

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

LENOVO (UNITED STATES) INC.,

Petitioner,

v.

UNIVERSITY OF ROCHESTER
(INTELLECTUAL VENTURES I, LLC),

Patent Owner.

Case No. IPR2024-01225

U.S. Patent No. 7,089,443

PATENT OWNER'S SUR-REPLY

TABLE OF CONTENTS

PATENT OWNER’S EXHIBIT LIST	iii
I. Introduction.....	1
II. “Preferred embodiment” does not require rational ratio tethering (sharing a reference $\neq$ sharing a frequency relationship).....	3
III. “Independence” means freedom from other domain constraints—not merely “not literally derived from another domain’s clock”	7
IV. Georgiou’s multiplexer is a selector, not a generator—and its master-fraction scheme defeats “independently changeable”	8
V. Shenai teaches rational multiples for synchronous comms—opposite of GALS	11
VI. Antedation — Actual reduction to practice	12
A. Apparatus claims: field appropriate lab proof (including simulation) satisfies “works for intended purpose”	13
B. Method claims: steps were actually performed.....	16
VII. Shenai’s provisional entitlement fails (written description).....	17
VIII. Conclusion	18

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	Provisional Application No. 60/311,575, filed August 10, 2001
EX2007	Provisional Application No. 60/337,837, filed November 5, 2001
EX2008	Provisional Application No. 60/337,479, filed November 5, 2001
EX2009	Provisional Application No. 60/338,510, filed November 5, 2001
EX2010	SimpleScalar Tutorial
EX2011	U.S. Patent No. 5,634,041 "Rationally clocked communication interface" by Pratt et al. ("Pratt").
EX2012	A. Chakraborty and M. R. Greenstreet, "Efficient self-timed interfaces for crossing clock domains," Ninth International Symposium on Asynchronous Circuits and Systems, 2003. Proceedings., Vancouver, BC, Canada, 2003, pp. 78-88.
EX2013	Declaration of Dr. David Kaeli

I. Introduction

Petitioner's Reply re-packages the Petition around four propositions: (1) that the '443 patent's preferred embodiment requires rational-ratio tethering because the specification distributes a shared reference (or common external clock) to domain-local PLLs (Reply, 11); (2) that "independent" in the claims means only "not directly derived from another domain's clock," such that Georgiou's multiplexer-based selection among master-derived fractions qualifies as a clock that "separately generates" (Reply, 23–27; Petition, 68–70); (3) that Shenai's small-integer multiples supply the missing clocking and support the merits (Reply, 21–22); and (4) that Patent Owner's antedation fails for both apparatus and method claims, while §119(e) provisional entitlement rescues Shenai (Reply, 14–19).

First, on claim construction, the Sur-Reply explains why the intrinsic record controls. For example, the Summary of the '443 Patent identifies a globally-asynchronous, locally-synchronous (GALS) microarchitecture where "each functional block operates with a separately generated clock," the global network distributes only an external reference to each local PLL, and the independence of domain clocks removes global skew and cross-domain critical-path constraints while enabling per-domain Dynamic Voltage and Frequency Scaling (DVFS) via separate voltage inputs and controllable in-domain frequency circuits. EX1001, 4:29–33; 4:38–40; 4:41–56; 4:55–66. The Detailed Description distinguishes synchronization

schemes that restrict phase relationships or relative frequencies and affirms “no known phase relationship.” EX1001, 6:32–36; 6:71–7:3.

Second, on the merits, the Sur-Reply addresses the Petition’s two primary arguments. Georgiou describes a master clock feeding dividers that produce predetermined fractions (e.g., divide-by-two, divide-by-n), with “each multiplexer 470 output[ting] the selected clock frequency” “while maintaining synchronization and reliability.” EX1008, 8:37–47; 8:61–67; 10:48–51; Petition, 68–70. Shenai teaches clock/shifters operating at integer/half/quarter multiples “to facilitate synchronous communication between loads.” EX1004, ¶[0036]. These references frame the issues of selector vs. generator, rational fractions vs. independent changeability, and synchrony vs. GALS asynchrony that are dispositive for the challenged limitations (including “separately generating,” “independently changeable,” and the method requirement that domains operate asynchronously across while synchronous within). EX1001, 17:19–27; 17:63–18:1.

Third, on antedation, the Sur-Reply separates the standards for apparatus and method claims and identifies the record showing: (i) field-appropriate cycle-accurate/timing simulation that computes and executes per-domain V/f schedules; (ii) modeling of independent clocks with independent jitter to capture asynchronous crossings and clock-domain crossing timing; and (iii) the specification’s own “existence proof” characterization. EX1001, 12:27–35; 12:41–


44; 12:51–63; 12:64–71; Reply, 14–15; POR, 22–23. Finally, the Sur-Reply addresses §119(e): the provisional passages the Reply invokes do not describe a controller determining current targets or a microprocessor, and thus cannot confer priority under governing written-description law. Reply, 16–19; POR, 39–43; EX1004, ¶¶[0026], ¶¶[0031].

