” <i>Proceedings 14th Annual IEEE International ASIC/SOC Conference (IEEE Cat. No. 01TH8558)</i> . IEEE, 2001.
1010	Liljeberg, Pasi, Juha Plosila, and Jouni Isoaho. “Asynchronous interface for locally clocked modules in ULSI systems.” <i>ISCAS 2001. The 2001 IEEE International Symposium on Circuits and Systems (Cat. No. 01CH37196)</i> , IEEE, 2001.
1011	<i>Curriculum Vitae</i> of Dr. Nader Bagherzadeh
1012	D. M. Chapiro, <i>Globally Asynchronous Locally Synchronous Systems</i> , Ph.D. thesis, Stanford University, 1984.

Exhibit No.	Description
1013	Hsu, Chung-Hsing, Ulrich Kremer, and Michael Hsiao. “Compiler-directed dynamic voltage/frequency scheduling for energy reduction in microprocessors.” <i>Proceedings of the 2001 international symposium on low power electronics and design</i> , 2001.
1014	I. E. Sutherland, “Micropipelines,” <i>Communications of the ACM</i> , June 1989.
1015	Declaration of Sean O’Brien, University of Texas Libraries
1016	WIPO WO 2004/066092
1017	International Search Report App. No. PCT/UIS04/01700
1018	U.S. Patent No. 6,487,675
1019	U.S. Patent No. 5,309,561
1020	U.S. Patent App. No. 2004/0025069
1021	U.S. Patent App. No. 2004/0017234
1022	U.S. Provisional App. No. 60/315,655

Petitioner TCL Electronics Holdings Ltd. and TCL Industries Holdings Co., (“TCL”) requests *inter partes* review (“IPR”) of claims 1, 3, 4, 5, 8, 10, and 12 (the “Challenged Claims”) of U.S. Patent No. 7,089,443 (EX1001, “the ’443 Patent”).

1. INTRODUCTION

Globally-asynchronous, locally-synchronous (“GALS”) systems contain several independent synchronous blocks which operate with their own local clocks and communicate asynchronously with each other. *See generally* EX1005; EX1012. The ’443 Patent models GALS clocking architecture in a microprocessor. However, GALS clocking systems were known well before the ’443 patent’s earliest priority date and the application of GALS to microprocessors had already taken place. *Id.*

In addition to the known upside of applying GALS to microprocessors, it was known to use dynamic voltage scaling in order to improve power efficiency. *See e.g.*, EX1009, §I. Thus, all aspects of the claimed microprocessor and method of operating the microprocessor were well known and routine in the art.

The Office issued the ’443 Patent in haste. Although prior art was disclosed in an IDS describing microprocessors having multiple independently clocked domains in a GALS architecture, the Examiner allowed the ’443 claims without entering a single prior art rejection. The allowance was in error because

the Examiner failed to appreciate that the prior art cited in the IDS and in a related PCT search report disclose the very feature that the Examiner believed to be missing from the prior art, as described in the '443 Notice of Allowance.

The prior art applied below describes that feature—microprocessors having multiple independently clocked domains—and other features the Examiner believed to be missing from the prior art (e.g., clock frequency being dynamically changeable independently of the frequencies of the other clock signals). Accordingly, Petitioner respectfully requests that *inter partes* review of the '443 Patent be instituted, and the Challenged Claims be cancelled as unpatentable pursuant to 35 U.S.C. § 318(b). This Petition demonstrates that there is a reasonable likelihood that at least one of the Challenged Claims is unpatentable as obvious.

2. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED

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

Petitioner challenges the patentability of the Challenged Claims of the '443 Patent and requests that they be canceled.

B. 37 C.F.R. §42.104(b)(2): Grounds for Challenge

The claims are challenged based on the following references, none of which were considered by the Examiner during prosecution. The '443 Patent is

subject to pre-AIA law,¹ as the patent was filed on January 23, 2004, and the earliest priority application was filed on January 23, 2003.

- **Iyer** (EX1005) is a conference paper from the IEEE/ACM International Conference on Computer-Aided Design, held November 10-14, 2002. EX1004, ¶¶44-51. It was publicly available no later than November 10, 2002. *Id.* Accordingly, Iyer qualifies as prior art under at least §102(a)-(b).

- **Soerensen** (EX1006) is a United States Patent Application filed on August 29, 2002, and claiming priority to Provisional Application 60/315,655 (EX1022), filed on August 29, 2001.² Soerensen therefore qualifies as prior art under at least §102(e).

¹ All citations in this Petition to 35 U.S.C. §§102 and 103 are to the pre-AIA law.

² At least Soerensen's Claim 1 and all of Soerensen's disclosures relied upon in this petition find support in the Provisional Application. *See* EX1003, ¶68; EX1006 at 8, Cl. 1; EX1022 at p. 110 ("one of the most effective ways of minimizing dynamic power consumption is to reduce the applied voltage to the minimum value required for the CMOS design to perform the required function(s) at any point in time"); *id.* ("it is possible to set the voltage (dynamically) to the minimum value required to allow the device to operate at that clock frequency, achieving minimal

- **Hemani** (EX1007) is a conference paper from the 36th Design Automation Conference (DAC '99) held on June 21-25, 1999, in New Orleans, Louisiana. It was publicly available no later than June 1, 1999. EX1004, ¶¶76-83. Accordingly, Hemani qualifies as prior art under at least §102(a)-(b).

- **Moore** (EX1008) is a conference paper from the Fifth ACiD-WG Workshop, held in Nauchâtel (Switzerland) in 2001. EX1004, ¶¶84-88. It was publicly available no later than August 16, 2002. *Id.* Accordingly, Moore qualifies as prior art under at least §102(a)-(b).

- **Njolstad** (EX1009) is a conference paper from the 14th Annual IEEE International ASIC/SOC Conference. EX1004, ¶¶68-75. It was publicly available no later than September 12, 2001. *Id.* Accordingly, Njolstad qualifies as prior art under at least §102(a)-(b).

- **Liljeberg** (EX1010) is a conference paper from the IEEE International Symposium on Circuits and Systems held May 6-9, 2001. EX1004, ¶¶52-59. It was

dynamic power consumption”); *id.* at p. 247 (“the baseband may have six active clock domains”); *id.* at p. 170 (“frequencies can be changed dynamically . . . in conjunction with dynamically varying the internal logic voltage levels, to optimize power consumption”); *id.* at p. 169 (showing a parallel schematic).

publicly available no later than May 1, 2001. *Id.* Accordingly, Liljeberg qualifies as prior art under at least §102(a)-(b).

Petitioner requests cancellation of the Challenged Claims on the following grounds:

Ground	Basis	Claims	Proposed Statutory Rejection
I	§ 103	1, 3, 4, 5, 8, 10, 12	Obvious over Iyer
II.A	§ 103	1, 3, 4, 5, 8, 10, 12	Obvious over Njolstad
II.B	§ 103	5, 10	Obvious over Njolstad in view of Liljeberg
III.A	§ 103	1, 3, 4, 5, 8, 10, 12	Obvious over Soerensen
III.B	§ 103	1, 3, 4, 5, 8, 10, 12	Obvious over Soerensen in view of Hemani
III.C	§ 103	5, 10	Obvious over Soerensen in view of Moore

C. 37 C.F.R. § 42.104(b)(3): Claim Construction

In the parallel litigation between Petitioner and Patent Owner, Petitioner has contended that in the context of Claim 8(d), which reads “operating the microprocessor such that each domain operates synchronously, while the domains operate asynchronously relative to one another”, the term “synchronously” means “using the same clock with a fixed phase relationship”; and “asynchronously” means “using independent clocks with no known phase

relationship.” Patent Owner disputes this construction in the District Court. The prior art satisfies both parties’ constructions of these terms.

D. 37 C.F.R. § 42.104(b)(4): How the Claims Are Unpatentable

A detailed explanation of how the Challenged Claims are unpatentable is provided below in Section 6.

E. 37 C.F.R. § 42.104(b)(5): Evidence Supporting Challenge

A list of exhibits is provided at the beginning of this Petition. The relevance of this evidence and the specific portions supporting this challenge are provided below in Section 6. Petitioner submits the declaration of Dr. Nader Bagherzadeh (EX1003) in support of this Petition. 37 C.F.R. § 1.68.

3. OVERVIEW OF THE ’443 PATENT

A. Purported Invention

The ’443 Patent is directed to a “straightforward extension” of the prior art—a microprocessor design utilizing multiple clock domain (MCD) microarchitecture and a globally asynchronous, locally synchronous (GALS) clocking style. EX1001 at 4:52-53. Each functional block within this microprocessor operates with its own clock frequency, and queue structures manage communication between domains.

As the '443 Patent acknowledges, the concept of using separate clock domains to address clock distribution and wire delay issues was well understood in the field. EX1001 at 1:31-37; EX1003, ¶62. The concept of GALS was also widely understood to date back to at least Chapiro's 1984 PhD thesis. *See, e.g.* EX1005 at §1.1 ("GALS systems have been studied in detail by Chapiro in his 1984 thesis"); EX1012 (Chapiro). At the time of the '443 Patent's prosecution, this concept had already been implemented in mass production. EX1001 at 1:55-58 ("several microprocessor systems run the memory bus off a different clock than the processor core").

The '443 Patent also claims the use of queues for interdomain synchronization. EX1001, Cl. 5, 10. Queues decouple producer and consumer domains by allowing the producer to continue writing even if the consumer is not yet ready to receive. EX1003, ¶115. The data simply goes into the queue until the consumer is ready to read it. *Id.* A basic iteration is the first-in-first-out (FIFO) queue. *Id.*, ¶52. While there are other means for interdomain synchronization, queues are straightforward and effective because they reduce latency and optimize throughput. *Id.*, ¶124. Predictably, queues were well known in the art. For instance, a 1989 paper titled *Micropipelines*, which the Iyer reference calls "the classic introduction to asynchronous design," illustrates a FIFO queue approach in detail. EX1005, §1.1; Ex. 1014 at 722 ("FIFOs often connect senders and receivers that

have separate clocks.”). The ’443 Patent itself acknowledges this fact, noting that “[o]ur design uses *existing* queue structures” to facilitate interdomain communication. EX1001 at 5:2-3. Thus, queues cannot comprise the inventive concept of the ’443 Patent.

The ’443 Patent also proposes the use of dynamic voltage scaling in each domain in order to lower energy consumption. However, this concept is well worn within the ’443 Patent’s own references. *See, e.g.*, EX1001 at 3:14; EX1013 at §1.2 (“In the context of dynamic voltage scaled microprocessors, voltage scheduling is a problem that assigns appropriate clock speeds to a set of tasks, and adjusts the voltage accordingly”). And to the extent that the controlling algorithm is the purported inventive concept, this too is disclosed by the ’443 Patent’s references. *See id.* (“Researchers have proposed many ways of determining "appropriate" clock speeds through on-line and off-line algorithms”). Finally, the ’443 Patent models several assumption-based outputs relating to performance overhead and energy-delay product, but these models are not required by the ’443 Patent’s Claims. *See generally* EX1001.

