

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

TCL ELECTRONICS HOLDINGS LTD. and TCL INDUSTRIES HOLDINGS
CO., LTD.,

Petitioners,

v.

UNIVERSITY OF ROCHESTER
(INTELLECTUAL VENTURES I, LLC),

Patent Owner.

Case No. IPR2024-01245
U.S. Patent No. 7,089,443

PATENT OWNER'S PRELIMINARY RESPONSE

TABLE OF CONTENTS

PATENT OWNER’S EXHIBIT LIST	v
I. INTRODUCTION	1
II. SPECIFICATION OF THE '443 PATENT	2
A. Summary Of The Invention.....	2
B. Detailed Description	5
III. THE CLAIMED INVENTION HAD BEEN CONCEIVED AND ACTUALLY REDUCED TO PRACTICE BY FEBRUARY, 2002	5
A. Legal Principles	5
B. The Paper Authored By The Inventors And Published In February 2002 Demonstrates Conception By Providing A Full Written Description Of The Claimed Invention	6
C. The Simulation Described In The Paper Embodied The Invention In A Tangible Form.....	16
D. The Simulation Described In The Paper Showed That The Invention Worked For Its Intended Purpose	18
IV. CLAIMS OF ‘443 PATENT.....	19
V. CLAIM CONSTRUCTION.....	20
A. The Specification Limits The Claims To Globally-Asynchronous Systems	20
B. The Specification Disparages Systems That Impose Significant Restrictions On The Possible Choices Of Frequencies	21
VI. PRIOR ART RELIED ON IN THE PETITION.....	25
A. Iyer	25
B. Njolstad.....	25
C. Liljeberg.....	26
D. Soerensen.....	26
E. Hemani	27

F.	Moore	27
VII.	THE PETITION FAILS TO DEMONSTRATE A REASONABLE LIKELIHOOD OF SUCCESS ON THE MERITS	27
A.	Ground I Fails Because Iyer Published After The Invention Was Actually Reduced To Practice	27
B.	Ground IIA Fails Because Njolstad’s Pipelined Architecture Relies on Synchronized Clock Rates and Does Not Teach Independent Clock Frequency Generation as Required by the Challenged Claims	28
C.	Ground IIB Fails Because Liljeberg Does Not Cure the Shortcomings of Njolstad.....	34
D.	Ground IIIA Fails	34
1.	Soerensen Is Not Prior Art	35
2.	Soerensen’s Reliance on Rational Frequency Relationships Fails to Teach or Render Obvious the Independent Clock Frequency Generation Required by the Challenged Claims	42
E.	The Petition Fails to Demonstrate a Reasonable Likelihood of Success Under Ground IIIB.....	47
1.	Soerensen Is Not Prior Art	47
2.	Petitioner Fails To Demonstrate A Reasonable Likelihood Of Success In Combining Soerensen with Hemani.	47
a)	Hemani Made Clear That There Was No Expectation Of Success.....	48
b)	The Petition Failed To Prove A Reasonable Expectation Of Success .	52
3.	Ground IIIB Also Failed To Prove A Motivation To Combine.....	54
4.	Neither Soerensen Alone Nor In Combination With Hemani Teaches Or Inherently Discloses The Independent Clock Frequency Generation Required By The Challenged Claims	55
F.	Ground IIIC Fails Because Moore Does Not Cure the Shortcomings of Soerensen.	60
VIII.	SUMMARY OF ARGUMENTS ON THE MERITS	60
IX.	INSTITUTION SHOULD BE DENIED UNDER THE <i>FINTIV</i> FACTORS....	61

A.	Factor 1 Weighs Against Institution Because A Stay Has Not Been Granted In The Parallel Litigation And There Is No Evidence That A Stay Would Be Granted If This IPR Was Instituted	62
B.	Factor 2 Weighs Against Institution Because The Parallel Litigation Has A Trial Date Scheduled Well Before The Projected Statutory Deadline For The Board’s Final Written Decision	62
C.	Factor 3 weighs against institution because the parties will have invested significant effort in the Parallel Litigation by the time the Board’s institution decision is due.....	63
D.	Factor 4 weighs against institution because there is substantial overlap between the Parallel Litigation and this IPR	65
E.	Factor 5 weighs against institution because the patent owner is the same and the Petitioner are the same parties as in the Parallel Litigation.....	67
F.	Factor 6 weighs against institution because there are no other circumstances that favor institution, and the merits of Petitioner’s challenge is weak.....	68
X.	CONCLUSION	69

PATENT OWNER'S EXHIBIT LIST

EX2001	Declaration of Dr. David Kaeli
EX2002	Resume of Dr. David Kaeli
EX2003	Declaration of David H. Albonesi
EX2004	L. F. G. Sarmenta, G. A. Pratt, and S. A. Ward. Rational Clocking. In Proceedings of the International Conference on Computer Design, Austin, Tex., October 1995
EX2005	Semeraro et al, "Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scaling," High Performance Computer Architecture (HPCA), Feb. 2, 2002
EX2006	Docket from Parallel Litigation, Case No. 6:23-cv-00309 (WD Tex.)
EX2007	Claim Construction Order and Memorandum dated September 10, 2024 in Case No. 6:23-cv-00309 (WD Tex.), Document 44
EX2008	Scheduling Order dated March 5, 2024 in Case No. 6:23-cv-00309 (WD Tex.), Document 26
EX2009	United States District Courts — National Judicial Caseload Profile, Dec. 2023 Edition
EX2010	Defendants' Final Invalidity Contentions in in Case No. 6:23-cv-00309 (WD Tex.), dated September 23, 2024
EX2011	SimpleScalar Tutorial

I. INTRODUCTION

The Petition (Paper #2) challenges claims 1, 3, 4, 5, 8, 10 and 12 (“Challenged Claims”) of U.S. Patent No. 7,089,443 (“’443 Patent,” EX1001) under six grounds. Petition, at 5. The Petition fails to present a reasonable likelihood of success in connection with any of the grounds.

The ’443 Patent describes the design and implementation of a microprocessor architecture utilizing a globally asynchronous, locally synchronous (GALS) clocking architecture in which each domain includes a separate clock for generating a clock signal at a frequency for that domain. EX1001, 4:13-16; 40-45. Among other features, the frequency in each domain is dynamically changeable independently of the frequencies of the clock signals generated for other domains. *Id.*, claim 1. The transformative nature of this architecture has been widely recognized, including the receipt of a prestigious Test of Time Award in 2023. EX2003, ¶9.

On the merits, the Petition is deficient for at least the following reasons:

First, there is overwhelming documentary evidence that the claimed invention was conceived and actually reduced to practice by February 2002, predating at least the Iyer (EX1005) and Soerensen (EX1006) references upon which Grounds 1, III.A, III.B and III.C rely.

Second, the invention’s novel features, including separate and independently adjustable clock domains, are clearly distinguished from prior art such as Njolstad

and Soerensen that relied on restrictive synchronization techniques such as Sarmenta's "rational factor" frequency relationships.

Third, with respect to Ground III.B, the Petition failed to acknowledge, let alone address, Hemani's statement that running partitions at different speeds may cause "deadlock." EX1007, at page 6. Given this statement in Hemani, it is clear that there was neither a motivation or reasonable expectation of success in combining Hemani and Soerensen.

On the issue of discretionary denial, Petitioner waited almost a full year after service of the complaint in parallel patent infringement litigation to file this challenge. As a result, the parallel litigation (in which Petitioner is advancing the same prior art relied on in the Petition) has already reached an advanced stage, with a trial scheduled to occur long before any Final Written Decision would be issued in these proceedings. Based on these facts, and as explained more fully below, all of the *Fintiv* factors favor a discretionary denial of the Petition.

II. SPECIFICATION OF THE '443 PATENT

A. Summary Of The Invention

As explained in the Summary of the Invention of the '443 Patent:

... the present invention is directed to a multiple clock domain (MCD) microarchitecture, which uses a globally-asynchronous, locally-synchronous (GALS) clocking style. In an MCD microprocessor each functional block operates with a separately generated clock, and

synchronizing circuits ensure reliable inter-domain communication. EX1001, 4:13-19 (emphasis added).

As explained next in the Summary of the Invention, the inventors identified several "potential" advantages of the MCD design, including "lower energy:"

Although the inter-domain synchronization increases the number of clock cycles required to run a given application, an MCD microprocessor affords a number of potential advantages over a singly clocked design: ...

The independence of each local domain clock implies no global skew requirement...

Using separate voltage inputs, external voltage regulators, and controllable clock frequency circuits in each clock domain allows for finer grained dynamic voltage and frequency scaling, and thus lower energy, than can be achieved with single clock, single-core-voltage systems. EX1001, 6:22-45. (emphasis added).

Next in the Summary of the Invention, the inventors explained that they implemented their MCD design and found that it achieved significant performance advantages:

In the present application, we describe an initial implementation of an MCD microprocessor that is a straightforward extension of a singly-clocked synchronous dynamic superscalar design. By accurately modeling inter-domain synchronization, we characterize the performance and energy costs of the required synchronization circuitry. We then explore the potential benefits of per-domain dynamic voltage and frequency scaling. Our results demonstrate a 20% average improvement in energy-delay product for a set of benchmarks that includes both compute and memory-bound applications. ...

Performance results for applications drawn from standard benchmark suites suggest that the division of the processor into

multiple domains incurs an average baseline performance cost of less than 4%. At the same time, by scaling frequency and voltage in different domains dynamically and independently, we can achieve an average improvement in energy-delay product of nearly 20%. By contrast, global voltage scaling to achieve comparable performance degradation in a singly clocked microprocessor achieves an average energy-delay improvement of only 3%. EX1001, 4:51-5:14 (emphasis added).

At the end of the Summary of The Invention, the inventors noted that "the invention" was described in an earlier paper by the inventors:

The following paper describes the invention and is hereby incorporated by reference in its entirety into the present disclosure: Semeraro et al, "Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scaling," *High Performance Computer Architecture (HPCA)*, Feb. 2, 2002. EX1001, 5:22-27.

The above referenced paper is referred to hereinafter as "the Paper," and attached as EX2005. As explained more fully in Section III, *infra*, the salient portions of the '443 Patent specification were lifted almost verbatim from the Paper. The findings of the inventors set forth in the Paper and reproduced in the '443 Patent were nothing short of groundbreaking. In 2023, the Paper describing this work was awarded the Technical Committee on Computer Architecture's 2022 Test of Time Award at the 29th IEEE International Symposium on High-Performance Computer Architecture. EX2003, ¶9. The Test of Time Award "recognizes the most influential papers published in prior sessions of the International Symposium of High Performance Computer Architecture (HPCA) each of whom have had significant impact in the field." EX2003, ¶10. The Paper has been cited 575 times since it was

published.