II. “Preferred embodiment” does not require rational ratio tethering (sharing a reference ≠ sharing a frequency relationship)

The Reply contends Patent Owner’s construction “excludes the preferred embodiment” because the specification distributes a common external clock to per domain PLLs. Reply, 11. The intrinsic record says the opposite. The ‘443 Patent explained that schemes that “restrict the ... relative frequencies of the clocks” were problematic because they “impose significant restrictions on the possible choices of frequencies.” EX1001, 6:32-36 (emphasis added). In contrast to prior systems that restricted the relative frequencies of the clocks, the Summary defines the “present invention” as an MCD/GALS microarchitecture in which “each functional block operates with a separately generated clock” and stressed “[t]he independence of each local clock domain.” EX1001, 4:29–33; 4:41–56 (emphasis added). The Summary then spells out what independence yields: “no global clock skew requirement” and design freedom because “[t]he designers of each domain are no longer constrained by the speeds of critical paths in other domains.” EX1001, 4:41–56. The Detailed Description further explains:

... These values assume a 1 GHz, on-chip clock generated from a common external 100 MHz clock source. Despite the common use of the external clock, because the local clock sources are independent, ... EX1001., 10:37-42 (emphasis added).

The diagram below (created by Patent Owner) illustrates an example of the ‘443 Patent’s disclosure of the common use of an external clock for independent clock sources:




Each local clock is independently changeable because it can scale up or down the common external clock with no restrictions, and there are no relative frequency relationships between the local clocks in the different domains. EX1001, 6:32-36 (explaining why schemes that “restrict the ... relative frequencies of the clocks” were problematic); EX2013, ¶34 (explaining that the ‘443 Patent allows “each frequency domain to adopt their own independent frequency, offering near infinite resolution in terms of the clock frequency selected for each domain”) and ¶47 (“the frequency of each clock in a domain can be changed independently of other clocks”).

The paragraph the Reply invokes from the Preferred Embodiment is entirely consistent: it identifies four on chip domains (plus a memory interface as a fifth, external domain), and notes an on-chip domain “may vary somewhat in size and still be covered by a single clock”—i.e., within that domain—while inter domain interfaces are handled asynchronously. EX1001, 5:7–19. Nothing there imposes a restriction on the changeability of a frequency within a given domain, or imposes any relative frequency relationships between the clocks in the different domains. Reply, 11; POR, 60.

Just as importantly, the specification distinguishes and criticizes the very class of rational ratio schemes Petitioner asks the Board to adopt. It states: “Some synchronization schemes restrict the phase relationship and relative frequencies of the clocks” and so “impose significant restrictions on the possible choices of frequencies,” expressly invoking the rational clocking literature. EX1001, 6:32–36; see also 3:54–56. And it draws the governing line: “Our design specifically recognizes the overhead associated with independent clocks with no known phase relationship.” EX1001, 6:71–7:3; POR, 26, 44–45. Sharing a reference to a local PLL does not equate to sharing a frequency relationship; the patent embraces no known phase relationship across domains.

By contrast, the references relied on in the Petition are aimed at synchrony. Georgiou describes a master clock feeding dividers whose outputs are “some

predetermined fraction” (e.g., “divide by two,” “divide by n”), and “each multiplexer 470 outputs the selected clock frequency.” EX1008, 8:37–47; 8:61–67. Georgiou’s objective is to adjust operating points “while maintaining synchronization and reliability.” EX1008, 10:48–51. Shenai likewise runs loads at integer/half/quarter multiples of a reference to facilitate synchronous communication. EX1004, ¶[0036]. Those are precisely the rational ratio regimes the ‘443 patent contrasts with its GALS approach. POR, 51–53, 59; Reply, 21–22. The diagram below (created by Patent Owner) illustrates the differences between the independently changeable (*unrestricted*) clocks disclosed and claimed in the ‘443 Patent, and the *restricted* clocks with *restrictive relative frequency relationships* disclosed in Shenai and Georgiou:



Two intrinsic anchors reinforce Patent Owner’s reading. First, the method claims require domains to operate synchronously within a domain and

asynchronously across domains—flatly inconsistent with Georgiou’s “maintaining synchronization” and with Shenai’s preference for synchronous inter-load communication. EX1008, 10:48–51; EX1004, ¶¶0036]. Second, the Detailed Description models independent, unphased clocks (including independent jitter) and treats the clock-domain crossing (CDC) overhead as the necessary cost of having no known phase relationship—the opposite of rationally phased schemes. EX1001, 6:71–7:3; 12:51–63; POR, 35.

As a result, the preferred embodiment uses local PLLs fed by a shared external clock, yet rejects tethered frequency relationships across domains; that is fully consistent with “separately generated” and “independently changeable” domain clocks under GALS.

III. “Independence” means freedom from other domain constraints—not merely “not literally derived from another domain’s clock”

Petitioner’s gloss reduces “independent” to “not directly *derived* from another domain’s clock.” Reply, 24–26. The ‘443 patent doesn’t use the terms “derive” or “derived” anywhere, let alone in connection with its discussion of independent domain clocks. Instead, the ‘443 Patent explains that independently changeable clocks are ones that do not “*restrict the ... relative frequencies of the clocks*” or “impose significant *restrictions on the possible choices of frequencies*.” EX1001, 6:32-36. The intrinsic evidence makes clear that such schemes are outside the scope

of domains with clocks that are “changeable independently” from the frequencies of the clock signals in the other domains, as claimed.