In essence, the ’443 Patent’s purported invention is at best a combination of known techniques of (a) clock domain separation and (b) dynamic voltage and frequency scaling, in the context of microprocessor design. The

incremental nature of these improvements and the reliance on well-established principles suggest that the patent does not present a significant advancement over the prior art. A person of ordinary skill in the art would have recognized that the purported benefits of the MCD microarchitecture were achievable through existing methods, and thus would question the purported innovation of the '443 Patent.

B. Relevant Prosecution History

The '443 Patent was filed on January 23, 2004 as Application No. 10/762,550, but claims priority to U.S. Provisional Application 60/441,759 ("759 provisional") filed on January 23, 2003. WO 2004/066092 ("PCT application"), EX1016, also claims priority to the '759 provisional application and too was filed on January 23, 2004.

An International Search Report, EX1017, ("search report") was entered on December 6, 2004 in the PCT application. The search report identified two anticipatory references: U.S. Patent No. 6,487,675 to Sager, EX1018, and U.S. Patent No. 5,309,561 to Overhouse, EX1019, which teach microprocessor domains. Sager expressly teaches a microprocessor having two or more "domains" clocked at different frequencies (e.g., at 4:3-20) and, although it does not use the term "domain," Overhouse discloses a microprocessor having fast and slow sections clocked by signals having different frequencies (e.g., at 4:27-32). A

POSITA would have understood the two microprocessor sections in Overhouse to be microprocessor domains. EX1003, ¶65.

A few months after the search report was entered in the PCT application, an IDS was filed during prosecution of the '443 patent disclosing the references from the search report, but not the underlying search report. *See* EX1002 (March 4, 2005 IDS). Surprisingly, in contrast to the search report which identified multiple anticipatory references, the '443 Patent was subject to a first action allowance. *See* May 19, 2005 Notice of Allowance ("NOA"). EX1002. The NOA identifies the closest prior art as (i) Gary (PGPub 2004/0025069), EX1020, and (ii) Tam (PGPub 2004/0017234), EX1021, but notes that neither Gary nor Tam describes "domains in a microprocessor." *See* EX1002. That this perceived deficiency was sufficient to conclude that the '443 claims were allowable is surprising. At least Sager and Overhouse teach microprocessors having multiple domains clocked by signals having different frequencies.

Like Sager and Overhouse, the prior art applied in each Ground presented below describes microprocessor domains. In addition to describing microprocessor domains, the prior art applied in the Grounds describe dynamically controlling clock frequency and voltage in each domain separately.

No other portion of the prosecution history is relevant for purposes of this Petition.

C. Priority Date

Petitioner uses January 23, 2003 as the priority date for the '443 Patent, but reserves the right to challenge the provisional filing date in this IPR or any other proceeding where the '443 priority date is raised. The '443 Patent is subject to pre-AIA law.

4. PERSON OF ORDINARY SKILL IN THE ART

A person of ordinary skill in the art ("POSITA") at the time of the '443 Patent would have at least a Master's degree in electrical engineering, computer engineering, computer science, or a related field, with at least two or more years of experience in the field of processor design, with an allowance for additional education or experience that might substitute for those requirements. EX1003, ¶38.

5. SUMMARY OF THE PRIOR ART

A. Iyer (EX1005)

Iyer is an Institute of Electrical and Electronics Engineers (IEEE) conference paper titled "Power Efficiency of Voltage Scaling in Multiple Clock, Multiple Voltage Cores." See EX1005 at Cover. Iyer discusses how "[b]y exploiting the flexibility of independent dynamic voltage scaling the various clock

domains, the power efficiency of GALS designs can be improved.” *Id.* Iyer outlines the design and modeling of a globally asynchronous locally synchronous (GALS) microprocessor where each domain is clocked independently.

Iyer is analogous prior art, including because (a) it is in the same field as the ’443 Patent, and (b) is directed to the same problem as the ’443 Patent in improving communication and power efficiency in a microprocessor. *See id.*; EX1003, ¶67.

B. Soerensen (EX1006)

Soerensen is a United States Patent Application Publication titled “Dynamic Voltage Control Method and Apparatus.” *See* EX1006 at Cover. Soerensen discusses using a dynamic power controller which identifies a clock frequency requirement of a processor and determines a voltage requirement to support the clock frequency requirement in each independent domain. The control logic transitions the processor to a power state defined by the clock frequency requirement and the voltage requirement.

Soerensen is analogous prior art because (a) it is in the same field as the ’443 Patent, and (b) it addresses the same problem as the ’443 Patent in improving communication and power efficiency in an on-chip processor. *See id.*; EX1003, ¶69.

C. Hemani (EX1007)

Hemani is a paper titled “Lowering Power Consumption in Clock by Using Globally Asynchronous Locally Synchronous Design Style.” EX1007 at Cover. Moore discusses partitioning a microprocessor into a plurality of individually clocked domains which are internally synchronous but asynchronous to one another. Hemani is focused on reducing power consumption and also disclosed means for inter-domain communication.

Hemani is analogous prior art because (a) it is in the same field as the ’443 Patent, and (b) it addresses the same problem as the ’443 Patent in improving communication and power efficiency in a microprocessor. *See id.*; EX1003, ¶71.

D. Moore (EX1008)

Moore is a paper titled “Channel Communication Between Independent Clock Domains.” EX1008 at Cover. Moore discusses reliable low-latency channel communication between clock domains using an asynchronous interconnect and queues for interdomain communications. Like Iyer, these techniques are implemented in a GALS (Globally Asynchronous Locally Synchronous) system, where each clock domain operates independently.

Moore is analogous prior art because (a) it is in the same field as the ’443 Patent, and (b) it addresses the same problem as the ’443 Patent in improving

communication and power efficiency in independently clocked domains of a GALS microprocessor. *See id.*; EX1003 ¶73.

E. Njolstad (EX1009)

Njolstad is a paper titled “A Socket Interface for GALS Using Locally Dynamic Voltage Scaling for Rate-Adaptive Energy Saving.” *See* EX1009 at Cover. Njolstad discusses a low-power amendment to the GALS paradigm using locally dynamic voltage scaling (LDVS) adapted to the actual processing rate requirements for each module. What Njolstad terms GALS-LDVS is a microprocessor-focused design with a plurality of independent clock domains which communicate asynchronously. *See id.* Each domain’s frequency and voltage is scalable in order to facilitate maximal power efficiency. *See id.*

Njolstad is analogous prior art because (a) it is in the same field as the ’443 Patent, and (b) it addresses the same problem as the ’443 Patent in improving communication and power efficiency in a GALS microprocessor. *See id.*; EX1003 ¶75.

F. Liljeberg (EX1010)

Liljeberg is a paper titled “Asynchronous Interface for Locally Clocked Modules in ULSI Systems.” *See* EX1010 at Cover. Liljeberg discusses the construction of globally asynchronous locally synchronous (GALS) systems using

an architecture that employs interface units to manage communication between locally clocked modules in an ultra large-scale integration (ULSI)—i.e., microprocessor—system. *See generally* EX1010. Liljeberg emphasizes the use of dual-port RAM as a FIFO queue for data transfer between asynchronous and synchronous domains, allowing each module to use an independent local clock. *Id.*

Liljeberg is analogous prior art because (a) it is in the same field as the '443 Patent, and (b) it addresses the same problem as the '443 Patent by facilitating communication and power efficiency in a microprocessor. *See id.*; EX1003, ¶77.

6. SPECIFIC GROUNDS FOR PETITION

A. Ground I: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Iyer (EX1005)

1. Claim 1

Iyer renders obvious Claim 1. *See* EX1003, ¶78.

a. Limitation 1[pre]


“A multiple clock domain microprocessor comprising:”

To the extent the preamble is limiting, Iyer teaches this limitation. EX1003, ¶79. For example, Iyer discloses the “design [of] a microprocessor in a GALS style with multiple clock domains.” EX1005 at §1.

b. Limitation 1[a]

“a plurality of domains;”

Iyer teaches this limitation. EX1003, ¶80. Iyer teaches a globally asynchronous, locally synchronous (GALS) design. EX1005 at §1. A GALS design entails “the use of several distinct local clocks (or clock domains)”. *Id.* Iyer models, for instance, “five clock domains”—in other words, a plurality of domains. EX1005 at §4.2. This plurality of domains can be seen, e.g., denoted with dotted lines on the right side of Iyer’s Figure 1:



EX1005 at 4, Fig. 1.

c. Limitation 1[b]

- i. ***“for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain ...”***

Iyer teaches this limitation. EX1003, ¶81. For instance, Iyer teaches “several independent synchronous blocks which operate with their own local clocks.” EX1005 at §1; *see also id.* at §3.2 (“various regions are clocked using different clock signals independent of each other”); *id.* at §3 (“Local clock generation: Each clock domain in a GALS system needs its own local clock generator; ring oscillators have been proposed as a viable clock generation scheme”); *id.* §1 “clocks (or clock domains) [] running at different frequencies”). In sum, for each of Iyer’s plurality of domains—alternately referred to as *synchronous blocks*—the clock is separately generated.

- ii. ***“... the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains; and”***

Iyer teaches this limitation. EX1003, ¶82. As discussed with respect to limitation 1[b](i), the frequencies in each of Iyer’s domains are independent of the other domains. Iyer further discloses that each domain is dynamically changeable. For instance, Iyer teaches that “the speed and voltage of each synchronous block are adapted dynamically.” EX1005 at §3.4. Iyer teaches both slowing down and increasing each domain’s frequency independently. *Compare* EX1005 at §3.4

("[w]e can slow down synchronous blocks that are off the critical path of the design while keeping the others running at full speed") *with id.* ("if the occupancy rises above a certain threshold . . . we switch the clock back to full speed").

d. Limitation 1[c]

"for each of the plurality of domains, a voltage input for receiving a voltage which is dynamically changeable independently of the voltages applied to said others of the plurality of domains."

Iyer teaches this limitation. EX1003, ¶83. As a general matter, Iyer notes that "dynamic voltage scaling (DVS) has been widely studied and implemented commercially." EX1005 at §1.1. Iyer specifically teaches "dynamically variable voltages for *each* synchronous block". EX1005 at §1 (emphasis added). In an entire section dedicated to "Dynamic Clock and Voltage Management" Iyer further discloses that "slower clock domains" can "operate at a smaller supply voltage," plainly teaching that voltage is dynamically changeable on an independent—that is, domain by domain—basis. EX1005 at §3.4. Dynamically changing both frequency *and voltage* for each domain independently is central to Iyer's teachings. *See, e.g.*, EX1005 at §2.3 (disclosing the ability to "dynamically adapt the speed and corresponding voltage for each synchronous block").