B. Detailed Description

The Detailed Description of '443 Patent distinguished the MCD processor disclosed in the '443 Patent from prior synchronization schemes that restricted the relative frequencies of the clocks:

Some synchronization schemes restrict the ... relative frequencies of the clocks, thereby eliminating the need for hardware arbitration [27]. Unfortunately, these schemes impose significant restrictions on the possible choices of frequencies. ... EX1001, 6:32-36 (emphasis added).

As discussed more fully below, reference [27] referred to in the above quotation corresponded to a paper by Sarmenta (EX2004) that described systems that placed significant restrictions on the possible choices of frequencies. As Dr. Kaeli explains “[t]he ‘443 invention overcomes these limitations, offering independence in clock frequency generation, allowing each frequency domain to adopt their own independent frequency, offering near infinite resolution in terms of the clock frequency selected for each domain.” EX2003, ¶38.

III. THE CLAIMED INVENTION HAD BEEN CONCEIVED AND ACTUALLY REDUCED TO PRACTICE BY FEBRUARY, 2002

A. Legal Principles









A patent owner can antedate an asserted prior art reference by showing conception and actual reduction to practice prior to the date of the reference.

Medtronic, Inc. v. Teleflex Life Scis. Ltd., 86 F.4th 902, 906 (Fed. Cir. 2023). "Conception can be proven with evidence that a full description" of the invention existed "in written materials disclosed to persons" in the United States. *Scott v. Koyama*, 281 F. 3d 1243, 1247 (Fed. Cir. 2002). "Reduction to practice in the United States requires that the invention be embodied in tangible form in the United States," and "that the invention, as conceived, was shown to work for its intended purpose." *Id.*; *In re Steed*, 802 F. 3d 1311, 1316 (Fed. Cir. 2015). In connection with the requirement that the invention be embodied in a tangible form, it is well settled that a computer program with executable instructions stored in memory (e.g., on a computer disk) is embodied in a tangible form. *In re Beauregard*, 53 F.3d 1583 (Fed. Cir.1995).

B. The Paper Authored By The Inventors And Published In February 2002 Demonstrates Conception By Providing A Full Written Description Of The Claimed Invention

As noted above, the '443 Patent cited and incorporated by reference the Paper (i.e., Semeraro et al, "Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scaling," High Performance Computer Architecture (HPCA), Feb. 2, 2002). EX1001, 5:22-27; EX2005. The Paper was published in *Proceedings Eighth International Symposium on High Performance Computer Architecture*, pp. 29-40 in February, 2002. EX2003, ¶5. The authors of the Paper coincide identically with the named inventors on the '443 Patent.

Id. As shown below, with the exception of reference numerals, all of the figures contained in the '443 Patent are substantively identical to corresponding figures in the Paper. EX2003, ¶18.

Figure from '443 Patent	Corresponding Figure in Paper
 Figure 1	 Figure 1. Multiple clock domain processor block diagram
 Figure 2	 Figure 2. Queue structure
 Figure 3	 Figure 3. Full flag
 Figure 4	 Figure 4. Synchronization timing




Figure 5




Figure 5. Performance degradation results




Figure 6




Figure 6. Energy savings results




Figure 7




Figure 7. Energy-delay improvement results






Figure 8A





Given the overlap in the figures, it should not be surprising that the wording of the '443 Patent closely tracks that of the Paper. For example, the first three paragraphs of the Introduction section of the Paper are reproduced virtually verbatim as the first three paragraphs in the Background of the '443 Patent (EX1001, 1:25-2:14; EX2005, at pages 1-2), and the 31 references listed at the end of the Paper are

listed identically at the end of the Background of the '443 Patent. EX1001, 2:17-4:4; EX2005, at 11-12. In addition, as shown below, the salient aspects of the Summary of the Invention of the '443 Patent (EX1001, 4:6-5:21) are taken verbatim from the Introduction section of the Paper (EX2005, 2):

Summary of Invention of '443 Patent	Corresponding Language In Introduction of Paper
It is another object of the invention to provide an approach that allows for aggressive future frequency increases, maintains a synchronous design methodology, and exploits the trend towards making functional blocks more autonomous. To achieve the above and other objects, the present invention is directed to a multiple clock domain (MCD) microarchitecture, which uses a globally-asynchronous, locally-synchronous (GALS) clocking style. In an MCD microprocessor each functional block operates with a separately generated clock, and synchronizing circuits ensure reliable inter-domain communication. Thus, fully synchronous design practices are used in the design of each domain. Although the inter-domain synchronization increases the number of clock cycles required to run a given application, an MCD microprocessor affords a number of potential advantages over a singly clocked design: ...	An approach that allows for aggressive future frequency increases, maintains a synchronous design methodology, and exploits the trend towards making functional blocks more autonomous, is ... a <i>multiple clock domain (MCD)</i> microarchitecture, which uses a <i>globally-asynchronous, locally-synchronous (GALS)</i> clocking style. In an MCD microprocessor each functional block operates with a separately generated clock, and synchronizing circuits ensure reliable inter-domain communication. Thus, fully synchronous design practices are used in the design of each domain. Although the inter-domain synchronization increases the number of clock cycles required to run a given application, an MCD microprocessor affords a number of potential advantages over a singly clocked design: ...

Using separate voltage inputs, external voltage regulators, and controllable clock frequency circuits in each clock domain allows for finer grained dynamic voltage and frequency scaling, and thus lower energy, than can be achieved with single clock, single-core-voltage systems. ... In the present application, we describe an initial implementation of an MCD microprocessor that is a straightforward extension of a singly-clocked synchronous dynamic superscalar design.	Using separate voltage inputs, external voltage regulators, and controllable clock frequency circuits in each clock domain allows for finer grained dynamic voltage and frequency scaling, and thus lower energy, than can be achieved with single clock, single-core-voltage systems. ... In this Paper, we describe an initial implementation of an MCD microprocessor that is a straightforward extension of a singly-clocked synchronous dynamic superscalar design.
--	---

Below is a mapping demonstrating a description of each element of claims 1-7 of the '443 Patent in the Paper. While not shown, the mappings for claims 8, 9, 10, 11 and 12 of the '443 Patent correspond to the mappings below for claims 1, 4, 5, 6 and 3, respectively.

Claim 1

<u>Claim Element</u>	<u>Disclosure in Paper</u>
A multiple-clock-domain microprocessor comprising:	"An approach that allows for aggressive future frequency increases, maintains a synchronous design methodology, and exploits the trend towards making functional blocks more autonomous, is a <i>multiple clock domain</i> (MCD) microarchitecture, which uses a globally-asynchronous, locally-synchronous (GALS) clocking style. In

	an <u>MCD microprocessor</u> each functional block operates with a separately generated clock, and synchronizing circuits ensure reliable inter-domain communication. ... In this Paper, we describe an initial implementation of an MCD microprocessor ..." EX2005, 2 (emphasis added).
a plurality of domains;	"We describe an alternative approach, which we call a Multiple Clock Domain (MCD) processor, in which the chip is divided into <u>several (coarse-grained) clock domains</u>, within which independent voltage and frequency scaling can be performed." EX2005, Abstract (emphasis added).
for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains; and	"We describe an alternative approach, which we call a Multiple Clock Domain (MCD) processor, in which the chip is divided into several (coarse-grained) clock domains, <u>within which independent voltage and frequency scaling can be performed</u>." EX2005, Abstract (emphasis added). "Figures 8 and 9 both illustrate the value of using <u>different frequencies in different clock domains: by controlling these frequencies independently</u> we can maintain the required frequency in domains that are critical to performance, while aggressively scaling those domains that are less performance-critical." EX2006, 10 (emphasis added).
for each of the plurality of domains, a voltage input for receiving a voltage	"We describe an alternative approach, which we call a Multiple Clock Domain

which is dynamically changeable independently of the voltages applied to said others of the plurality of domains.	(MCD) processor, in which the chip is divided into several (coarse-grained) clock domains, <u>within which independent voltage and frequency scaling can be performed.</u>" EX2005, Abstract (emphasis added). "[B]y <u>scaling</u> frequency and <u>voltage in different domains dynamically and independently</u>, we can achieve an average improvement in energy-delay product of nearly 20%." EX2005, 11 (emphasis added).
---	---

Claim 2

<u>Claim Element</u>	<u>Disclosure in Paper</u>
2. The microprocessor of claim 1, wherein each said clock comprises a phase-locked loop,	"The global clock distribution network is greatly simplified, requiring only the distribution of the externally generated clock to <u>the local Phase Lock Loop (PLL) in each domain.</u> " EX2005, 2 (emphasis added).
and wherein the microprocessor further comprises means for receiving an externally generated clock signal and for supplying the externally generated clock signal to each said phase-locked loop.	"The global clock distribution network is greatly simplified, requiring only the <u>distribution of the externally generated clock to the local Phase Lock Loop (PLL) in each domain.</u> " EX2005, 2 (emphasis added).

Claim 3

<u>Claim Element</u>	<u>Disclosure in Paper</u>
----------------------	----------------------------

3. The microprocessor of claim 1, wherein there are at least four of said domains.	"We propose <u>four clock domains</u> , corresponding to the front end (including L1 instruction cache), integer units, floating point units, and load-store units (including L1 data cache and L2 cache)." EX2005, Abstract (emphasis added).
--	--

Claim 4

<u>Claim Element</u>	<u>Disclosure in Paper</u>
4. The microprocessor of claim 1, wherein the microprocessor is programmed to determine a slack in processing in one of the domains and to reduce the clock frequency and the voltage in said one of the domains to reduce the slack.	"Whenever an event in the dependence DAG has two or more incoming arcs, it is possible—in fact likely—that one arc will constitute the critical path and that the others will have “slack”. This slack indicates that the previous operation completed earlier than necessary. If all of the outgoing arcs of an event have <u>slack</u>, then we have an opportunity (assuming zero-cost scaling) to save energy by performing the event at a <u>lower frequency and voltage</u>." EX2005, 6 (emphasis added). "In particular, the slack associated with L1 data cache misses often allows our reconfiguration tool to scale the integer and floating point domains without significantly impacting overall performance (due to the fact that the available ILP is not sufficient to completely hide the miss latency), even when the utilization for these domains is high." EX2005, 9.

Claim 5

<u>Claim Element</u>	<u>Disclosure in Paper</u>
5. The microprocessor of claim 1, further comprising a queue for communication between at least two of the domains.	"The general <u>queue structure that we use for inter-domain communication</u> is shown in Figure 2." EX2004, 3 (emphasis added).

Claim 6

<u>Claim Element</u>	<u>Disclosure in Paper</u>
6. The microprocessor of claim 5, wherein the queue has a Full flag and an Empty flag, and wherein the microprocessor is programmed to prevent a write to the queue when the Full flag is asserted, until the Full flag is deasserted; and to prevent a read from the queue when the Empty flag is asserted, until the Empty flag is deasserted.	The general queue structure that we use for inter-domain communication is shown in Figure 2. The assertion of the <u>Full flag</u> indicates to the producer that it can <u>no longer write to the queue until the flag is deasserted</u> ..., while the <u>Empty flag</u> when asserted indicates that there is no valid data for the consumer to read from the queue. The consumer <u>waits until Empty is deasserted before reading again.</u> " EX2005, 3 (emphasis added).

Claim 7

<u>Claim Element</u>	<u>Disclosure in Paper</u>
7. The microprocessor of claim 6, wherein the queue is implemented as a dual-ported SRAM.	"This concurrency requires a dual-ported SRAM structure where simultaneous read and write cycles are allowed to <i>different</i> SRAM cells." EX2005, 4.

Given the significant overlap between the Specification of the '443 Patent and the Paper, as well as explicit support in the Paper for each element recited in the claims, there can be no dispute that the Paper demonstrates conception by providing a full description of the invention. *Scott*, 281 F. 3d 1243, 1247; EX2003, ¶17 ("The microarchitecture design described in the Paper is the same design implemented by the multiple-clock-domain microprocessor described and claimed in the '443 Patent") and ¶24 ("my Co-Inventors and I had conceived of the invention ... before the February 2002 publication of the Paper.")

C. The Simulation Described In The Paper Embodied The Invention In A Tangible Form

After describing, in detail, the claimed multiple clock domain microarchitecture (including all of the aspects set forth in the claim charts in Section III(B) above), the Paper described a detailed simulation that implemented the multiple clock domain microarchitecture prior to the publication of the Paper in February, 2002. *See*, EX2005, 5-8; EX2003, ¶16 ("The simulation described in the Paper was completed before publication of the Paper in February, 2002"). The simulation described in Paper used the SimpleScalar tool. As explained by Dr. Kaeli:

In 2002, SimpleScalar was a computer simulator that was commonly used by the computer architecture community to evaluate the performance and correctness of new architectural designs. Among other features, the SimpleScalar tool included a performance simulator that implemented the device at the microarchitecture level, including data paths, memory architecture and control units of the computer device. *See*, e.g.,

Paper at 5 (discussing use of SimpleScalar to implement modified "microarchitecture of the Alpha 21264 microprocessor" where the centralized Register Update Unit was split "into separate reorder buffer (ROB), issue queue, and physical register file structures.") The SimpleScalar tool provided for execution-driven simulations, where the device being implemented ran as an instrumented program on a host. EX2001, ¶42 (emphasis added).

Significantly, "[w]hen the SimpleScalar tool executed the simulation described in the Paper on a host, the host executing the simulation embodied all of the elements of claim 1 in tangible form." The same can be said for each of the other claims in the '443 Patent, namely, that the host executing the simulation embodied all of the elements of the claim in tangible form." EX2001, ¶43. By embodying all of the elements of the claims in tangible perform, execution of the simulation described in the Paper constituted an actual reduction to practice of the claimed invention by February, 2002. *Scott*, 281 F. 3d 1243, 1247 ("[r]eduction to practice ... requires that the invention be embodied in tangible form.")

Moreover, "separate and apart from the host used to implement the device during the simulation, the computer code (standing alone) used to implement the simulation embodied all of the elements of claim 1 (and the other claims of the '443 patent) in tangible form." EX2001, ¶43; *see, e.g., In re Beauregard*, 53 F.3d 1583 (computer program with executable instructions stored in memory (e.g., on a computer disk) is embodied in a tangible form). Thus, the existence of the computer

code used to run the simulation provides an independent basis for demonstrating that the invention claimed in the '443 was actually reduced to practice by February 2002.

D. The Simulation Described In The Paper Showed That The Invention Worked For Its Intended Purpose

The Introduction of the Paper described several "potential" advantages of the MCD design, including "lower energy." As explained in the Paper:

Although the inter-domain synchronization increases the number of clock cycles required to run a given application, an MCD microprocessor affords a number of potential advantages over a singly clocked design: ...

- Using separate voltage inputs, external voltage regulators, and controllable clock frequency circuits in each clock domain allows for finer grained dynamic voltage and frequency scaling, and thus lower energy, than can be achieved with single clock, single-core-voltage systems. EX2005, 2 (emphasis added).

As explained in the Paper, the performance results of the simulation demonstrated that the MCD design achieved the advantages that the inventors sought to recognize:

Performance results for applications drawn from standard benchmark suites suggest that the division of the processor into multiple domains incurs an average baseline performance cost of less than 4%. At the same time, by scaling frequency and voltage in different domains dynamically and independently, we can achieve an average improvement in energy-delay product of nearly 20%. By contrast, global voltage scaling to achieve comparable performance degradation in a singly clocked microprocessor achieves an average energy-delay improvement of only 3%. EX2005, 11.

Thus, the simulation proved that the intended purpose of the MCD design—lower energy—had been achieved. In other words, the simulation results proved that the invention worked for its intended purpose. EX2001, ¶44 ("[O]thers working in the field reading the Paper would have understood that the device implemented in the simulation was shown to work for its intended purpose"); *Scott*, 281 F. 3d 1243, 1247 (Fed. Cir. 2002) (reduction to practice requires "that the invention, as conceived, was shown to work for its intended purpose").

In summary, the Paper taken together with the declaration evidence included herewith, establishes that the invention claimed in the '443 Patent was conceived and actually reduced to practice by the time the Paper was published in February, 2002. Thus, any reference dated after that time is not prior art.

IV. CLAIMS OF '443 PATENT

The claims of the '443 Patent capture the MCD microprocessor architecture that the inventors designed and implemented. Claim 1 is representative and provides:

A multiple-clock-domain microprocessor comprising:
a plurality of domains;
for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains; and
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.

V. CLAIM CONSTRUCTION

The claim limitations at issue must be construed under *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005). Claim language is not read in isolation. *Phillips*, 415 F.3d at 1315. Rather, the claim language, when read in light of the Specification, elucidates the meaning of the claim language.

A. The Specification Limits The Claims To Globally-Asynchronous Systems

When a patentee describes the features of the "present invention," it is well-settled that "this description limits the scope of the invention." *Pacing Techs., LLC v. Garmin Int'l, Inc.*, 778 F.3d 1021 (Fed. Cir. 2015). In this regard, the "Summary of the Invention" of the '443 Patent states the "present invention" uses a globally-asynchronous clocking style:

... the present invention is directed to a multiple clock domain (MCD) microarchitecture, which uses a globally-asynchronous, locally-synchronous (GALS) clocking style. In an MCD microprocessor each functional block operates with a separately generated clock, and synchronizing circuits ensure reliable inter-domain communication. ... an MCD microprocessor affords a number of potential advantages over a singly clocked design: ...

The independence of each local domain clock implies no global skew requirement... EX1001, 4:13-31 (emphasis added).

Thus, the claim language reciting "for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals

generated for others of the plurality of domains," requires that the separate clocks in the plurality of domains be globally-asynchronous. As Dr. Kaeli explains, the "Patent '443 overcomes earlier limitations of GALS systems by developing novel methods to support using independent clock and independent voltage controls, overcoming issues related to inter-domain synchronization and communication." EX2001, ¶39.

B. The Specification Disparages Systems That Impose Significant Restrictions On The Possible Choices Of Frequencies

The law is clear and well-settled that claims should not be read so broadly as to encompass a prior art structure distinguished or disparaged in the specification. *Indivior Inc. v. Dr. Reddy's Labs*, 752 Fed. Appx. 1024 (Fed. Cir. 2018) (finding that claims did not cover solely using conventional top air drying to produce films because the specification disclaimed such drying by disparaging its use); *SciMed Life Sys. v. Advanced Cardiovascular Sys.*, 242 F.3d 1337 (Fed. Cir. 2001), *citing*, *Tronzo v. Biomet, Inc.*, 156 F.3d 1154, 1159 (Fed. Cir. 1998) (Specification distinguished prior art as inferior and touted advantages of a conical shaped cup for use in an artificial hip device; "such statements make clear that the '589 patent discloses *only* conical shaped cups and nothing further"); *Ekchian v. Home Depot, Inc.*, 104 F.3d 1299, 1304, 41 U.S.P.Q.2D (BNA) 1364, 1368 (Fed. Cir. 1997) ("Since, by distinguishing the claimed invention over the prior art, an

applicant is indicating what the claims do not cover, he is by implication surrendering such protection.”).

As recognized by the Federal Circuit, even “where the language of the claims, read without reference to the specification, might be considered broad enough to encompass a feature, that feature is deemed to be outside the reach of the claims where the specification makes clear that the invention does not include the particular feature.” *Chi. Bd. Options Exch., Inc. v. Int’l Sec. Exch., LLC*, 677 F.3d 1361 (Fed. Cir. 2012), *citing*, *Honeywell Int’l, Inc. v. ITT Indus., Inc.*, 452 F.3d 1312, 1319 (Fed. Cir. 2006) (quoting *SciMed*, 242 F.3d 1337, 1341).

Here, the specification of the '443 Patent disparages prior art synchronization schemes that restrict the relative frequencies of the clocks:

Some synchronization schemes restrict the ... relative frequencies of the clocks, thereby eliminating the need for hardware arbitration [27]. Unfortunately, these schemes impose significant restrictions on the possible choices of frequencies. ... EX1001, 6:32-36 (emphasis added).

The Specification explains that Reference [27] is “L. F. G. Sarmenta, G. A. Pratt, and S. A. Ward. Rational Clocking. In Proceedings of the International Conference on Computer Design, Austin, Tex., October 1995.” EX1001 3:54-56. (EX2004) (hereinafter, “Sarmenta”).

Sarmenta describes an “approach [that] maintains a known phase relationship between clocks whose frequencies are related by a rational factor, and exploits the predictability of their relative phases to algorithmically time communications

without run-time arbitration contests.” EX2004, Abstract. The “rational factor” refers to the ratio of two small integers (M/N), representing the relationship between the clock frequencies of two clocked digital subsystems. EX2001, ¶36. This rational relationship ensures periodic alignment of the clocks, creating predictable moments when the active edges of both clocks coincide (and are therefore synchronized). EX2001, ¶36. For instance, if one clock operates at frequency (f_m) and the other at other at (f_n), their frequencies satisfy the relationship ($f_m/f_n = M/N$). This periodic alignment defines a “coincidence cycle,” during which the relative positions of clock edges are synchronized, enabling efficient transfers between subsystems. EX2001, ¶36.