2. Claim 3

“The microprocessor of claim 1, wherein there are at least four of said domains.”

Iyer teaches this limitation.³ EX1003, ¶84. Iyer teaches, for instance, “we have chosen to have five clock domains in the GALS design.” EX1005 at §3.2. Iyer also notes, for example, that the 21264 Alpha processor has six clock domains. *Id* at §3.1. Ultimately, as a POSITA would have understood and as Iyer expresses, the number of domains is at the discretion of the designer, and thus this claim is disclosed and, in any event, does not add any non-obvious elements. *Id.*; EX1003, ¶84.

3. Claim 4

a. Limitation 4[a]

“The microprocessor of claim 1, wherein the microprocessor is programmed to determine a slack in processing in one of the domains and . . .”


Iyer teaches this limitation. EX1003, ¶85. As a threshold matter, Iyer’s core invention is “an adaptive approach for GALS processors which is able

³ For each dependent claim of each ground, it should be understood that the teachings of the underlying claim(s) on which the claim depends is incorporated by reference, and the Petition’s analysis focuses for sake of brevity only on the additional limitation(s) added in the dependent claim.

to tune the operating voltage of different clock domains, *according to the usage* of that clock domain.” EX1005 at §2.4. In other words, Iyer is principally focused on reducing underutilization, or “slack.” EX1003, ¶85.

In a GALS system, slack may be determined by monitoring the occupancy of each issue queue, where lower occupancy indicates the presence of slack. EX1003, ¶86. Iyer expressly teaches determining slack in this way, disclosing that “we monitor the occupancy of each issue queue and we reduce the frequency to a lower value only when the queue occupancy drops below a low threshold for *N* consecutive cycles.” EX1005 at §3.4.

Claim 4 only claims the determination of slack in “one of the domains.” EX1001 at 17:27-31. Considering, for example, one domain in Iyer—the Iyer reference teaches an integer issue queue as part of Domain 3:




Integer issue queue thresholds	7/9
FP issue queue thresholds	3/5
Memory issue queue thresholds	5/8
Threshold for clock pausing	30% of the largest clock period
Aggressive mode	FP 2.4 slowdown Mem 1.2 slowdown
Moderate mode	Int 1.8 slowdown FP 1.2 slowdown Mem 1.2 slowdown

Table 2. Microarchitecture organization

EX1005 at Fig. 1 (cropped), Table 2 (cropped and annotated).

Iyer then teaches monitoring occupancy in each issue queue—including this Domain 3—to determine slack:



EX1005, Fig. 3 (annotated).

Monitoring slack is a threshold prerequisite to the operation described in Figure 3. Specifically, low occupancy must be determined in order to reduce the domain's clock frequency to low speed. EX1003, ¶89. Thus, Iyer plainly discloses a microprocessor programmed to determine slack in one of the domains.


b. Limitation 4[b]

“ . . . to reduce the clock frequency and the voltage in said one of the domains to reduce the slack.”

Iyer teaches this limitation. EX1003, ¶¶90-91. As noted above with respect to limitation 4[a], Iyer teaches monitoring slack in order to determine whether a reduction in clock frequency is appropriate. And as Iyer notes and a

POSITA would have known, “slower clock domains” can “operate at a smaller supply voltage.” EX1005 at §3.4. For this reason, Iyer’s “method for dynamic clock management involves changing the operating *frequencies (and supply voltages)*.” EX1005 at §3.4 (emphasis added); *see also id.* at §2.3 (“synchronous blocks whose speed may be gracefully scaled down, can also run at a lower voltage, thus providing additional power savings”).

Additionally, while Claim 4 appears to separately claim *frequency* reduction and *voltage* reduction, these are in fact intertwined. Iyer gives the formula $f_i \propto \frac{(V_{dd,i} - V_t)^\alpha}{V_{dd,i}}$ where the voltage and clock speed associated with each clock domain i are $V_{dd,i}$ and f_i , respectively, the threshold voltage is V_t , and α (sensitivity) is “is a technology-dependent factor, which is 1.2 to 1.6 for current technologies.” *Id.* EX1005 at §2.4. In the simplest iteration, Iyer discloses reducing clock frequency from full speed f_{high} to low speed f_{low} in response to low occupancy (slack):



EX1005 at Fig. 3.


The relationship at full speed is $f_{high} \propto \frac{(V_{dd,high} - V_t)^\alpha}{V_{dd,high}}$. At low speed, the relationship is $f_{low} \propto \frac{(V_{dd,low} - V_t)^\alpha}{V_{dd,low}}$. Given that $f_{low} < f_{high}$ it follows that $\frac{(V_{dd,low} - V_t)^\alpha}{V_{dd,low}} < \frac{(V_{dd,high} - V_t)^\alpha}{V_{dd,high}}$. Therefore, $V_{dd,low} < V_{dd,high}$, in other words, a reduction in frequency ultimately implicates some reduction in voltage. EX1003, ¶92. A POSITA would also have understood generally that in the context of dynamic voltage and frequency scaling $V \propto \sqrt{f}$. *Id.* Finally, a POSITA would have understood that the purpose of Iyer's invention is to enable—wherever possible—a reduction in a domain's frequency, and consequentially that domain's voltage. *See id.* Therefore, voltage will be reduced when frequency is reduced. *Id.* Conversely, a POSITA would also have known that frequency cannot be substantially increased without a corresponding increase in supply voltage, or the signal will deteriorate.

Id. Thus, Iyer implicitly and expressly teaches scaling both voltage and frequency in response to slack.

4. Claim 5

“The microprocessor of claim 1, further comprising a queue for communication between at least two of the domains.”

Iyer teaches this limitation. EX1003, ¶93. This limitation only claims one queue between at least two domains, and Iyer teaches, for instance, a queue between domains 2 and 5:



EX1005 at Fig. 1 (cropped).

Iyer discloses several “issue queues in the design” including “one for loads and stores (clock domain 5).” EX1005 at §3.2. Iyer also teaches other integer queues, including “one for integer instructions (clock domain 3)” and “one for

floating-point instructions (clock domain 4).” *Id.* Iyer further teaches that their “approach uses asynchronous FIFOs [] between clock domains.” EX1005 at §3. As a POSITA would have understood, a first in first out system (FIFO) is a well-known type of queue. EX1003, ¶94.

5. Claim 8

a. Limitation 8[pre]

“A method of operating a microprocessor, the method comprising:”

As discussed with respect to limitation 1[pre], Iyer teaches a microprocessor. §VI.A.1. Iyer further discloses a “method” of operating said microprocessor. *See, e.g.*, EX1005 at §3.4 (“Figure 3 shows the outline of our method.”).

b. Limitation 8[a]

“providing a plurality of domains in the microprocessor;”

As discussed with respect to Claim limitation 1[a], Iyer teaches providing a plurality of domains in the microprocessor. §VI.A.1.

c. Limitation 8[b]

“clocking each of the domains separately at a clock frequency”

As discussed with respect to Claim limitation 1[b], Iyer teaches clocking each of the domains separately at a clock frequency. §VI.A.1.

d. Limitation 8[c]

“applying a voltage to each of the domains separately;”

As discussed with respect to Claim limitation 1[c], Iyer teaches applying a voltage to each of the domains separately. §VI.A.1.

e. Limitation 8[d]

“operating the microprocessor such that each domain operates synchronously, while the domains operate asynchronously relative to one another; and”


Iyer teaches this limitation. EX1003, ¶99. Iyer teaches “using a Globally Asynchronous Locally Synchronous (GALS) organization,” EX1005 at Abstract, and focuses on the design of “a microprocessor in a GALS style.” *Id.* at §1. Notably, operating a processor in this manner was already well known in the art for almost two decades. EX1003, ¶99; EX1005 at §1.1 (“GALS systems have been studied in detail by Chapiro in his 1984 thesis”); *see, e.g.*, EX1012.

f. Limitation 8[e]

“dynamically controlling the clock frequency and the voltage in each of the plurality of domains independently of the clock frequencies and the voltages in others of the plurality of domains.”

Iyer teaches this limitation. EX1003, ¶100. As discussed with respect to claim limitations 1[b] and 1[c], Iyer teaches the ability to dynamically control the clock frequency and voltage in each domain independently of the other domains. §VI.A.1. Iyer further teaches the method of dynamically changing

frequency and voltage. See EX1005 at §1 (“[h]ow can we ***exploit*** the extra flexibility offered by independent clock domains in a GALS processor and what is the impact of ***using*** a Multiple Clock Dynamic Voltage (MCDV) scheme”) (emphasis added). Iyer also teaches that “Figure 3 shows the outline of our ***method.***” EX1005 at §3.4 (emphasis added). Figure 3 teaches dynamically controlling clock frequency:



EX1005, Fig. 3.

See also *id.* at §3.4 (“method for dynamic clock management involves changing operating frequencies (and supply voltages)”).

6. Claim 10

“The method of claim 8, wherein step (d) comprises providing a queue for communication between at least two of the domains.”

As discussed with respect to Claim 5, Iyer teaches providing a queue for communication between at least two of the domains. §VI.A.4.

7. Claim 12

“The microprocessor of claim 8, wherein there are at least four of said domains.”

As discussed with respect to Claim 3, Iyer teaches the method of operating the microprocessor wherein there are at least four domains. §VI.A.2.

B. Ground II.A: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Njolstad (EX1009)

1. Claim 1

a. Limitation 1[pre]⁴


To the extent the preamble is limiting, Njolstad teaches this limitation. EX1003, ¶103. For instance, Njolstad teaches a digital signal processor (DSP) module, which a POSITA would have understood to be a type of microprocessor. *See* EX1009 at §IV.C; EX1003, ¶103. Njolstad further discloses that its teachings are intended to further existing work on scaling in a “microprocessor.” EX1009 at §III. Njolstad also discloses multiple clock domains. For instance, Njolstad teaches a “globally-asynchronous, locally-synchronous paradigm.” EX1009 at Abstract. As a POSITA would have understood, a GALS paradigm contains multiple clock

⁴ *See* Appendix of Challenged Claims at the end of this Petition for corresponding claim language of each claim limitation.

domains. EX1003, ¶103. Njolstad discloses, for instance, “local clock generation” in each domain. EX1009 at Abstract, *see also id.* at §IV.D.

b. Limitation 1[a]

Njolstad teaches this limitation. EX1003, ¶104. As discussed with respect to limitation 1[pre], Njolstad’s GALS paradigm includes a plurality of domains. §VI.A.1.a. This plurality of clock domains can be seen, for instance, in Njolstad’s Figure 3, which shows four domains with four different voltages (V_{DDi}) and four different frequencies (f_{CLKi}):




EX1009, Fig. 3.

c. Limitation 1[b]

i. Limitation 1[b][i]

Njolstad teaches this limitation. EX1003, ¶105. Njolstad discloses a “locally generated clock rate” for each domain, noting that the “local clock is generated by a ring oscillator.” EX1009 at §IV.D. The clocks for separately


generating a clock signal at a frequency for each domain are again visible in Figure 3, as f_{CLK1} , f_{CLK2} , f_{CLK3} , and f_{CLK4} , respectively:



EX1009, Fig. 3.

ii. Limitation 1[b][ii]

Njolstad teaches this limitation. EX1003, ¶106. In a given domain, the local frequency (“Clk gen.”) is based on the local processing demands (“RATE REQUIREMENT_i”) and scaled accordingly. EX1003, ¶106; EX1009 at §IV.D (“[a] clock signal with the correct frequency [is] given by the logical depth of the synchronous module”). For instance, a given domain or “module” is shown in Njolstad’s Figure 5:



EX1009, Fig. 5 (annotated).