The specification makes clear that Sarmenta is an example of synchronization scheme that impose “significant restrictions on the possible choices of frequencies, and clearly disparages and disclaims/disavows these types of restrictive schemes.” EX1001, 6:32-36. To overcome the limitations of systems like Sarmenta, the '443 Patent explains that the invention provides “[t]he designers of each domain [to be] no longer constrained by the speeds of critical paths in other domains, affording them greater freedom in each domain to optimize the tradeoffs among clock speed, latency, and the exploitation of application parallelism via complex hardware structures.” EX1001, 6:34-39. As Dr. Kaeli explains, “[t]he ‘443 invention overcomes these limitations offering independence in clock frequency generation,

allowing each frequency domain to adopt their own independent frequency, offering near infinite resolution in terms of the clock frequency selected for each domain.” EX2001, ¶38.

Accordingly, the claims should do not be read to cover schemes like Sarmenta’s where frequencies are related by a “rational factor” that significantly restricts the possible choices of frequencies to those related by the "rational factor." In addition to being expressly disavowed, such systems are outside of the scope of the "present invention" which is limited to systems that use a globally-asynchronous clocking style (*id.*, EX1001, 4:13-19), because in systems like Sarmenta where the frequencies are related by a “rational factor,” there will be a periodic alignment during which the relative positions of clock edges are synchronized. EX2001, ¶37. Thus, the claim language reciting "for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains " does not cover synchronized systems like Sarmenta that maintain global synchronization by restricting the possible choice of frequencies to those related by a "rational factor."

VI. PRIOR ART RELIED ON IN THE PETITION

A. Iyer

Iyer is the sole reference cited in Ground I. According to the Petition, Iyer was published in the IEEE/ACM International Conference on Computer Aided Design, 2002. The conference occurred 10-14 November 2002, in San Jose, CA, USA. Iyer states “Semeraro et al. [13] have also used a FIFO-based communication mechanism for their multiple clock domain implementation.” EX1005 at §3.3. Iyer’s identifies Semeraro et al. [13] as “G. Semeraro, G. Magklis, R. Balasubramonian, D. H. Albonesi, S. Dwarakadas, and M. L. Scott, “Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scaling,” in Proc. Intl. Symp. on High Performance Computer Architecture (HPCA), 2002,” which corresponds to the Paper by the inventors discussed above. EX1005, §References; EX2005.

B. Njolstad

Njolstad is the sole reference cited in Gound IIA. The paper presents a model for a wideband radio access system that uses varying clocks, voltages, and device technologies to implement flow control and reduce power usage. The system is structured as a sequence of stages or buffers, where clock rates and voltages are dynamically adjusted to ensure that each stage processes data at a rate compatible with the next, preventing overruns. EX1005, §IV. D, EX2001, ¶ 55-56.

C. Liljeberg

Liljeberg is combined with Njolstad to form Petitioner's Gound IIB. Liljeberg, is a paper titled "Asynchronous Interface for Locally Clocked Modules in ULSI Systems," that introduces an architecture that facilitates communication between locally clocked modules using interface units. The design employs elastic burst-mode data transfer and features interface units equipped with dual-ported memory to manage data and control signal transfer between asynchronous and synchronous domains. EX2001, ¶ 58-59.

D. Soerensen

Soerensen is the sole reference cited in Gound IIIA. Soerensen, titled "Dynamic Voltage Control Method and Apparatus," outlines a method for dynamically selecting a clock frequency for a microprocessor and determining the voltage required to support the chosen frequency. Figure 2 illustrates how each clock domain relies on a system clock signal, with clock signals generated and distributed via select, divider, and gating logic through multiple clock trees. Significantly, in Soerensen, "the clock signals distributed throughout the processor are synchronized." EX1006, [0033] (emphasis added). Soerensen further explains clock control mechanisms that restrict the ability to adjust clock frequencies within various domains to frequencies achievable by dividing the system clock signal by factors such as 2 or 4. EX1006, ¶0038, EX2001, ¶ 48-50.

E. Hemani

Hemani is combined with Soerensen to form Petitioner's Gound IIIB. Hemani, in the paper titled "Lowering Power Consumption in Clock by using Globally Asynchronous Locally Synchronous Design Style," focuses on optimizing power consumption in a design where all system blocks are clocked at the same frequency." EX1007, 5 (emphasis added). Hemani's experiments demonstrated that when all SBs are clocked at the same frequency, the power consumed by the clock function decreases as a function of partitioning. EX1007, Fig. 4. . EX2001 at 55-56.

F. Moore

Moore is combined with Soerensen to form Petitioner's Gound IIIC. In the paper titled "Channel Communication Between Independent Clock Domains," Moore presents methods for achieving reliable and low-latency communication across clock domains using a combination of clock pausing and self-calibrating delay lines. The approach is designed for system-on-a-chip architectures. EX2001, ¶ 53-54.

VII. THE PETITION FAILS TO DEMONSTRATE A REASONABLE LIKELIHOOD OF SUCCESS ON THE MERITS

A. Ground I Fails Because Iyer Published After The Invention Was Actually Reduced To Practice

Ground I relies on Iyer (EX1005), which the Petition alleges was publicly available "no later than November 10, 2002." Pet., 3. As discussed in Section III

above, the invention claimed in the '443 Patent was both conceived and actually reduced to practice prior to February 2002, which predates the alleged publication date of Iyer. Iyer expressly recognizes that it post-dates the work of the inventors of the '443 Patent by citing the Paper of the inventors as a reference. See, Ex 1005, §3.3 (“Semeraro et al. [13] have also used a FIFO-based communication mechanism for their multiple clock domain implementation.”)

Since Iyer is not prior art, Ground I fails.

B. Ground IIA Fails Because Njolstad’s Pipelined Architecture Relies on Synchronized Clock Rates and Does Not Teach Independent Clock Frequency Generation as Required by the Challenged Claims

The Petition relies on Figure 3 of Njolstad to argue that it discloses “separately generating a clock signal at a frequency for each domain,” with each frequency being “dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains.” Pet., 29–31. However, a closer examination of Njolstad, particularly Figure 3 (reproduced below), reveals a fundamentally different structure and operation.





Figure 3: A model of a wideband radio access system using multiple supply voltages and logic level shifters.

As explained by Dr. Kaeli, Figure 3 depicts a pipelined architecture for a wireless wideband radio access system, which relies on synchronized clock rates to maintain system performance and avoid processing failures. This *interdependence* between clock rates across processing stages is explicitly described in Njolstad and underscores the system's inability to support the dynamic *independence* of clock frequencies as required by the Challenged Claims. EX2001, ¶56.

The Petition alleges that “[t]he clocks for separately generating a clock signal at a frequency for each domain are again visible in Figure 3, as fCLK1, fCLK2, fCLK3, and fCLK4, respectively.” Pet., 29-30. However, Njolstad’s Figure 3 shows a pipelined architecture for a wireless wideband radio access system. As explained by Dr. Kaeli, in pipelined architecture, input is received, and the input is then serially

processed through a series of stages to generate an output that is much like an assembly line in a factory. Accordingly, each respective processing stage receives input from a prior stage, performs processing, and generates an output that is subsequently used by another stage. EX2001, ¶57. Dr. Kaeli's explanation of a pipelined architecture is consistent with Njolstad's description of Fig. 3. For example, Njolstad states, "in Figure 3 many tasks are performed at a fixed rate (or at a temporally fixed rate)," and "[e]ach processing element ... receives an input data set and produces an output data set at this fixed rate." EX1005, §IV. D. This alignment between Dr. Kaeli's explanation and Njolstad's description of Figure 3 confirms that Njolstad's pipelined architecture operates through *interdependent* (synchronized) processing stages, each constrained by fixed-rate requirements, rather than *independently* generated clock frequencies as required by the Challenged Claims. EX2001, ¶62.

Njolstad emphasizes the critical importance of maintaining *synchronized* rate requirements between processing elements in a pipelined architecture to ensure balanced operation and prevent system failures. EX2001, ¶55-57. For example, Njolstad explains that "[t]he fixed rate requirement or the fixed average rate requirement can dynamically be configured, by the special control signal RATE-REQUIREMENT_i" as shown in Figure 5 (reproduced below). EX1005, at §IV. D.



Dr. Kaeli explains that the rate requirement for each processing element in the pipeline is critically important because the rate at which data is generated in one stage cannot be faster than the data consumed in the subsequent stage. Otherwise, Njolstad's pipelined architecture for a wireless wideband radio access system would suffer from the processing equivalent of Lucy and Ethel from I Love Lucy comically failing on the production line at the chocolate factory. EX2001, ¶57

In this famous scene (depicted below), the upstream production rate of chocolates is increased, and Lucy and Ethel comically can't keep up with the rate at which chocolates are passing their processing station.



Just like Lucy and Ethel, in a computer-pipelined architecture, if the processing rate is increased in an upstream processor, there needs to be a corresponding increase in the downstream processing element. Otherwise, the processor malfunctions, and the results are less comical than in the case of I Love Lucy. EX2001, ¶57

Njolstad appreciated the fact that pipelined architectures require a balance between the upstream and downstream process by stating, “[f]rom 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.” EX1005, §IV. D . In fact, the Petition quotes Njolstad as explaining, “[t]he 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.” Pet., 31. (internal markings and emphasis omitted).

Therefore, Njolstad does not teach a globally-asynchronous architecture including a plurality of domains each of which “separately generat[es] a clock signal... the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains,” as recited in the Challenged Claims. EX2001, ¶61. Instead, as shown in Fig. 3 (reproduced above), increasing Fclk1 from a steady state value will cause an increase in Fclk2. Otherwise, stage two in the pipeline would overflow like Lucy and Ethel. Similarly, decreasing Dclk1 would cause a corresponding decrease in Fclk2. Otherwise, stage 2 in the pipeline would be wasting energy. Accordingly, Fclk2 is dependent on Fclk1, whereas the Challenged Claims expressly require that the clock frequency in each domain be “dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains.” EX2001, ¶62.

By conflating Njolstad’s pipelined architecture with the independent clock generation required by the Challenged Claims, the Petition mischaracterizes Njolstad’s teachings. As discussed above, Njolstad’s pipelined architecture inherently depends on coordinated clock signals to balance the flow of data between processing stages. Any deviation from this coordinated balance would lead to system

instability, as vividly illustrated by Dr. Kaeli's analogy to the classic I Love Lucy chocolate factory scene.

This critical flaw in the Petitioner's argument establishes that Njolstad fails to teach or suggest the independent clock frequency generation required by the Challenged claims. Accordingly, for at least these reasons, the Petitioner has failed to meet its burden regarding Ground 2A.

C. Ground IIB Fails Because Liljeberg Does Not Cure the Shortcomings of Njolstad

The Petition does not allege, nor does Liljeberg teach, "separately generating a clock signal... the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains," as recited in the independent Challenged Claims. Accordingly, the Petitioner has failed to meet its burden with regard to Ground IIB for at least the reasons discussed with regard to Ground IIA above.

D. Ground IIIA Fails

Ground IIIA fails because Soerensen does not qualify as prior art against the '443 Patent. Even if it were prior art, Soerensen relies on rational frequency relationships that inherently limit clock domain independence, a restriction that the '443 Patent specifically improves upon. Accordingly, Ground IIIA cannot support a finding of invalidity.

1. Soerensen Is Not Prior Art

As explained above, the invention claimed in the '443 Patent was conceived and actually reduced to practice by February, 2002. Soerensen was filed on August 29, 2002, which is after the invention date of the '443 Patent. While Soerensen claims priority to Provisional Application No. 60/315,655, filed August 29, 2001 (EX1022) (Soerensen Provisional Application), the passages from Soerensen relied on in the Petition are not entitled to the benefit of the filing date of the Soerensen Provisional Application because the subject matter was not disclosed in the manner provided by 35 U.S.C. § 112.

As explained by the Federal Circuit, “A reference patent is only entitled to claim the benefit of the filing date of its provisional application if the disclosure of the provisional application provides support,” in compliance with § 112 for both the claims in the reference patent and the disclosures from the reference patent relied upon for invalidity. *See, Dynamic Drinkware, LLC v. Nat'l Graphics, Inc.*, 800 F.3d 1375, 1382 (Fed. Cir. 2015). In order to satisfy the written description requirement of §112, “one skilled in the art, reading the original disclosure, must immediately discern” the material at issue. *Purdue Pharma L.P. v. Faulding Inc.*, 230 F.3d 1320, 1323 (Fed. Cir. 2000) (emphasis added). “[A] description that merely renders the invention obvious does not satisfy the [written description] requirement.” *Ariad*

Pharms., Inc. v. Eli Lilly & Co., 598 F.3d 1336, 1352 (Fed. Cir. 2010) (*en banc*); see also *Lockwood v. Am. Airlines, Inc.*, 107 F.3d 1565, 1571-72 (Fed. Cir. 1997).

The Soerensen Provisional Application is 257 pages long, and consists of a collection of disparate technical documents that were appended together. The Soerensen Provisional Application begins with the technical Specification for a "Digital Baseband Processor," (EX1022, at 1-54), followed by a document entitled "Monaco MCU/DSP - DMA Controller Module Specification," (*id.*, at 55-102), an invention disclosure entitled "Monaco Clock Control - short r01.doc," (*id.*, at 103-106), another invention disclosure entitled "Simplified Dynamic Voltage Control Intellectual Property Disclosure," (*id.*, at 107-112), a document entitled "High-Speed Data Logging and Functions (HSL)," (*id.*, at 113-117), a document entitled "Monaco Patent: PLL - Clock before Lock Technical Note," (*id.*, at 118-123), a document entitled "eGSP Patent-able items?" (*id.*, at 124-136), a document entitled "Monaco Improving MCU performance Technical note," (*id.*, at 137-142), a document entitled "Cache Support for Real-time Multithreading IP Disclosure," (*id.*, at 143-147), and a document entitled "Monaco Digital Baseband Design Specification" (*id.*, at 148-257).

The Petition at note 2 argues that Soerensen's claims are supported by the Provisional Patent Application, but provides little evidence that such is the case. For example:

- At page 45, the Petition quotes Soerensen 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’)—but provides no supporting citation for this subject matter in the Soerensen Provisional Application.
- At pages 45-46, the Petition states: "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,"—but provides no supporting citation for this subject matter in the Soerensen Provisional Application.
- At page 46, the Petition states: "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 90b’ and ‘clock domain 90c’ 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]"—but provides no supporting citation for this subject matter in the Soerensen Provisional Application.

- At page 47, the Petition relies on Figure 5 of Soerensen—but provides no citation to any corresponding figure in the Soerensen Provisional Application.

Rather than going limitation by limitation and citing support in the Soerensen Provisional Application for each passage of Soerensen (EX1006) relied on in the Petition, the Petition attempts a "hand wave" at note 2, which states:

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 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).

These citations fail to align with the specific passages of Soerensen relied on in the Petition. In addition, as explained below, citations cobbled together from three disparate documents in the Soerensen Provisional Application are inadequate to satisfy the written description requirement. *Purdue Pharma*, 230 F.3d 1320, 1323 ("one skilled in the art, reading the original disclosure, must immediately discern" the material at issue"); *Ariad Pharms.*, 598 F.3d 1336, 1352 ("[A] description that

merely renders the invention obvious does not satisfy the [written description] requirement.”)

The first and second citations relied on by Petitioner at page 110 of the Soerensen Provisional Application are part of a 4-page invention disclosure describing dynamically varying the clock frequency and voltage in a phone as a whole, rather doing so independently from domain to domain, as claimed. EX2001, ¶72. In fact, the invention disclosure that includes page 110 (see EX1022, at 107-112) makes no mention of there even being separate domains within the microprocessor.

The third citation relied on by Petitioner (at page 247 of the Soerensen Provisional Application) is part of a separate document entitled "Monaco Digital Baseband Processor Design Specification." EX1022, at 148-257. The 4-page invention disclosure that included page 110 does not mention or reference the "Monaco Digital Baseband Processor Design Specification." The Petition relies on the passage on page 247 which states that “the baseband may have six active clock domains.” Pet., note 2. However, the disclosure that surrounds this passage (on page 247) does not describe independently controlling clock frequency and voltage from processing domain to processing domain, or a separate clock for each domain as claimed.

The fourth citation relied on by Petitioner (at page 170 of the Soerensen Provisional Application) is part of a still further separate document that states that "[t]he DCLK, DSCLK, and BCLK frequencies can be dynamically changed." However, the passage does not describe this functionality in the context of a processor with separate clocks (as opposed to just clock signals) in each domain; nor does it state that the clock frequencies are varied dynamically independently of each other, as claimed. In fact, in the diagram below the passage, DCLK and DSCLK are treated as the same signal (which means that they do not vary independently of each other), and the diagram is so illegible that it does not show how BCLK operates in relation to DCLK and DSCLK. The discussion below the diagram describes 9 operational states, none of which describe clock frequencies are varied independently of each other from domain to domain, as claimed. For example, the description of State 2 (at the top of page 171) states:

This state ramps to the highest possible performance/energy level, with both the DSP (DCLK) and MCU (BCLK) cores running at maximum clock rate possible. Note that the cores wakeup and begin interrupt processing before the system PLL has acquired phase lock. Over time, less performance is required; therefore, the dynamic voltage control mechanism steps the clock rate down, allowing the supply voltage to drop. EX1022, 171.

This disclosure is describing both the DSP (DCLK) and MCU (BCLK) cores running at the same maximum clock rate possible, and when less performance is required, the dynamic voltage control mechanism steps "the" clock rate down. The fact that

the description is referring to "*the*" clock rate in the singular, indicates that both the DSP (DCLK) and MCU (BCLK) are adjusted down together at a common rate, rather than being adjusted independently of each other.

This passages relied on in the Soerensen Provisional Application fall far short of what is required to meet the exacting written description requirement, namely, that "one skilled in the art, reading the original disclosure, must immediately discern," the material at issue. *Purdue Pharma*, 230 F.3d 1320, 1323. "[A] description that merely renders the invention obvious," is insufficient. *Ariad Pharms.*, 598 F.3d 1336, 1352. Four unconnected quotes from a 257-page provisional application fail to satisfy these requirements.

Petitioner appears to be asking the Board to read the Soerensen Provisional Application through the lens of what was described in Soerensen (EX1006) in order to conclude that there is support for the relied on subject matter. Applying such a standard would be error. The Soerensen Provisional Application standing alone (not read through the lens of Soerensen (EX1006)) must satisfy the requirements of §112. *Ariad Pharms.*, 598 F.3d 1336, 1351 (inquiry limited to four corners of provisional application).

In short, since all of the portions of Soerensen (EX1006) relied on the Petition are not fully supported under §112, the relied on portions of Soerensen (EX1006)

are only entitled to a priority date of August 29, 2002, which is after the invention date of the '443 Patent which occurred by February, 2002. See, Section III, *supra*.

2. Soerensen's Reliance on Rational Frequency Relationships Fails to Teach or Render Obvious the Independent Clock Frequency Generation Required by the Challenged Claims

Unlike the globally-asynchronous system that is the subject of the '443 Patent, both the prior art systems disparaged in the '443 Patent and Soerensen utilize rational frequency relationships to achieve synchronization. EX1006, [0033] ("the clock signals distributed throughout the processor are synchronized.") (Emphasis added). Moreover, like the prior art systems disparaged in the specification, Soerenson imposes significant restrictions on frequency selection—which is the exact problem '443 Patent solves. See Section V(B), *supra*. Soerensen's disclosures fail to provide the independently dynamic frequency control across clock domains required by the Challenged Claims.

The Federal Circuit has held that “[t]he patentee may demonstrate intent to deviate from the ordinary and accustomed meaning of a claim term by including in the specification expressions of manifest exclusion or restriction, representing a clear disavowal of claim scope.” *Thorner v. Sony Comput. Entm't Am. LLC*, 669 F.3d 1362, 1366 (Fed. Cir. 2012) (Internal Citation Omitted). In the situation where the inventor's intention is to disavow the scope, “the inventor's intention, as expressed in the specification, is regarded as dispositive.” *Phillips v. AWH Corp.*, 415 F.3d

1303, 1316 (Fed. Cir. 2005). In order “[t]o disavow claim scope, the specification must contain ‘expressions of manifest exclusion or restriction, representing a clear disavowal of claim scope.’” *Cont’l Circuits LLC v. Intel Corp.*, 915 F.3d 788, 797 (Fed. Cir. 2019) (internal citation omitted). For example, “an inventor may disavow claims lacking a particular feature when the specification distinguishes or disparages prior art based on the absence of that feature.” *Poly-America, L.P. v. API Indus., Inc.*, 839 F.3d 1131, 1136-37 (Fed. Cir. 2016). In the present case, the specification deliberately and unambiguously disavows an embodiment that imposes a “rational factor” requirement on selection of the frequencies.

The Patent states, “[s]ome synchronization schemes restrict the phase relationship and relative frequencies of the clocks, thereby eliminating the need for hardware arbitration [27].” EX1001 6:33-35. In the '443 Patent, citation [27] corresponds to Sarmenta (EX2004). Sarmenta describes an “approach maintains [that] a known phase relationship between clocks whose frequencies are related by a rational factor, and exploits the predictability of their relative phases to algorithmically time communications without run-time arbitration contests.” EX2004, Abstract. As explained in Section VB above, the Challenged Claims do not cover synchronized systems like Sarmenta that maintain global synchronization by restricting the possible choice of frequencies to those related by a “rational factor.”