The local “rate requirement can *dynamically* be configured” such that “the locally generated clock rate is just high enough to perform the actual processing task” at that rate. EX1009 at §IV.D (emphasis added). As Njolstad notes, from “a system point of view, faster operation of the processing element than given by the rate requirement does not give any better system performance, while slower operation than the requirement can of course not be tolerated.” EX1009 at §IV.D. For this reason, Njolstad implements a dynamically changeable frequency for each domain. EX1003, ¶107.


d. Limitation 1[c]

Njolstad teaches this limitation. EX1003, ¶108. As a POSITA would have understood and as Njolstad expressly discloses, voltage and frequency are

varied together. *Id.*; *see also* EX1009 at, *e.g.*, §IV.D (“scale down the local supply voltage such that the locally generated clock rate is just high enough to perform the actual processing task”). For instance, frequency cannot be substantially increased without increasing voltage, or the clock signal integrity will deteriorate. EX1003 at ¶108. And decreasing clock frequency without decreasing voltage would defeat the energy-saving focus of the technology. *Id.*; EX1009 at §III.C (“it is only by local supply voltage selection rate-optimal energy saving can be attained”). For this reason, Njolstad’s core focus relates to “locally dynamic voltage scaling (LDVS) adapted to the actual processing rate requirements” for each domain. EX1009 at Abstract. Notably, Njolstad dedicates an entire section this concept. *See* EX1009 at §IV.C. (“C. Means to generate different local supply voltages”). As Njolstad’s concludes, “[w]e have described GALS-LDVS” with “locally rate-adapted dynamic voltage scaling.” EX1009 at §V; *see also id.* at §IV.D (further discussion of this concept).

2. Claim 3

Njolstad teaches this limitation. EX1003, ¶109. Njolstad’s Figure 3, for instance, shows four clock domains:



EX1009, Fig. 3 (annotated).

3. Claim 4


a. Limitation 4[a]

Njolstad teaches this limitation. EX1003, ¶110. Njolstad teaches that its “rate requirement can dynamically be configured, by the special control signal RATE-REQUIREMENT_i” in order to ensure that there is not “faster operation of the processing element than” required. EX1009 at §IV.D. As a threshold matter, this signal is itself a determination of slack, because the rate requirement dynamically resolves under-utilization of processing capacity. Ex. 1003, ¶110. Njolstad also teaches reducing clock frequency “when no data processing is performed.” EX1009 at §IV.D. Thus, as further discussed with respect to limitation 4[b], below, Njolstad expressly and implicitly discloses monitoring slack in a given domain. *Id.*

b. Limitation 4[b]

Njolstad teaches this limitation. EX1003, ¶111. Njolstad teaches monitoring slack, and further teaches dynamically adjusting frequency and voltage in response to said slack. EX1003, ¶112. In fact, Njolstad explicitly discloses that “[t]he control signal RATE-REQUIREMENT_i can be used to *scale down* the local supply *voltage* such that the locally generated *clock rate* is *just high enough* to perform the actual processing task and to produce output data at the given rate.” EX1009 at §IV.D. This disclosure leaves no room for doubt that Njolstad teaches (i) monitoring slack and reducing (ii) local frequency and (iii) local voltage in order to reduce or eliminate said slack.

As Njolstad notes, slack monitoring and reduction is executed “locally”, *id.*—that is, independently for each domain:




EX1009, Fig. 5 (annotated).

In sum, within a given clock domain, the rate requirement signal is determined by monitoring slack. EX1003, ¶113. This signal determines whether scaling is required with respect to the frequency generated (“Clk gen.”) and corollary the supply voltage (“Vdd gen.”). *Id.*; EX1009, Fig. 5.

4. Claim 5

Njolstad renders obvious this limitation. EX1003, ¶114. As would have been obvious to a POSITA, and as Njolstad expressly notes, inter-domain communication in a GALS paradigm requires “some given protocol, such that e.g., the processing does not start before the module has input data ready.” EX1009 at IV.D. The simplest and most well-known mechanism for enabling asynchronous inter-domain communication is the use of buffer—which is a type of queue—such that the producer and consumer domains may operate at distinct frequencies. EX1003, ¶114. Njolstad, in particular, teaches toward the straightforward implementation of a queue, because the data transfer shown is unidirectional, as shown in Figure 4:



EX1009, Fig. 4.

Thus, as a POSITA would have understood, a basic FIFO queue would manage the asynchronous operation of producer and consumer domains by simply buffering data until the consumer is ready to process it. EX1003, ¶115. The '443 Patent acknowledges that this implementation would have been obvious to a POSITA, noting that its own purported invention simply coopts existing queue designs. EX1001 at 5:1 (“Our design uses *existing* queue structures in a superscalar processor core”). In fact, queues had been the standard mechanism for enabling asynchronous communication in processors for over a decade. EX1003, ¶115. For instance, a 1989 paper titled *Micropipelines*, which Iyer refers to as “the classic introduction to asynchronous design,” illustrates a FIFO queue approach in detail. EX1005 at §1.1; EX1014 (*Micropipelines*) at 722 (“FIFOs often connect senders and receivers that have separate clocks.”). Thus, a POSITA would have been well aware of this approach and motivated to adopt it in accordance with Njolstad’s teaching that “some given protocol” is required for asynchronous communication. EX1009 at IV.D. And given the widespread and longstanding implementation of such queues, a POSITA would have reasonably expected to be successful in implementing said queue. EX1003, ¶115.

5. Claim 8

**a. Limitation 8

```
[pre]
```**

To the extent the preamble is limiting, Njolstad teaches this limitation. EX1003, ¶116. As discussed with respect to limitation 1

```
[pre]
```

, Njolstad teaches a microprocessor. §VI.B.1. Njolstad further teaches a “method” and “means” for operating said microprocessor. *See, e.g.*, EX1009 at IV.C.

b. Limitation 8[a]

Njolstad teaches this limitation. EX1003, ¶117. As discussed with respect to Claim limitation 1[a], Njolstad teaches providing a plurality of domains in the microprocessor. §VI.B.1.

c. Limitation 8[b]

Njolstad teaches this limitation. EX1003, ¶118. As discussed with respect to Claim limitation 1[b], Njolstad teaches clocking each of the domains separately at a clock frequency. §VI.B.1.

d. Limitation 8[c]

Njolstad teaches this limitation. EX1003, ¶119. As discussed with respect to Claim limitation 1[c], Njolstad teaches applying a voltage to each of the domains separately. §VI.B.1.

e. Limitation 8[d]

Njolstad teaches this limitation. EX1003, ¶120. Njolstad discloses a method of operating a microprocessor in a “globally-asynchronous locally-synchronous paradigm.” EX1009 at §I.

f. Limitation 8[e]

Njolstad teaches this limitation. EX1003, ¶121. As discussed with respect to Claim limitation 1[b] and 1[c], Njolstad discloses dynamically controlling the frequency and voltage in each domain independently. §VI.B.1.

6. Claim 10

Njolstad renders obvious this limitation. EX1003, ¶122. As discussed with respect to Claim 5, Njolstad renders obvious providing a queue for communication between at least two of the domains. §VI.B.4.

7. Claim 12

Njolstad renders obvious this limitation. EX1003, ¶123. As discussed with respect to Claim 3, Njolstad teaches the method of operating the microprocessor wherein there are at least four domains. §VI.B.2.

C. Ground II.B: Claims 5 and 10 Are Rendered Obvious By Njolstad (EX1009) in view of Liljeberg (EX1010)

1. Motivation to Combine Njolstad and Liljeberg

To the extent that it would not already have been obvious to a POSITA to implement Njolstad's teachings with a queue for asynchronous communication, this limitation is expressly disclosed by Liljeberg and would have been obvious to incorporate into Njolstad. EX1003, ¶124. Liljeberg is focused on solving the very same issues Njolstad addresses. *Compare* EX1009 at §I ("fully synchronous designs will eventually not be feasible on the chip level"), *with* EX1010 at §1 ("the system-on-chip design strategy brings new challenges to manage system level timing"). Njolstad and Liljeberg also offer aligned solutions. *Compare* EX1009 at §I ("it is imperative to depart from the fully synchronous chip design paradigm, and replace it with a globally-asynchronous locally-synchronous paradigm"), *with* EX1010 at §1 ("A viable alternative to overcome the problems of a globally clocked system is the globally asynchronous locally synchronous approach"). Furthermore, both references are focused on implementation in a CMOS integrated circuit. *See* EX1009 at Fig. 3 (depicting a CMOS model); EX1010 at §4 ("4. CMOS Implementation"). Thus, a POSITA would have reasonably expected to be successful in integrating the teachings of Liljeberg into Njolstad. Furthermore, whereas Njolstad is focused on dynamically scaling independent domains, Liljeberg is dedicated to inter-domain communication.

EX1003, ¶124. Both focuses must be addressed to enable a coherent and energy efficient GALS design, and thus a POSITA would have been motivated to combine these teachings. *Id.*


Furthermore, although Njolstad discloses a handshake protocol between domains, as a POSITA would understand, as a matter of flow control, implementing buffers in addition to a full handshake mechanism between producer and consumer domains is advantageous in minimizing latency and optimizing throughput. EX1003 ¶124. While Njolstad's handshake ensures data integrity and synchronization, a queue such as that taught by Liljeberg allows the producer to continue sending data without waiting for the consumer to be immediately ready, thus reducing latency. *Id.* Queues also help manage bursty traffic by absorbing spikes in data generation and decoupling the producer and consumer, allowing each to operate at their own optimal speeds. *Id.* For this reason, a POSITA would have been motivated to employ a queue irrespective of whether there is a full handshake protocol in place. *Id.*

2. Claim 5

Njolstad and Liljeberg teach or render obvious this limitation. EX1003, ¶125. As would have been obvious to a POSITA, and as Njolstad expressly notes, inter-domain communication in a GALS paradigm requires “some given protocol, such that e.g., the processing does not start before the module has

input data ready.” EX1009 at IV.D. The simplest and most well-known mechanism for enabling asynchronous inter-domain communication is the use of a queue—such that the producer and consumer domains may operate at distinct frequencies. EX1003, ¶125. As a POSITA would have known, a queue may manage the asynchronous operation of producer and consumer domains by simply buffering data until the consumer is ready to process it. EX1003, ¶125.