The “rational factor” in the Paper refers to the relationship between the clock frequencies of two independently clocked digital subsystems. These frequencies are related by the ratio of two integers, (M/N) , where (M) and (N) are small integers. This rational relationship ensures that the clocks align periodically, creating predictable points in time where the active edges of both clocks coincide. EX2001, ¶36.

For example, in Sarmenta, if one clock operates at a frequency (f_m) and the other at (f_n) , their frequencies satisfy the relationship $(f_m/f_n = M/N)$. This periodic alignment allows the system to define a “coincidence cycle,” during which the relative positions of the clock edges are predictable. As Dr. Kaeli explains, “This rational relationship ensures that the clocks align periodically, creating predictable points in time where the active edges of both clocks coincide.” EX2001, ¶36.

However, as the Patent observes when discussing Sarmenta, “these schemes impose significant restrictions on the possible choices of frequencies.” EX1001 6:35-37. To overcome the limitations of systems like Sarmenta, the Patent explains that invention provides “[t]he designers of each domain [to be] no longer constrained by the speeds of critical paths in other domains, affording them greater freedom in each domain to optimize the tradeoffs among clock speed, latency, and the exploitation of application parallelism via complex hardware structures.” EX1001

2:34-39. Accordingly, while Sarmenta introduces a synchronization approach leveraging rational frequency relationships, it imposes significant restrictions on frequency selection, as acknowledged in the Patent. As Dr. Kaeli explains, “The ‘443 invention overcomes these limitations, offering independence in clock frequency generation, allowing each frequency domain to adopt their own independent frequency, offering near infinite resolution in terms of the clock frequency selected for each domain.” EX2001, ¶36

The Petition states, "Soerensen's Figure 2, for instance, shows an example plurality of clock domains 90a, 90b, and 90c." Pet., 43-44. As shown in the Patent Owner's annotated Figure 2 below:

Soerensen's clock domain 90a (blue) operates at a full clock;
Soerensen's clock domain 90b (green) operates at either full clock or $\frac{1}{2}$ clock; and
Soerensen's clock domain 90c (red) operates at either full clock, $\frac{1}{2}$, $\frac{1}{4}$, $\frac{1}{6}$, $\frac{1}{8}$, $\frac{1}{10}$, $\frac{1}{12}$ or $\frac{1}{14}$.




Fig. 2

Specifically, Par. [0038] of Soerensen states:

[0038] Additionally, clock control 110 may divide the frequency of clock signal 84 in order to provide the components of the processor with clocks signals having reduced frequencies in order to meet the clock frequency requirement of the processor. In particular, dividers 62 and 64 allow the frequency of clock signal 84 to be reduced in frequency. For example, divider 62 may provide the clocked components operating in clock domain 90 b with a clock signal having half the frequency of clock signal 84 when the clock signal from the system PLL multiplier is selected. Similarly, divider 64 may provide the clocked components operating in clock domain 90 c with a clock signal having a frequency reduced by factors of 2, 4, 6, 8, 10, 12, and 14. Dividers 62 and 64 may be, for instance, divide-by-n counters or equivalent components for reducing the frequency of a clock signal. The various divider factors selected for each clock domain may be selected, for example, by divider select 66.

Accordingly, Soerensen teaches operating domain 70b at M integers (i.e., 1 and 2) and operating domain 70c at N integers (i.e., 1, 2, 4, 6, 8, 10, 12, and 14). Therefore, Soerensen teaches a system that requires the same “rational factor” as Sarmenta and, thus, like Sarmenta, “impose[s] significant restrictions on the possible choices of frequencies.” Ex.1001 6:35-37; EX2001, ¶36. Since such systems are outside the scope of the Challenged Claims, Petitioner has failed to meet its burden of establishing that Soerensen renders the challenged claims obvious under Ground IIA.

E. The Petition Fails to Demonstrate a Reasonable Likelihood of Success Under Ground IIIB

1. Soerensen Is Not Prior Art

Ground IIIB relies on the combination of Hemani and Soerensen. For the reasons set forth in Section VI(D)(1), *supra*, Soerensen is not prior art. For this reason alone, Ground IIIB fails.

2. Petitioner Fails To Demonstrate A Reasonable Likelihood Of Success In Combining Soerensen with Hemani.

Even if Soerensen was prior art, Ground IIIB would still fail because a POSITA would not have had a reasonable expectation of success. A motivation to combine references by itself is insufficient to prove obviousness. Rather,

obviousness requires proof of "a motivation to combine the references with a reasonable expectation of success." *Honeywell Int'l Inc. v. Mexichem Amanco Holding S.A. de C.V.*, 856 F.3d 1348, 1355 (Fed. Cir. 2017) (emphasis added).

Significantly, a finding that "one would no more have expected failure than success is not a valid ground for holding an invention to have been obvious." *Id.* With respect to the expectation of success requirement, "[t]he standard is not whether the patent owner can persuasively show that one of ordinary skill would have expected failure. Rather, the burden is on the [challenger] to show that one of ordinary skill would have had...a reasonable expectation of success." *Id.* (citations omitted). As explained by the Federal Circuit, "[u]npredictability of results equates more with nonobviousness rather than obviousness." *Id.*; *see also, Teva Phars. USA, Inc. v. Corcept Therapeutics, Inc.*, 18 F.4th 1377 (Fed. Cir. 2021); and *Eli Lilly & Co. v. Teva Pharms. Int'l Gmbh*, 8 F.4th 1331 (Fed. Cir. 2021).

As explained below, the Petition has failed to meet its burden of proving either a motivation to combine or a reasonable expectation of success based on Hemani and Soerensen.

a) Hemani Made Clear That There Was No Expectation Of Success

Hemani attempted to answer whether a globally-asynchronous, locally-synchronous (GALS) architecture would lower the power consumption of the clock function. EX1007, Title. As Hemani explained, "Power consumption in the clock,

as mentioned earlier, comes from two sources, one is due to the load seen at the clock terminals - the registers, and the second is due to the clock wire itself. " EX1007, 3.

While Hemani disclosed separate local clocks in each block (Hemani referred to each block as an "SB") to save power, Hemani did not explore an embodiment where the frequency of the clocks would vary from SB to SB. Instead, in Hemani's experiments, "[a]ll SBs are clocked at the same frequency." EX1007, 5 (emphasis added). Hemani's experiments demonstrated that when all SBs are clocked at the same frequency, the power consumed by the clock function decreases as a function of partitioning. EX1007, Fig. 4.

Because all of Hemani's SBs were clocked at the same frequency, Hemani did not need to evaluate whether a GALS architecture could successfully function when the frequency of the clocks varied from SB to SB. Nor did Hemani need to evaluate the performance costs associated with exchanging data between SB's operating at different clock frequencies, or whether those performance costs could be justified by the power savings achieved with local clocks that have separately variable frequencies. EX2001, ¶52. Instead, as explained below, Hemani identified a potentially fatal problem with the operation of the SBs at different clock frequencies, stating that "it may cause deadlock to occur." EX1007, 6.

Soerensen was directed to a standard processor with a single clock that produced a clock signal that was distributed to different domains. Unlike the

globally-asynchronous, locally-synchronous (GALS) architecture in Hemani, in Soerensen "the clock signals distributed throughout the processor are synchronized." EX1006, [0033] (emphasis added). In that synchronized context, Soerensen taught that power consumption could be reduced by scaling the clock frequency based on the computational demands on the processor.

What was unknown were two questions that neither Hemani nor Soerensen had answered, namely: (1) Will a globally-asynchronous, locally-synchronous (GALS) architecture become deadlocked if the clock frequencies vary from SB to SB, and, if not (2) given the performance costs associated with exchanging data between SB's operating at different clock frequencies, are the performance costs justified by the power savings achievable through the use of local clocks with frequencies that vary independently from domain to domain. Hemani could not predict the answers. Rather, Hemani cautioned that while running partitions at different speeds allowed "cheaper and cooler logic to be used," such a design will result in "a performance penalty and may cause deadlock." In this regard, the penultimate paragraph of Hemani stated:

This Paper was concerned mainly with power consumption and savings, but there are further consequences of the GALS architecture that deserve closer analysis. The GALS architecture allows SBs to run at different clock speeds. We are in the process of experimenting with other large designs where partitions have a natural need to run at grossly different speeds, and thus have the potential for exploiting the GALS characteristic. Clock skew constraints are limited to the SB

boundaries and thus smaller than in the GS case which results in more effective computation time allowing cheaper and cooler logic to be used. On the other hand, the introduction of global protocol signals will increase the total area required causing the average wire length to grow and thus increasing the performance and power penalty. The protocol exchange degrades the communication speed between blocks adding more performance penalty and may cause deadlock to occur. EX1007, 6 (emphasis added).

In the absence of evidence that it would work, the inventors of the '443 Patent designed an MCD processor with separate clocks whose frequencies varied from domain to domain. While the design afforded "a number of *potential* advantages over a singly clocked design," including improvements in "*both performance and energy efficiency*," whether or not the design would work (e.g., not result in deadlock) and whether the advantages were outweighed by the performance penalties, were *unknown*. Ex.1001, 4:23-50 (emphasis added).

To test their design, the inventors identified a performance metric known as the "energy-delay product" (EPD)—which measures how well a system balances energy usage with performance.¹ The inventors then used a simulator to implement their MCD processor where the clock frequency varied (independently) from domain to domain. The implementation worked (i.e., it did not result in deadlock). At the same time, the inventors simultaneously and collectively quantified the impact of (i)

¹ EPD is calculated by multiplying normalized energy consumption by normalized execution time.

the performance penalty associated with moving data between domains operating at different clock frequencies, together with (ii) the power savings associated with the operation of separate local clocks and the ability to scale clock frequency for a given block based on the computational demand in that block.

Based on their implementation, the inventors found that an MCD processor with separate clocks whose frequencies vary from domain to domain not only worked, but resulted in "a 20% average improvement in energy-delay product." EX1001, 4:58-59. As explained by Dr. Kaeli, "[t]he scale of the improvement was surprising." EX2001, ¶59; *Honeywell*, 856 F.3d 1348, 1355 ("[u]npredictability of results equates more with nonobviousness rather than obviousness."). The findings of the inventors were nothing short of groundbreaking. As noted above, the Paper evidencing the inventors' work was awarded the "Test of Time Award" by the IEEE, and has been cited 575 times since it was published. International Symposium on High-Performance Computer Architecture. EX2003, ¶9. The Test of Time Award "recognizes the most influential papers published in prior sessions of the International Symposium of High Performance Computer Architecture (HPCA) each of whom have had significant impact in the field." EX2003, ¶10. The Paper has been cited 575 times since it was published.

b) The Petition Failed To Prove A Reasonable Expectation Of Success

Despite relying on Hemani in its mappings, the Petition failed to acknowledge,

let alone address, Hemani's penultimate statement that, in a GALS architecture where SBs are allowed to run at different clock speeds, "[t]he protocol exchange degrades the communication speed between blocks adding more performance penalty and may cause deadlock to occur." EX1007, 6 (emphasis added). It was incumbent upon the Petitioner to explain why, faced with this statement from Hemani, a POSITA would have had a reasonable expectation that a GALS architecture where SBs are allowed to run at different clock speeds would work (not deadlock).

The Petition offers only the following conclusory statement on the question of whether a reasonable expectation of success existed. From the Petition:

... 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. Pet., 55-56 (emphasis added).

As an initial matter, the premise of the underlined statement is wrong: Soerensen does not relate to lowering power consumption in globally-asynchronous microprocessors. See, EX1006, [0033] ("the clock signals distributed throughout the processor are synchronized.") (emphasis added); EX2001, ¶71. Moreover, even if

the two references "relate" to the same subject, that fact alone is not probative of whether a POSITA would have had a reasonable expectation of success, i.e., that the combination would work and not deadlock. EX2001, ¶72. The Petition is devoid of any evidence that a POSITA would have had a reasonable expectation that the combination of Hemani and Soerensen would have yielded a successful result as opposed to deadlock. *Honeywell*, 856 F.3d 1348, 1355 ("a finding that 'one would no more have expected failure than success is not a valid ground for holding an invention to have been obvious.'")

3. Ground IIIB Also Failed To Prove A Motivation To Combine

Absent reasonable guidance that an MCD processor with separate clocks whose frequencies vary from domain to domain would actually work (not result in deadlock) but also achieve an improvement in energy-delay product, a POSITA would not have been motivated to combine Hemani and Soerensen as argued in the Petition. EX2001, ¶72. While Soerensen taught that, when synchronized signals from a single clock are distributed to different domains, the frequencies of the distributed signals could be scaled to save power, Soerensen was not directed to a globally-asynchronous, locally-synchronous (GALS) architecture. See, EX1006, [0033] ("the clock signals distributed throughout the processor are synchronized." (emphasis added)). Hemani, which was directed to a GALS architecture, explained that varying clock frequencies from domain to domain "may cause deadlock to

occur." EX1007, at 6. These teachings would not have motivated a POSITA to import the frequency scaling techniques from Soerensen's synchronized environment into Hemani's globally-asynchronous architecture, as doing so "may cause deadlock to occur." EX1007, at 6.

Finally, the proposed combination of Hemani and Soerensen is overly simplistic. That combination ignores the complexities and tradeoffs associated with the MCD architecture set forth in the Challenged Claims by looking only to the power savings disclosed in Soerensen associated with scaling the clock frequency in a given block as a reason for combining the references, while ignoring the performance costs of simultaneously operating different blocks at different clock frequencies and the fact that the design could deadlock. However, a POSITA would have needed to understand the complete picture, e.g., that the power savings benefit outweighed the performance costs, before being motivated to combine the references as argued in the Petition. Absent an understanding of the complete picture (which was neither known nor predictable), a POSITA would not have been motivated to combine Hemani and Soerensen as argued in the Petition.

4. Neither Soerensen Alone Nor In Combination With Hemani Teaches Or Inherently Discloses The Independent Clock Frequency Generation Required By The Challenged Claims

Assuming a POSITA were motivated to make the Petitioner's proposed combination, neither Soerensen nor the proposed combination of Soerensen with

Hemani teaches or inherently discloses the claimed feature of “separately generating a clock signal...the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains.” Instead, both Soerensen alone and the combination with Hemani impose significant frequency restrictions, namely Sarmenta’s “rational factor,” which contradicts the claimed invention. See Section VI(B), *supra*.

The Petition states that Soerensen teaches that “various frequencies required by the processor are often generated from the system clock by a phase locked loop (PLL) multiplier.” Pet., 58. (emphasis by PO). As discussed above, Soerensen’s PLL multiplier divides the system clock by integers (i.e. 1, 2, 4, 6, 8, 10, 12 and 14) and as a result maintains Sarmenta’s “rational factor” scheme which is outside the scope of the Challenged Claims. See Section V(B), *supra*. In contrast to the invention of the '443 Patent, Soerensen “impose[s] significant restrictions on the possible choices of frequencies.” EX1001, 6:35-37.

The additional teachings of Hemani do not cure the shortcomings of Soerensen. The Petition notes that Hemani explains that the possibilities are “either to have an independent local clock generator or to use a global reference clock signal with frequency multiplier in SBs” and “[i]n the former instance, Hemani teaches local clock generators based on ring oscillators.” Pet., 59. Accordingly, the Petition alleges that “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.” Pet., 59. As a result, the Petition concludes that a POSITA would be motivated to swap Soerensen’s phase locked loop (PLL) multiplier for Hemani’s ring oscillators to generate the clock for Soerensen’s clock domains 70b (green) and 70c (red) as shown in Soerensen’s Figure 2 (annotated below).




Fig. 2

Assuming *arguendo* that a POSITA would be motivated² to swap Soerensen's phase locked loop (PLL) multiplier for Hemani's ring oscillators to generate the clocks domains 70b and 70c, the resulting system would still not “separately generat[e] a clock signal...the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains,” as recited in the Challenged Claims.

Instead, the proposed combination of Hemani and Soerensen would utilize Hemani's ring oscillator to operate domain 70b at M integers (i.e., 1 and 2) and Hemani's ring oscillator to operate domain 70c at N integers (i.e., 1, 2, 4, 6, 8, 10, 12, and 14). Accordingly, the proposed combination would again utilize Sarmenta's “rational factor.” Therefore, the Petitioner's proposed combination, like Soerensen alone, “impose[s] significant restrictions on the possible choices of frequencies,” — a scheme that the '443 Patent disparaged. EX1001, 6:35-37; EX2001, ¶72.

To the extent Petitioner alleges that the proposed combination of Hemani and Soerensen inherently results in a system that does not utilize Sarmenta's “rational factor” because it uses a ring oscillator, the Federal Circuit directly prohibits this vague usage of inherency. Specifically, the Court explains “the use of inherency, a

² Dr. Kaeli contends that “that the Petitioner is working backwards from the ‘443 patent's claims, looking for disclosures in disparate references, without meaningfully considering whether the reference they site would, in fact, be combined in the first place.” Ex. 2001, ¶71

doctrine originally rooted in anticipation, must be carefully circumscribed in the context of obviousness,” and “[t]he mere fact that a certain thing may result from a given set of circumstances is not sufficient to establish inherency.” *Par Pharm., Inc. v. TWi Pharm., Inc.*, 773 F.3d 1186, 1194-95 (Fed. Cir. 2014) (internal markings and citations omitted). The Court has “emphasized that the limitation at issue necessarily must be present in order to be inherently disclosed by the reference.” *Southwire Co. v. Cerro Wire LLC*, 870 F.3d 1306, 1310-11 (Fed. Cir. 2017) (internal citation omitted). Accordingly, “[a] party must...meet a high standard in order to rely on inherency to establish the existence of a claim limitation in the prior art in an obviousness analysis.” *Millennium Pharms., Inc. v. Sandoz, Inc.*, 862 F.3d 1356, 1367 (Fed. Cir. 2017).

In the present case, Petitioner has not met this standard of showing that the mere inclusion of a ring oscillator must necessarily result in operating domain 70b at something other than M integers (i.e., 1 and 2) and operating domain 70c at something other than N integers (i.e., 1, 2, 4, 6, 8, 10, 12, and 14). Instead, the record shows that domains 70b and 70c are operated at M and N integers, respectively. EX2001, ¶72.

As a result, Hemani, alone or in combination with Soerensen, fails to teach, inherently or expressly, “separately generating a clock signal...the frequency being dynamically changeable independently of the frequencies of the clock signals

generated for others of the plurality of domains,” as recited in the Challenged Claims. For this further reason, Petitioner has failed to meet its burden of establishing that the Challenged Claims are invalid under Ground IIIB.

F. Ground IIIC Fails Because Moore Does Not Cure the Shortcomings of Soerensen.

The Petition does not allege, nor does Moore teach, “separately generating a clock signal...the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains,” as recited in the independent Challenged Claims. Accordingly, the Petitioner has failed to meet its burden with regard to Ground IIIC for at least the reasons discussed with regard to Gound IIIA above.

VIII. SUMMARY OF ARGUMENTS ON THE MERITS

The Petition falls well short of demonstrating a reasonable likelihood of success in establishing the invalidity of the Challenged Claims. Each Ground advanced by the Petitioner suffers from critical deficiencies, including reliance on inapplicable or insufficiently supported prior art, mischaracterizations of the references, and a failure to address key technical distinctions and legal requirements. The Petitioner has not provided persuasive evidence to overcome the invention's demonstrated novelty and non-obviousness. Moreover, the Petition fails to establish either a motivation to combine references or a reasonable expectation of success, as

required under 35 U.S.C. § 103. For these reasons, the Board should find that the Petitioner has not met its burden on the merits, and the institution of the IPR should be denied.

IX. INSTITUTION SHOULD BE DENIED UNDER THE *FINTIV* FACTORS

On August 3, 2023, Petitioner was served with a Complaint filed by Patent Owner alleging infringement of the '439 Patent in *Intellectual Ventures I LLC v. TCL Electronics Holdings Ltd. et al.*, WDTX-6-23-cv-00309 (the “Parallel Litigation.”) EX2006, at 7, #16. Petitioner waited until August 2, 2024—1 day shy of a full year after being served with the Complaint—to file the present IPR. As a result, the Parallel Litigation has already reached an advanced stage, with the court issuing a claim construction order on September 10, 2024, and the parties having served their final infringement and invalidity contentions on September 23, 2024. EX2007 (Claim Construction Order) and EX2008 (Scheduling Order). The case is so advanced that fact discovery is scheduled to be closed on February 24, 2025—a few days after the anticipated date of the institution decision in this IPR on February 18, 2025. The trial is scheduled for July 28, 2025 (EX2006), nearly 7 months before the statutory deadline for issuing a Final Written Decision on February 18, 2026.

The precedential order in *Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (PTAB Mar. 20, 2020) (“*Fintiv*”) identifies six factors for the Board to consider in deciding whether to exercise discretion to deny institution based on parallel

proceedings involving the same patent. As discussed below, the application of the present facts to these factors demonstrates that there are strong reasons for exercising discretion to deny institution here.