To this end, Liljeberg articulates a detailed implementation of the cross-boundary communication “protocol” expressed by Njolstad. *Id.* Liljeberg teaches transferring data “to the local synchronous module through the FIFO queue.” EX1010 at §3. This can be seen, for instance, in Liljeberg’s Figure 2, wherein the queue is implemented as dual-ported RAM that bridges the asynchronous domain and the synchronous domain:



EX1010, Fig. 2 (annotated).

In its own words, Liljeberg “present[s] an asynchronous design approach to interface locally clocked modules” with “FIFO queue [] interface units.” EX1010 at §6. Thus, Njolstad in view of Liljeberg teaches a queue for communication between at least two domains. EX1003, ¶¶126-28.

3. Claim 10

Njolstad in view of Liljeberg renders obvious this limitation. EX1003, ¶129. As discussed with respect to Claim 5, Njolstad in view of Liljeberg teaches providing a queue for communication between at least two of the domains. §VI.C.2.

D. Ground III.A: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Soerensen (EX1006)


1. Claim 1

a. Limitation 1[pre]

To the extent the preamble is limiting, Soerensen teaches this limitation. EX1003, ¶130. Soerensen teaches a processor, and further teaches that the “processor may, for instance, be fabricated on a single semiconductor chip,” i.e., a microprocessor. EX1006 at [0028]; EX1003, ¶130. Further, Soerensen teaches multiple clock domains. *See, e.g.*, EX1006 at [0037] (“clock signals in each clock domain”); *id.* (“each of clock domains [] 70a, 70b, and 70c, respectively”). Soerensen teaches, for instance, that “at least one processor has a plurality of clock domains, each clock domain providing a clock signal . . . in each of the plurality of clock domains.” EX1006 at 8, Cl. 6.

b. Limitation 1[a]

Soerensen teaches this limitation. EX1003, ¶131. As discussed with respect to Claim 1[pre], Soerensen teaches a “plurality of clock domains.” EX1006 at 8, Cl. 6. And Soerensen’s Figure 2, for instance, shows an example plurality of clock domains 90a, 90b, and 90c:



EX1006, Fig. 2 (cropped and annotated).

c. Limitation 1[b]

i. Limitation 1[b][i]

Soerensen teaches this limitation. EX1003, ¶132-33. Soerensen teaches separately generating clock signals such that each domain may have a separate frequency. EX1006 at [0040] (“For example, clock tree 50a drives the clock signal of clock domain 90a and distributes it to the low level components of

DSP **22a**. Similarly, clock tree **50e** drives the clock signal of clock domain **90c** and distributes it to the low level logic components of MCU **24a**.”). Soerensen is clear that each domain receives its own signal at an independent frequency. *See* EX1006 at [0038] (“For example, clock control **110** includes multiplexers in each of clock domains (e.g., multiplexers **70a**, **70b**, and **70c**, respectively)”). This generates an independent and asynchronous frequency in each domain. *Id.* To this end, Soerensen claims a “plurality of clock domains, *each clock domain providing a clock signal* at a plurality of frequencies.” Soerensen at 8, Cl. 6 (emphasis added).


i. **Limitation 1[b][ii]**

Soerensen teaches this limitation. EX1003, ¶133. Soerensen discloses, for instance, “dynamically adjust[ing] the frequency of the clock signals provided . . .” EX1006 at [0023]. Soerensen further discloses that this dynamic frequency scaling is executed independently for each domain. EX1006 at [0081] (“the MCU [domain] may finish its task and require only a 45 MHz clock in order to meet the computational demands of the current processing state.”). In this instance, Soerensen teaches that “the change in frequency required by the MCU (e.g., 45 MHz) may be generated by adjusting the divider factor of the appropriate clock domain.” *Id.* at [0082]. Thus, the frequency in, e.g., the MCU domain is varied dynamically and independently of the frequencies generated for the plurality of other domains. As another, example, Soerensen notes that “[a]t some later instant,

the DSP may need to process another high priority task.” *Id.* at [0083]. In such an instance, the software will “identify the new clock frequency requirement of the DSP, for instance, 240 MHz” and adjust that domain’s frequency accordingly. *Id.* Soerensen thus emphasizes that the frequency adjustments within each domain are made independently of the other domains. As a further example, Soerensen discloses that, for instance, the frequencies within “clock domain 90*b*” and “clock domain 90*c*” may each be reduced by disparate amounts in order to provide the requisite frequencies to the subsidiary processor components in each domain. EX1006 at [0038].

d. Limitation 1[c]

Soerensen teaches this limitation. EX1003, ¶134. For instance, Soerensen teaches that “by monitoring the clock frequency requirement” the voltage can be “*dynamically controlled*” to reduce slack. EX1006 at [0048] (emphasis added). As a POSITA would have understood, voltage is varied proportionally to the frequency variation in each domain. EX1003, ¶134. Predictably, Soerensen claims “determining the voltage requirement based on the at least one associated frequency of the clock domains.” EX1006 at 9, Cl. 11. The relationship between a domain’s frequency and the requisite voltage is also outlined in Soerensen’s Figure 5:



EX1006, Fig. 5.

As Soerensen notes, “when going up to a higher power state, the voltage may be stepped up before the clock frequency is increased.” EX1006 at

[0064]. Conversely, when “going to a lower power state, frequency may be first decreased before reducing the voltage” *Id.*; *see also id.* at [0058] (“the clock frequency requirement is evaluated in order to determine a voltage requirement necessary to support the new clock frequency requirement”). Thus, Soerensen teaches the dynamic variance of both frequency and, consequentially, voltage in each domain. EX1003, ¶135.

2. Claim 3

Soerensen teaches this limitation. EX1003, ¶136. Soerensen expressly discloses such an embodiment. EX1006 at [0043] (“Plot **300** shows exemplary frequency values of four clock domains of a system clocking scheme”).

3. Claim 4

a. Limitation 4[a]

Soerensen teaches this limitation. EX1003, ¶137 Soerensen teaches that a “processor is considered to be consuming excess power when the computational tasks, functions, and operations of the processor may be performed, and/or the processing demands met, at a lower power state.” EX1006 at [0017]. The state described by Soerensen—frequency in excess of the requirement—is a state of slack, and reducing such slack is a primary focus of Soerensen’s disclosures. EX1003, ¶137. To be clear, Soerensen’s use of the term “processor” is intended to encompass “any apparatus that performs logic operations,

computational tasks, control functions, etc.” EX1006 at [0017]. A POSITA would have understood this to include, for instance, an instruction unit or execution unit, at least because such units perform logic operations and computational tasks. EX1003, ¶137-38.

Soerensen repeatedly discloses determining slack in at least one domain. *See, e.g.*, EX1006 at [0058] (“a DPC *monitors* the clock *frequency requirement* of the processor”); *id.* at [0092] (“a comparator may map information provided about the clock *frequency requirement* of the processor”); *id.* at [0026] (“*identifying* a clock *frequency requirement*”). Where there is a “delta” between the requirement and the frequency given, “there may be opportunities to reduce both the clock signal frequencies and the voltage.” EX1006 at [0061], [0027]. This teaching recurs throughout Soerensen’s disclosures. *See, e.g., id.* at [0041] (“recognizing the clock frequency requirement . . . configure the frequency distribution to reduce the intervals when *excess clock frequencies* are being generated”).

b. Limitation 4[b]

Soerensen teaches this limitation. EX1003, ¶139. As Soerensen notes, its invention is directed to reducing power consumption in processors, among other things. EX1006 at [0002]. And as discussed with respect to Claim limitation 4[a], Soerensen teaches monitoring slack *and reducing* frequency and voltage in

response to slack. EX1006 at [0027] (“opportunities to reduce both the clock signal frequencies and the voltage”). A POSITA would have recognized that the ultimate point of reducing frequency in response to slack is to ultimately reduce voltage and thus reduce power consumption. EX1003, ¶139; EX1006 at [0018] (“power consumption of a device is related to the square of the voltage supplied to the device’s processor and proportional to the frequency at which the processor operates”). Unsurprisingly, Soerensen teaches that “by monitoring the clock frequency requirement of a processor, the voltage supplied to the processor can be dynamically controlled such that the processor spends less time supplying a voltage to the processor in excess of the requirement.” EX1006 at [0048].

4. Claim 5

Soerensen renders obvious this limitation. EX1003, ¶140. As discussed with respect to Claim limitation 1[pre], Soerensen discloses multiple asynchronous clock domains necessitating interdomain data transfer. *Id.*; *see also* EX1006 at [0028] (“various subsystems of the processor are connected to a system bus **30** in order to communicate and/or transfer data between the various components”). A POSITA would have understood that the system bus requires a mechanism for coherent transfer between domains, including a queue. EX1003, ¶140. It would have been obvious for a POSITA to implement a queue for interdomain transfer because queues were well known in the art and further

because a POSITA would have recognized that Soerensen was compatible with a queue. EX1003, ¶140. By way of example, Hemani—a 1999 publication—outlines a POSITA’s approach to the “communication refinement step” and advocates for the consideration of a queue in any compatible asynchronous construct. EX1007 at §2; EX1003, ¶140. The ’443 Patent itself acknowledges that this implementation would be obvious to a POSITA, noting that its own purported invention simply coopts existing queue designs. EX1001 at 5:1 (“Our design uses *existing* queue structures in a superscalar processor core”). Thus, a POSITA would find it obvious and be motivated to implement a queue in order to decouple consumer and producer domains in Soerensen’s design, and further reasonably expect to be successful in doing so. EX1003, ¶140.

5. Claim 8

a. Limitation 8[pre]

To the extent the preamble is limiting, as discussed with respect to Claim limitation 1[pre], Soerensen teaches a microprocessor. EX1003, ¶141; §VI.D.1. Soerensen further teaches that its invention relates “to methods” for dynamically controlling said microprocessor. *See, e.g.*, EX1006 at [0002]; *id.* at [0027] (disclosing the method of “identifying a clock frequency requirement of a processor and adjusting the voltage supplied”).

b. Limitation 8[a]

Soerensen teaches this limitation. EX1003, ¶142. As discussed with respect to Claim limitation 1[a], Soerensen teaches providing a plurality of domains. §VI.D.1.

c. Limitation 8[b]


Soerensen teaches this limitation. EX1003, ¶143. As discussed with respect to Claim limitation 1[b], Soerensen teaches clocking each of the domains separately. §VI.D.1.

d. Limitation 8[c]

Soerensen teaches this limitation. EX1003, ¶144. As discussed with respect to Claim limitation 1[c], Soerensen teaches applying a voltage to each of the domains separately. §VI.D.1.

e. Limitation 8[d]

Soerensen teaches this limitation. EX1003, ¶145-46. As discussed with respect to Claim limitation 1[a], Soerensen teaches a plurality of domains. §VI.D.1. Soerensen further teaches that each domain is synchronous, while the domains are asynchronous relative to one another. The dotted lines in Figure 2 outline domains, within which all components are synchronous:



EX1006, Fig 2 (annotated).