A. Factor 1 Weighs Against Institution Because A Stay Has Not Been Granted In The Parallel Litigation And There Is No Evidence That A Stay Would Be Granted If This IPR Was Instituted

No stay has even been requested in the Parallel Litigation. In addition, even if this IPR was instituted (which would not occur until on or around February 18, 2025) it is unlikely that the Court (Judge Albright) would grant a stay, particularly given the very advanced stage of the Parallel Litigation at that point and the proximity to the trial date. *See, generally, Continental Intermodal Grp. v. Sand Revolution LLC*, No. 7:18-cv-00147 (W.D. Tex. July 22, 2020) (Judge Albright); *Multimedia Content Mgmt. LLC v. Dish Network*, No. 6:18-CV-00207-ADA, 2019 U.S. Dist. LEXIS (W.D. Tex. May 30, 2018) (Judge Albright) (denying stay).

Thus, the first factor weighs against institution.

B. Factor 2 Weighs Against Institution Because The Parallel Litigation Has A Trial Date Scheduled Well Before The Projected Statutory Deadline For The Board's Final Written Decision

As noted above, the trial in the Parallel Litigation is scheduled for July 28, 2025 (EX2006), nearly 7 months before the statutory deadline for issuing a Final Written Decision in this matter on February 18, 2026.

With respect to factor 2, the Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation (June 21, 2022) directs the Board to consider the median time from filing to case disposition statistics. The median time from filing to case disposition in the Western District of Texas (where the Parallel Litigation is pending) is 28.3 months (EX2009, 27). Given that the Parallel Litigation was filed on April 26, 2023 (EX2006 at 6, #1), the 28.3 month median time to disposition indicates disposition of the Parallel Litigation by the beginning of September, 2025—about six months before the statutory deadline for issuing a Final Written Decision in this matter on February 18, 2026.³

Thus, the second factor weighs against institution.

C. Factor 3 weighs against institution because the parties will have invested significant effort in the Parallel Litigation by the time the Board’s institution decision is due

Factor 3 relates to investment in the Parallel Litigation by the court and the parties. *Fintiv*, Paper 11 at 9. This factor is judged from the date of the institution decision, which is expected to be no later than February 18, 2025. *Id.*

³ The Petition (at 74) cites the USPTO Interim Procedure. However, the Petition neglected to perform the timing analysis based on trial statistics mandated by the Interim Procedure, at 3 (“when considering the proximity of the district court’s trial date to the date when the PTAB final written decision will be due, the PTAB will consider the median time from filing to disposition of the civil trial for the district in which the parallel litigation resides.”)

In the Parallel Litigation, Patent Owner served Preliminary Infringement contentions (February 12, 2024), Petitioner (Defendant) served Preliminary Invalidity Contentions (April 8, 2024), the parties exchanged claim terms for construction (April 22, 2024), the parties exchanged proposed claim constructions (May 6, 2024), the parties completed claim construction briefing with Petitioner filing an Opening Claim Construction Brief (May 27, 2024), Plaintiff filing a responsive Claim Construction Brief (May 27, 2024), Petitioner (Defendant) filing a Reply Claim Construction Brief (July 1, 2024), and Plaintiff filing a Sur-Reply Claim Construction Brief (July 15, 2024). EX2008, at 2. In addition, the Court held a Claim Construction Hearing on August 8, 2024, and issued its Claim Construction Order (EX2007) on September 10, 2024. EX2006, at 10, #44.

The parties and the Court have already made a substantial investment in the Parallel Litigation, which will grow even larger by the expected date of the institution decision on February 18, 2025. By that time, the parties will be only days away from the scheduled close of fact discovery on February 24, 2025. EX2006, at 8.

Thus, the third factor weighs against institution.

D. Factor 4 weighs against institution because there is substantial overlap between the Parallel Litigation and this IPR

Rather than offer a full *Sotera*⁴ stipulation that would have barred Petitioner from pursuing any ground raised *or that could have been reasonably raised in this IPR* if this IPR is instituted, Petitioners offered a much more limited *Sand*⁵ stipulation stating that, if this IPR is instituted “Petitioners will not pursue in the Texas Litigation *the same grounds* asserted in the Petition.” Pet., at 75. Under the present facts, this stipulation is so narrow as to be meaningless. To gain two bites at the apple, Petitioners have laid the groundwork to advance the identical (or cumulative) prior art and virtually identical arguments in the Parallel Proceedings as advanced in these proceedings, by rearranging the combinations to technically get around the narrow *Sand* stipulation. Below are portions of Petitioner’s invalidity contentions in the Parallel Litigation, which for example rely upon proposed combinations of Iyer and Njolstad; Iyer, Hemani and Njostad; Iyer and Liljeberg; Iyer, Hemani and Liljeberg; Soerensen and Iyer; Soerensen and Liljeberg; Soerensen

⁴ *Sotera Wireless, Inc. v. Masimo Corporation*, IPR2020-01019, Paper 12, at 18-19 (P.T.A.B. December 1, 2020)(stipulating “that, if IPR is instituted, [Petitioner] will not pursue in the District Court Litigation any ground raised or that could have been reasonably raised in an IPR.”)

⁵ *Sand Revolution II, LLC v. Continental Intermodal Group—Trucking LLC*, IPR2019-01393, Paper 24 at 11-12 (P.T.A.B. June 16, 2020).

and Njolstad; and Soerensen, Liljeberg and Njolstad -- none of which would be precluded by the *Sand* stipulation offered by Petitioner.

'443 Patent Claims	Ground
1, 4, 5, 8, 10	Obvious over Iyer* ²
1, 4, 5, 8, 10	Obvious over Njolstad*
1, 4, 5, 8, 10	Obvious over Njolstad in view of Liljeberg*
1, 4, 5, 8, 10	Obvious over Soerensen*
1, 4, 5, 8, 10	Obvious over Soerensen in view of Hemani*
1, 4, 5, 8, 10	Obvious over Soerensen in view of Moore*
1, 4, 5, 8, 10	Obvious over Iyer in view of Njolstad
1, 4, 5, 8, 10	Obvious over Iyer in view of XScale and Hemani
1, 4, 5, 8, 10	Obvious over Iyer in view of Hemani and Njolstad
1, 4, 5, 8, 10	Obvious over Iyer in view of Pillai
1, 4, 5, 8, 10	Obvious over Iyer in view of Liljeberg

1, 4, 5, 8, 10	Obvious over Iyer in view of Hemani and Liljeberg
1, 4, 5, 8, 10	Obvious over Iyer in view of Pillai and Njolstad
1, 4, 5, 8, 10	Obvious over Soerensen in view of Iyer
1, 4, 5, 8, 10	Obvious over Soerensen in view of Liljeberg
1, 4, 5, 8, 10	Obvious over Soerensen in view of Njolstad
1, 4, 5, 8, 10	Obvious over Soerensen in view of Liljeberg and Njolstad

EX2010, at 16-17.

Because the *Sand* stipulation will essentially allow Petitioner to advance substantially the arguments made here in the district court, the stipulation does little to mitigate the duplication of efforts that will occur in the PTAB and the district court should this IPR be instituted. While the Board in *Sand* found such a limited stipulation carried only “marginal” weight, the stipulation should carry even less weight here given Petitioner’s plan to advance rearranged versions of the same prior art combinations in the district court. *Sand Revolution*, IPR2019-01393, Paper 24 at 11.

Because the limited *Sand* stipulation offered by Petitioner does not mitigate the potential for significant overlap between this IPR and the Parallel Litigation, there are particularly strong “concerns of inefficiency and the possibility of conflicting decision,” if this IPR is instituted. *Fintiv*, Paper 11 at 12; *see Eclinicalworks, LLC et al. v. Decapolic LLC*, IPR2022-00229, 2022 Pat. App. LEXIS 2274, Paper 10 at *13 (PTAB Apr. 13, 2022) (finding factor 4 weighs strongly in favor of denial where the ground in the Petition was included in the preliminary invalidity contentions in parallel case).

Thus, the fourth factor weighs against institution.

E. Factor 5 weighs against institution because the patent owner is the same and the Petitioner are the same parties as in the Parallel Litigation

Patent Owner and Petitioners are the same parties as in the Parallel Litigation.

Thus, the fifth factor weighs against institution.

F. Factor 6 weighs against institution because there are no other circumstances that favor institution, and the merits of Petitioner’s challenge is weak

As discussed above, Petitioners’ challenge is weak. Accordingly, the Petition not only fails to meet the “reasonable likelihood” standard under 35 U.S.C. § 314(a) but fails to meet the higher “compelling merits” standard under Factor 6. See *CommScope Technologies LLC v. Dali Wireless, Inc.*, IPR2022-01242, Paper 23 at 4-5 (P.T.A.B. Feb. 27, 2023) (precedential).

Thus, the sixth factor weighs against institution.

All *Fintiv* factors weigh against institution of the IPR. Under §314(a) and, based on the totality of the *Fintiv* factors, institution should be denied.

X. CONCLUSION

For the above reasons, Patent Owner respectfully requests that the Board deny institution of the Petition.

Respectfully submitted,

Dated: November 19, 2024

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
University of Rochester
(Intellectual Ventures I LLC)

CERTIFICATE OF COMPLIANCE WITH WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), I certify that this Patent Owner Preliminary Response complies with the type-volume limits of 37 C.F.R. § 42.24(b)(1) because it contains 13,967 words, excluding the parts of this Patent Owner's Preliminary Response that are exempted by 37 C.F.R. § 42.24(a), according to the word processing system used to prepare this Patent Owner's Preliminary Response.

Respectfully submitted,

Dated: November 19, 2024

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
University of Rochester
(Intellectual Ventures I LLC)

CERTIFICATE OF SERVICE

The undersigned certifies that pursuant to 37 C.F.R. § 42.6(e), a copy of the foregoing **PATENT OWNER'S PRELIMINARY RESPONSE** was served via email (as consented to by counsel) on November 19, 2024 to lead and backup counsel of record for Petitioners as follows:

Russell E. Levine, P.C. (Reg. No. 32,153)
KIRKLAND & ELLIS LLP
333 West Wolf Point Plaza
Chicago, IL 60654

Joshua Glucoft (Reg. No. 67,696)
KIRKLAND & ELLIS LLP
2049 Century Park East
Suite 3700
Los Angeles, CA 90067

Kat Li (Reg. No. 64,857)
Michael T. Dinnella (Reg. No. 82,580)
KIRKLAND & ELLIS LLP
401 Congress Avenue

Email: TCL-IV-IPR@kirkland.com

Dated: November 19, 2024

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
University of Rochester
(Intellectual Ventures I LLC)