Soerensen further teaches that each domain is asynchronous relative to the others, because each of multiplexers 70a, 70b, and 70c may generate a different frequency independently. EX1003, ¶146. For instance, Soerensen discloses that “clocked components operating in domain **90b**” may have “a clock signal having half the frequency of clock signal **84**.” EX1006 at [0038]. Critically, there is no requirement that the domains maintain a fixed ratio relative to one

another. *Id.* Thus, each domain in Soerensen is internally synchronous, and fully asynchronous relative to the others. EX1003, ¶146.

f. Limitation 8[e]

Soerensen teaches this limitation. EX1003, ¶147. As discussed above with respect to Claim limitations 1[b] and 1[c], Soerensen teaches the ability to dynamically adjust both frequency and voltage in each domain independently. §VI.D.1. Soerensen further discloses the method of operating the processor in this manner. Indeed, Soerensen’s disclosures are directed to “a method of dynamic power control.” Soerensen at [0006]. In line with this inventive concept, Soerensen’s Figure 5 outlines “illustrates a method of providing dynamic power control.” *Id.* at [0014], Fig. 5.

6. Claim 10

Soerensen renders obvious this limitation. EX1003, ¶148. As discussed with respect to Claim 5, Soerensen renders obvious the method of providing a queue for inter-domain communication. §VI.D.4.

7. Claim 12

As discussed with respect to Claim 3, Soerensen teaches the method of operating the microprocessor wherein there are at least four domains. §VI.D.2.

E. Ground III.B: Claims 1, 3, 4, 5, 8, 10, and 12 Are Rendered Obvious By Soerensen (EX1006) in view of Hemani (EX1007)

1. Motivation to Combine Soerensen and Hemani

A POSITA would be motivated to combine the teachings of Soerensen and Hemani because they are complementary disclosures focused on the same issue. EX1003, ¶150-51. For instance, Hemani is titled *Lowering Power Consumption in Clock by Using Globally Asynchronous Locally Synchronous Design Style*, and Soerensen's "invention relates to reducing the power consumption in processors" by "determining the clock frequency requirement." EX1007 at Cover; EX1006 at [0002]. Both teachings disclose dividing a processor into distinct clock regions which are locally synchronous but asynchronous to one another. EX1003, ¶151. Soerensen, like the '443 patent, refers to these blocks as "clock domains" while Hemani refers to these regions as "synchronous blocks." EX1006 at [0037]; EX1007 at §1. The core means of achieving their respective inventions are similarly uniform. *Compare* EX1006 at [0038] (describing "clock signals having reduced frequencies" in domains with slack), *with* EX1007 at §1 ("the frequency at which each SB is clocked can be tailored to local needs").

Because Hemani provides a broad overview of means by which a GALS processor may be divided into multiple clock domains, and Soerensen provides a detailed description of the dynamic control of such a processor, a POSITA would be motivated to combine these teachings in developing a

dynamically scalable GALS processor. EX1003, ¶150-51. And because both teachings relate to lowering power consumption in globally asynchronous, locally synchronous microprocessors, a POSITA would reasonably expect to be successful in integrating the teachings of each reference. *Id.*

2. Claim 1

a. Limitation 1[pre]

To the extent the preamble is limiting, Soerensen and Hemani render it obvious. EX1003, ¶152. As noted above, Soerensen teaches a processor, and further teaches that “processor **20** may, for instance, be fabricated on a single semiconductor chip,” i.e., a microprocessor. EX1006 at [0028]; EX1003, ¶152. To the extent that Soerensen does not explicitly disclose the “microprocessor” claim language, Hemani discloses that its teachings are directed to “high performance microprocessors.” EX1007 at §1. It would have been obvious to a POSITA that Soerensen’s systems would have been implemented on Hemani’s “microprocessor,” because Hemani is explicit regarding the mechanical aspects to which GALS systems such as Soerensen were intended to be applied. *See* EX1003, ¶152.

Further, Soerensen teaches multiple clock domains. *See, e.g.,* EX1006 at [0037] (“clock signals in each clock domain”); *id.* (“each of clock domains [] 70a, 70b, and 70c, respectively”). Similarly, Hemani—which refers to clock


domains as *synchronous blocks*—teaches multiple “synchronous blocks” with “independent local clock generators.” EX1007 at §1.

b. Limitation 1[a]

Soerensen and Hemani render this limitation obvious. EX1003, ¶154.


Soerensen, for instance, teaches a “plurality of clock domains.” Soerensen at 8, Cl.

6. And Soerensen’s Figure 2 shows this plurality of clock domains:



EX1006, Fig. 2 (cropped and annotated).

Hemani's Figure 2 similarly shows a plurality of domains in a GALS system, which it again refers to as "synchronous blocks" (SBs):



EX1007, Fig. 2. These complementary teachings confirm the compatibility of these references in combination. *See* EX1003, ¶156.

c. Limitation 1[b]

i. Limitation 1[b][i]

Soerensen and Hemani render this limitation obvious. EX1003, ¶157. Soerensen teaches, for instance, that the “various frequencies required by the processor are often generated from the system clock by a phase locked loop (PLL) multiplier.” EX1006 at [0033]. Soerensen notes that this means is “merely exemplary and is meant to illustrate one scheme by which the various clock signals required by a processor may be controlled and distributed.” *Id.* Hemani discloses

this method, as well as alternatives, teaching “there are two possibilities, either to have an independent local clock generator or to use a global reference clock signal with frequency multiplier in SBs.” EX1007 at §3.4. In the former instance, Hemani teaches “[l]ocal clock generators based on ring oscillators.” EX1007 at §3.4. In sum, Soerensen and Hemani disclose multiple means for using a clock to separately generate frequencies for each domain. EX1003, ¶157. As Hemani observes, a POSITA would have been aware of—and free to choose between—the two disclosed alternatives, which both achieved the same objective and were suitable. EX1007 at §3.4.

ii. **Limitation 1[b][ii]**

Soerensen and Hemani render this limitation obvious. EX1003, ¶158. As noted above, Soerensen discloses “dynamically adjust[ing] the frequency of the clock signals provided” EX1006 at [0023]. Soerensen further discloses that this scaling is executed independently for each domain. EX1006 at [0081] (“the MCU [domain] may finish its task and require only a 45 MHz clock in order to meet the computational demands of the current processing state”; *id.* (“the change in frequency required by the MCU (e.g., 45 MHz) may be generated by adjusting the [] appropriate clock domain.”).

Hemani, too, teaches “independent local clock generators” based on ring oscillators, and the ability to dynamically “manipulate the frequency of the

ring oscillator[s].” EX1007 at §3.4. This dynamic scaling achieves the same purpose—it “allows SBs to run at different clock speeds” and thus affects “power consumption and savings.” EX1007 at §5. Thus, each reference discloses local frequencies which are independently and dynamically changeable. EX1003, ¶159.

d. Limitation 1[c]

Soerensen and Hemani render this limitation obvious. EX1003, ¶160. As noted above, §II.A.1, Soerensen teaches that voltage can be “dynamically controlled.” EX1006 at [0048]. For instance, Soerensen—which teaches individual and independent frequencies in each domain—notes that “when going up to a higher power state, the voltage may be stepped up before the clock frequency is increased.” EX1006 at [0064]. And conversely, when “going to a lower power state, frequency may be first decreased before reducing the voltage” *Id.*

To the extent that Soerensen does not explicitly disclose this limitation, Hemani does. EX1003, ¶160-61. As a POSITA would have understood, the purpose of having independent and dynamically variable clock domains is to enable power savings by reducing the voltage in each domain whenever able, without sacrificing overall performance. *Id.* For this reason, a POSITA developing a GALS architecture such as Soerensen’s would have been motivated to reduce the voltage in each domain whenever the frequency is reduced in that domain. *Id.* Recall that Hemani discloses that “the frequency at which each SB is clocked *can*

be tailored to the local needs, thus reducing . . . overall power consumption” and that “[p]ower consumption in VLSI systems is reduced by reducing voltage.” EX1007 at §1. To this end, Hemani expressly teaches “dynamic voltage scaling of the power supply voltage at each SB containing a ring oscillator” in order to achieve an “energy efficient design.” EX1007 at §3.4.

3. Claim 3

Soerensen and Hemani render this limitation obvious. EX1003, ¶162. Soerensen expressly discloses such an embodiment. EX1006 at [0043] (“Plot **300** shows exemplary frequency values of four clock domains of a system clocking scheme”). Hemani, for its part, discloses up to 52 domains, which it refers to as synchronous blocks (SBs), as shown in the bottom right quarter of Table 1:

| Table 1: Design Data | | | |
|-----------------------------|------------------------|---------------|-----------------------|
| | Gate Count in Millions | Die Size (mm) | Number of Initial SBs |
| Design 1 | 1.125 | 7.5 X 7.5 | 24 |
| Design 2 | 2.09 | 10.9 X 10.9 | 46 |
| Design 3 | 3.04 | 12.8 X 12.8 | 52 |

EX1007, Table 1.

In sum, there is nothing inventive or non-obvious about designing four (or more) domains. EX1003, ¶163.

4. Claim 4

a. Limitation 4[a]

Soerensen and Hemani render this limitation obvious. EX1003, ¶164. Soerensen teaches with respect to slack that a “processor is considered to be consuming excess power when the computational tasks, functions, and operations of the processor may be performed, and/or the processing demands met, at a lower power state.” EX1006 at [0017]. Soerensen’s use of the term “processor” is intended to encompass “any apparatus that performs logic operations, computational tasks, control functions, etc.” EX1006 at [0017]. A POSITA would have understood this to include, for instance, domains comprising an instruction unit or execution unit, as discussed above. EX1003, ¶164. As discussed above with respect to limitation 1[b], Soerensen discloses a plurality of clock domains with independent frequencies. §VI.E.2 Soerensen further discloses a “controller adapted to determine a voltage requirement based on the clock frequency requirement.” EX1006 at 9, Cl. 12. In other words, Soerensen teaches monitoring and adapting to slack in a domain based on the “delta” between the existing frequency and voltage and the frequency and voltage “required.” *Id.*; *see also id.* at [0061]. As Soerensen further explains, “the clock frequency requirement may represent the minimum frequency that each clock domain requires in order to perform the functions and/or

computational tasks of the components operating in that clock domain.” *Id.* at [0059].

b. Limitation 4[b]

Soerensen and Hemani render this limitation obvious. EX1003, ¶165. As Soerensen notes preliminarily, its invention is directed to reducing power consumption in processors. EX1006 at [0002]. And as discussed with respect to Claim limitation 4[a], Soerensen teaches monitoring slack and reducing frequency and voltage in response to slack. A POSITA will recognize that the point of reducing frequency in response to slack is to ultimately reduce voltage and thus reduce power consumption. EX1003, ¶165; EX1006 at [0017] (“it is advantageous to reduce the amount of time in which the device’s processor is consuming excess power”). Unsurprisingly, Soerensen teaches that “by monitoring the clock frequency requirement of a processor, the voltage supplied to the processor can be dynamically controlled such that the processor spends less time supplying a **voltage** to the processor in excess of the requirement.” EX1006 at [0048]; *see also* EX1006 at [0041] (emphasis added) (“clock control **100** may configure the **frequency distribution** to reduce the intervals when excess clock frequencies are being generated”).

As noted throughout, the purpose of monitoring slack is to reduce frequency—and consequentially voltage—whenever able, and thus reduce power

consumption. EX1003, ¶166. Hemani, which is focused on “lowering power consumption” teaches “dynamic voltage scaling of the power supply voltage at *each* SB.” EX1007 at §3.4 (emphasis added). As would have been clear to a POSITA, voltage in a domain can only be reduced in response to some measure of slack. EX1003, ¶166. Reducing voltage without slack in the frequency would impermissibly degrade the clock signal. EX1003, ¶166. As Soerensen notes, several “methods of determining a voltage requirement from a clock frequency requirement” will be apparent to a POSITA. EX1006 at [0060]. But in any event, “power consumption is . . . proportional to the frequency at which the processor operates” and thus Hemani implicitly or expressly teaches monitoring and responding to slack. EX1006 at [0018]; *see also* EX1007 at §3.4; EX1003, ¶166.

5. Claim 5

Soerensen and Hemani render this limitation obvious. EX1003, ¶167. As discussed above, Soerensen and Hemani teach or render obvious Claim 1. §VI.E.2. And as discussed with respect to Claim limitation 1[pre], Soerensen discloses multiple asynchronous clock domains necessitating interdomain data transfer. *See also* EX1006 at [0028] (“various subsystems of the processor are connected to a system bus **30** in order to communicated and/or transfer data between the various components”). A POSITA would have understood that the system bus requires a mechanism for coherent transfer between domains, including

a queue. EX1003, ¶167-68. It would have been obvious for a POSITA to implement a queue for interdomain transfer because queues were well known in the art and Soerensen is compatible with a queue. *Id.* The '443 Patent itself acknowledges that this implementation would be obvious to a POSITA, noting that its own purported invention simply coopts existing queue designs. EX1001 at 5:1 (“Our design uses *existing* queue structures in a superscalar processor core”). This a POSITA would find it obvious and be motivated to implement a queue in Soerensen’s design. EX1003, ¶167.

To the extent that Soerensen does not expressly disclose this claim language, Hemani teaches this limitation, and for the reasons noted above, a POSITA would be motivated to employ Hemani’s queue structure. EX1003, ¶168. Specifically, Hemani describes multiple interdomain communication methods, including a “FIFO.” EX1007 at §2. As discussed with respect to Iyer and Liljeberg, a FIFO is a “first in first out” queue. EX1003, ¶168. Of the “existing queue structures” acknowledged by the '443 specification, EX1001 at 5:2, this queue was among the most well-known, and thus a POSITA would have been motivated to adapt it and would reasonably expect to be successful in its implementation. EX1003, ¶168. As further motivation, Hemani encourages the selection of the most energy efficient—or “cheapest”—structure, and thus where a FIFO is most energy

efficient, a POSITA would be especially motivated to employ such a queue. EX1007 at §2; EX1003, ¶168.

6. Claim 8

a. Limitation 8[pre]

To the extent the preamble is limiting, as discussed with respect to Claim limitation 1[pre], Soerensen and Hemani teach a microprocessor. §VI.E.2. Soerensen further teaches that its invention relates “to methods” for dynamically controlling said microprocessor. *See, e.g.*, EX1006 at [0002]; *id.* at [0027] (disclosing the method of “identifying a clock frequency requirement of a processor and adjusting the voltage supplied”). Hemani similarly teaches “suggested methods” of operation. *See* EX1007 at §3.4.

b. Limitation 8[a]

Soerensen and Hemani render this limitation obvious. EX1003, ¶170. As discussed with respect to Claim limitation 1[a], Soerensen and Hemani teach providing a plurality of domains. §VI.E.2.

c. Limitation 8[b]

Soerensen and Hemani render this limitation obvious. EX1003, ¶171. As discussed with respect to Claim limitation 1[b], Soerensen and Hemani teach or render obvious clocking each of the domains separately. §VI.E.2.

d. Limitation 8[c]


Soerensen and Hemani render this limitation obvious. EX1003, ¶172.

As discussed with respect to Claim limitation 1[c], Soerensen and Hemani teach or render obvious applying a voltage to each of the domains separately. §VI.E.2.

e. Limitation 8[d]

Soerensen and Hemani render this limitation obvious. EX1003, ¶173.

For instance, Soerensen's Figure 2 outlines domains, within which all components are synchronous:



EX1006, Fig 2.

Soerensen further teaches that each domain is asynchronous relative to the others, because each of multiplexers 70a, 70b, and 70c may generate a different frequency independently. *Id.* at [0037]; EX1003, ¶174.

Hemani is even more explicit, disclosing a “GALS architecture [] composed of large synchronous blocks (SBs) which communicate with each other on an asynchronous basis.” EX1007 at §3.4

f. Limitation 8[e]

Soerensen and Hemani render this limitation obvious. EX1003, ¶176. As discussed above with respect to Claim limitations 1[b] and 1[c], Soerensen and Hemani teach or render obvious the ability to dynamically adjust both frequency and voltage in each domain independently. §VI.E.2. Soerensen further discloses the method of operating the processor in this manner. Indeed, Soerensen’s disclosures are directed to “a method of dynamic power control.” EX1006 at [0006]. And Hemani “propose[s] methods” for operating and evaluating their proposed GALS architecture. EX1007 at §Abstract.

7. Claim 10

Soerensen and Hemani render this limitation obvious. EX1003, ¶177. As discussed with respect to Claim 5, Soerensen and Hemani teach or render obvious the method of providing a queue for inter-domain communication. §VI.E.5.

8. Claim 12

As discussed with respect to Claim 3, Soerensen and Hemani teach or render obvious the method of operating the microprocessor wherein there are at least four domains. §VI.E.3; EX1003, ¶178.

F. Ground III.C: Claims 5 and 10 Are Rendered Obvious By Soerensen (EX1006) in view of Moore (EX1008)

1. Motivation to Combine Soerensen and Moore

To the extent that Soerensen alone does not render the subject matter recited by claims 5 and 10 obvious, it is also taught by Moore. EX1003, ¶179-81. As noted above and by the '443 Patent itself, queues were well known in the art, and perhaps the most common form of queue was a first-in-first-out (FIFO) queue. *Id.*; EX1001 at 5:2. Moore expressly discloses a FIFO queue for communication between synchronous domains which are asynchronous relative to one another. EX1008 at §5. A POSITA would have been motivated to use Moore's teachings in combination with Soerensen because Soerensen does not expressly outline a preferred mechanism for enabling asynchronous interdomain communication, and Moore's disclosures are focused on this very issue. EX1003, ¶181. Indeed, Moore is principally concerned with "communication between independent clock domains" in "globally asynchronous locally synchronous (GALS) system[s]." *See* EX1008 at Abstract. Similarly, because Moore is specifically focused on

interdomain communication in a GALS design, a POSITA would have reasonably expected to be successful in implementing Moore's queue in Soerensen's design.

2. Claim 5

Soerensen in view of Moore renders obvious this limitation. EX1003, ¶182. As discussed above, Soerensen renders obvious Claim 1. §VI.D.1. And as discussed with respect to Claim limitation 1[pre], Soerensen discloses multiple asynchronous clock domains necessitating interdomain data transfer. *See also* EX1006 at [0028] ("various subsystems of the processor are connected to a system bus **30** in order to communicated and/or transfer data between the various components"). A POSITA would have understood that said system bus requires a mechanism for coherent transfer between domains, including a queue. EX1003, ¶183. It would have been obvious for a POSITA to implement a queue for interdomain transfer because queues were well known in the art and Soerensen does not teach toward an incompatible alternative. *Id.* Specifically, in a construct such as Soerensen's, with multiple asynchronous domains, a POSITA would be motivated to implement a queue in order to reduce latency and decouple producers and consumers. Ex. 1003, ¶181. As the '443 Patent admits, this approach was well-known. EX1001 at 5:2.

Moore describes in detail the implementation of a FIFO queue in a GALS system, noting that "[a]dding an asynchronous FIFO between the producer

and consumer (see Figure 5) decouples the [signal] path.” EX1008 at §5. Figure 5 itself is similarly explicit in teaching a queue between clock domains:




Figure 5: Channel communication with FIFO buffering

EX1008, Fig. 5.

Moore’s Figure 5 shows a channel communication system with FIFO buffering, where the FIFO queue is placed between the synchronous producer and the synchronous consumer, effectively decoupling the Request (“req”) and Acknowledge (“ack”) paths and facilitating asynchronous communication between different clock domains. This queue enables interdomain transfer between components operating at different clock frequencies without loss or corruption by temporarily storing data from the producer until the consumer is ready to process it.

Therefore, Soerensen in view of Moore renders obvious the microprocessor of Claim 1, further comprising a *queue* for communication between at least two of the domains.

3. Claim 10

Soerensen in view of Moore renders obvious this limitation. EX1003, ¶177. As discussed with respect to Claim 5, Soerensen in view of Moore renders obvious the method of providing a queue for inter-domain communication. §VI.D.1.

7. SECONDARY CONSIDERATIONS

There is no evidence of any evidence of secondary indicia that would support non-obviousness. EX1003, ¶¶187-194.

8. NO JUSTIFICATION FOR DISCRETIONARY DENIAL

A. The Board Should Not Decline Institution Pursuant to *Fintiv*

The Board should not exercise its discretion under 35 U.S.C. § 314(a) to deny institution in view of the Texas Litigation (*see* Mandatory Notices, *infra*) because the balance of the *Fintiv* factors discussed below weigh in favor of institution.

Fintiv Factor 1 (likelihood of a stay) is neutral. There are two patents at issue in the Texas Litigation (the '443 Patent at issue here and U.S. Patent No. 7,623,439), and Petitioners will be challenging (more than) all asserted claims of both patents in IPR. Institution of the two IPR petitions would thus enable the

Board to resolve core issues in, if not completely moot, the Texas Litigation. As such, the District Court likely would stay the case pending the outcome of the IPRs in view of the potentially substantial judicial efficiencies to be gained. Thus, this factor is neutral.

Factor 2 (proximity of trial date to IPR deadline) is also neutral. The District Court set trial in the Texas Litigation for July 28, 2025, but “scheduled trial dates are unreliable and often change” so “[a] court’s schedule trial date, therefore, is not by itself a good indicator of whether the district court trial will occur before the statutory deadline for a final written decision.” *See* Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation (June 21, 2022) (“Vidal Memo”) at 8. This factor is neutral.

Factor 3 (investment in parallel proceeding) weighs against discretionary denial. The Texas Litigation is in its early stages, and no discovery has been taken yet beyond initial infringement and invalidity contentions and document productions associated therewith. The parties and the District Court have therefore invested relatively minimal resources to date, such that this factor weighs against discretionary denial.

Factor 4 (overlap between issues) weighs against discretionary denial. Petitioners stipulate that, if the IPR is instituted, Petitioners will not pursue in the Texas Litigation the same grounds asserted in this Petition. This “mitigates to some degree the concerns of duplicative efforts.” *See Sand Revolution II, LLC v. Continental Intermodal Group—Trucking LLC*, IPR2019-01393, Paper 24 at 12 (P.T.A.B. June 16, 2020). Moreover, this Petition challenges unasserted Claims 3 and 12, such that their invalidity will not be addressed in that proceeding. This factor thus weighs against discretionary denial so that the Board can address these non-overlapping claims and grounds.

Factor 5 (same parties) weighs slightly in favor of denial because Petitioners and Patent Owner are the same parties in both this proceeding and in the Texas Litigation (as is common in most IPR proceedings).

Factor 6 (other circumstances impacting discretion) weighs against discretionary denial. As discussed above, the Grounds present compelling, meritorious challenges to the '443 claims. Each of the references are direct hits on the alleged distinctions between the prior art considered by the examiner and the claimed invention. The asserted prior art thus weighs against denial because the merits are so strong.

Overall, two *Fintiv* factors are neutral, three factors weigh against discretionary denial, and only one factor weighs slightly in favor of discretionary denial. On balance, therefore, the *Fintiv* factors weigh against denial.

B. The Board Should Not Decline Institution Pursuant to *General Plastic*

This is the first petition challenging any of the '443 claims. Thus, there is no basis for the Board to use its discretion to deny institution pursuant to *General Plastic*.

C. The Board Should Not Decline Institution Pursuant To § 325(d)

The Board should not deny this Petition under the two-part test for §325(d) denial provide for by *Advanced Bionics*. None of the grounds here involve the same or substantially the same prior art or arguments previously before the PTO. No reference was applied during prosecution of the '443 Patent. *See* '443 NOA and March 4, 2005 IDS. EX1002. Further, the grounds presented herein are not cumulative of any prior art rejection made during prosecution of the '443 Patent, as no prior art rejection was entered. As noted above, each of the Grounds describe key characteristics (i.e., microprocessor domains) which were not disclosed by the prior art characterized by the examiner in the NOA.

To the extent that Patent Owner argues that the prior art applied in the Grounds is cumulative to any of the references disclosed by the search report in the

PCT application and in the March 4, 2005 IDS—in particular Sager and Overhouse—those references did not form the basis of a rejection at the Office. *Microsoft Corp. v. Parallel Networks Licensing, LLC*, IPR2015-00486, Paper 10 at 14–15 (P.T.A.B. July 15, 2015) (instituting review despite reference having been listed on Information Disclosure Statements initialed by the Examiner). Further, under the *Advanced Bionics* step 2 and *Becton, Dickinson* factors (c), (e), and (f), Section VI shows that the Office made a material error by not applying a reference teaching multiple microprocessor domains in a rejection of any claims. The Examiner’s silence with respect to (i) Sager and Overhouse and (ii) the art applied in the current Grounds before allowing the ’443 claims reflects error and oversight.

9. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8(A)(1)

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

Petitioners TCL Electronics Holdings Ltd. and TCL Industries Holdings Co. are the real parties-in-interest.

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

The ’443 Patent is one of two asserted patents in Civil Action No. 6:23-cv-000309-ADA (W.D. Tex.), filed April 26, 2023 (“Texas Litigation”). The other asserted patent in the Texas Litigation, U.S. Patent No. 7,089,439, is the subject of a recently filed IPR, IPR2024-01244. This action may affect, or be affected by, the Texas Litigation.

C. 37 C.F.R. § 42.8(b)(3) and (4): Counsel and Service Information

| Lead Counsel | Back-Up Counsel |
|--|---|
| Russell E. Levine, P.C.
(Reg. No. 32,153)
rlevine@kirkland.com

<u>Postal and Hand-Delivery Address:</u>
KIRKLAND & ELLIS LLP
333 West Wolf Point Plaza
Chicago, IL 60654
Telephone: (312) 862-2466
Facsimile: (312) 862-2200

Email: TCL-IV-IPR@kirkland.com | Joshua Glucoft (Reg. No. 67,696)
David Levin (<i>pro hac vice</i> forthcoming)
KIRKLAND & ELLIS LLP
2049 Century Park East
Suite 3700
Los Angeles, CA 90067
Tel: (310) 552-4200

Kat Li (Reg. No. 64,857)
Michael T. Dinnella (Reg. No. 82,580)
KIRKLAND & ELLIS LLP
401 Congress Avenue
Austin, TX 78701
Tel: (512) 678-9100 |

TCL concurrently submits a Power of Attorney, 37 C.F.R. §42.10(b), and consents to electronic service directed to the following email address: TCL-IV-IPR@kirkland.com.

10. PAYMENT OF FEES UNDER 37 C.F.R. §42.103

The undersigned authorizes the Office to charge the fee set forth in 37 C.F.R. §42.15(a)(1) for this Petition to Deposit Account No. 506092. Review of 7 claims is requested. The undersigned further authorizes payment for any additional

fees that may be due in connection with this Petition to be charged to the above-referenced deposit account.

11. CERTIFICATION OF GROUNDS FOR STANDING

Petitioners certify that the '443 Patent is available for IPR and that Petitioners are not barred or estopped from requesting an IPR of the Challenged Claims on the grounds identified herein. Petitioners certify: (1) Petitioners are not the owner of the '443 Patent; (2) Petitioners (or any real party-in-interest) have not filed a civil action challenging the validity of any claim of the '443 Patent; (3) Petitioners filed this Petition within one year of the date they were served with a complaint asserting infringement of the '443 Patent; (4) estoppel provisions of 35 U.S.C. §315(e)(1) do not prohibit this IPR; and (5) this Petition is filed more than nine months after the '443 Patent was granted.

12. CONCLUSION

Petitioner requests institution of an IPR and cancellation of the Challenged Claims.

Dated: August 2, 2024

Respectfully submitted,

/s/ Russell E. Levine, P. C.

Russell E. Levine, P.C. (Reg. No. 32,153)

KIRKLAND & ELLIS LLP

333 West Wolf Point Plaza

Chicago, Illinois 60654

Telephone: (312) 862-2000

Joshua Glucoft (Reg. No. 67,696)
David Levin (*pro hac vice* forthcoming)
KIRKLAND & ELLIS LLP
2049 Century Park East
Suite 3700
Los Angeles, CA 90067
Tel: (310) 552-4200

Kat Li (Reg. No. 64,857)
Michael T. Dinnella (Reg. No. 82,580)
KIRKLAND & ELLIS LLP
401 Congress Avenue
Austin, TX 78701
Tel: (512) 678-9100

TCL-IV-IPR@kirkland.com

Attorneys for Petitioner

APPENDIX OF CHALLENGED CLAIMS

| Claim Limitation | Claim Language |
|-------------------------|---|
| 1[pre] | A multiple clock domain microprocessor comprising: |
| 1[a] | a plurality of domains; |
| 1[b]i | for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, |
| 1[b]ii | the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains; and |
| 1[c] | for each of the plurality of domains, a voltage input for receiving a voltage which is dynamically changeable independently of the voltages applied to said others of the plurality of domains. |
| 3 | The microprocessor of claim 1, wherein there are at least four of said domains. |
| 4[a] | The microprocessor of claim 1, wherein the microprocessor is programmed to determine a slack in processing in one of the domains and |
| 4[b] | to reduce the clock frequency and the voltage in said one of the domains to reduce the slack. |
| 5 | The microprocessor of claim 1, further comprising a queue for communication between at least two of the domains. |
| 8[pre] | A method of operating a microprocessor, the method comprising: |
| 8[a] | providing a plurality of domains in the microprocessor; |
| 8[b] | clocking each of the domains separately at a clock frequency; |
| 8[c] | applying a voltage to each of the domains separately; |
| 8[d] | operating the microprocessor such that each domain operates synchronously, while the domains operate asynchronously relative to one another; and |
| 8[e] | dynamically controlling the clock frequency and the voltage in each of the plurality of domains independently of the clock frequencies and the voltages in others of the plurality of domains. |
| 10 | The method of claim 8, wherein step (d) comprises providing a queue for communication between at least two of the domains. |
| 12 | The method of claim 8, wherein there are at least four of said domains. |

CERTIFICATION OF COMPLIANCE

The undersigned certifies that the foregoing Petition complies with the type-volume limitations as mandated in 37 C.F.R. § 42.24, totaling 12,856 words. Counsel has relied upon the word count feature provided by Microsoft Word. This word count does not include the items excluded by 37 C.F.R. § 42.24 as not counting towards the word limit.

Dated: August 2, 2024

/s/ Russell E. Levine, P.C.
Russell E. Levine, P.C. (Reg. No. 32,153)
KIRKLAND & ELLIS LLP
333 West Wolf Point Plaza
Chicago, Illinois 60654
Telephone: (312) 862-2000

Counsel for Petitioner

CERTIFICATE OF SERVICE ON PATENT OWNER

Pursuant to 37 C.F.R. §§ 42.6(e)(4)(i) *et seq.* and 42.105(b), the undersigned certifies that a copy of the foregoing document with accompanying Exhibits 1001-1022, as listed in Petitioner's Exhibit List on pages vi-vii of this document, was caused to be served by Federal Express, costs prepaid, to Patent Owner at the following address of record:

73552 - Schwabe, Williamson & Wyatt/SFC
1211 SW Fifth Ave. Suite 1900
Portland, OR

Dated: August 2, 2024

/s/ Russell E. Levine, P.C.
Russell E. Levine, P.C. (Reg. No. 32,153)
KIRKLAND & ELLIS LLP
333 West Wolf Point Plaza
Chicago, Illinois 60654
Telephone: (312) 862-2000

Counsel for Petitioner