Measured against that standard, Georgiou and Shenai pull the opposite direction: rigid, set fractions rather than independence. Georgiou’s clocks are master derived fractions; “each multiplexer 470 outputs the selected clock frequency” from divide by N inputs, and the system aims to “maintain[] synchronization.” EX1008, 8:37–47; 8:61–67; 10:48–51. Shenai uses integer/half/quarter multiples for synchronous inter load communications. EX1004, ¶[0036]. Those choices couple domains’ allowable frequencies to a shared master and target synchrony, not the ‘443’s GALS asynchrony with independently changeable domain clocks. POR, 21–22, 51–57; Reply, 21–27.

Accordingly, independence in the ‘443 is freedom from cross domain timing constraints and domain local Dynamic Voltage and Frequency Scaling (DVFS), not merely the absence of a direct feed through from another domain’s clock.

IV. Georgiou’s multiplexer is a selector, not a generator—and its master-fraction scheme defeats “independently changeable”

Petitioner’s mapping elevates multiplexer 470 to “the claimed ‘clock’ for separately generating a clock signal.” Petition, 68–69. That framing cannot be squared with Georgiou’s own description. Georgiou teaches a master clock 435 feeding clock dividers whose outputs are “some predetermined fraction of the master clock frequency”—e.g., “divide-by-two” and “divide-by-n”—and then states that

“each multiplexer 470 outputs the selected clock frequency.” EX1008, 8:37–47; 8:61–67. A device that selects among already-generated master-derived fraction signals is not a “clock ... separately generating” the domain’s clock. The claim language of the ‘443 Patent is the opposite: “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.” EX1001, 17:19–27. Petition, 68–70; POR, 54–57.

Two structural points make Petitioner’s theory untenable:

First, under Georgiou, generation occurs in the divider blocks driven by the master; the mux does not synthesize a new frequency, it chooses one of the divider’s precomputed outputs and forwards it. EX1008, 8:37–47; 8:61–67. Even if one loosely called the divider tree a “generator,” it is not a per-domain generator: the same bank of divider outputs is shared across units, and multiplexer 470 for each unit merely taps the shared set. That is not “for each ... domain, a clock for separately generating” that domain’s clock; rather, it is a global generator with local selection. EX1001, 17:19–27; Petition, 68–70.

Second, Georgiou’s allowed frequencies are tethered to each other—defeating the requirement that they be “independently changeable.” Because every unit’s clock must be one of the small-integer fractions of a master clock, the set of permissible

domain frequencies in each domain is constrained, and the relative frequencies across domains are similarly limited to a fixed set of relationships. Changing the master frequency rescales every available fraction; and the relative relationships among domains are fixed to the same rational set. That is the antithesis of “independently changeable” domain frequencies as required by the ‘443 Patent. EX1008, 8:37–47; EX1001, 17:19–27; POR, 54–57.

The specification’s intrinsic teaching underscores the difference: the ‘443 Patent distributes only an external reference to local PLLs and embraces “independent clocks with no known phase relationship,” rejecting rational-ratio constraints that “restrict the phase relationship and relative frequencies” and “impose significant restrictions” on frequency choices. EX1001, 4:38–40; 4:55–66; 6:32–36; 6:71–7:3. Georgiou is the opposite—a master-fraction scheme expressly designed to keep the system aligned.

Petitioner’s mapping also sidesteps the operational mismatch. Georgiou’s control loop explicitly adjusts operating points “while maintaining synchronization and reliability.” EX1008, 10:48–51. The ‘443 Patent, by contrast, discloses GALS behavior: synchronous within each domain and asynchronous across domains. EX1001, 17:63–18:1. A clocking fabric to maintain synchronization across units is not evidence of asynchronous inter-domain operation; it confirms that Georgiou teaches synchrony, not GALS. Petition, 68–70; POR, 54–57.

V. Shenai teaches rational multiples for synchronous comms—opposite of GALS

Petitioner’s reliance on Shenai shows a small integer multiple/division regime chosen to facilitate synchronous communication between loads: clock/shifter blocks “operate at an integer, half integer, or quarter integer multiple/division of a reference clock.” EX1004, ¶[0036]. That is the very type of phase/frequency restricted approach the ‘443 patent criticizes and distinguishes. “Some synchronization schemes restrict the phase relationship and relative frequencies of the clocks” and “impose significant restrictions on ... choices of frequencies.” EX1001, 6:32–36. The ‘443 Patent instead “recognizes the overhead associated with independent clocks with no known phase relationship.” EX1001, 6:71–7:3. Its Summary ties that choice to separate domain clocks and references only distribution to local PLLs—enabling domain local frequency control rather than master fraction selection. EX1001, 4:38–40; 4:41–49; 4:55–66.

Two consequences follow. First, clocks constrained by rational-multiple schemes are not changeable independently of the frequencies” of each other. Tethering each domain to small-integer fractions of a master reference in order to maintain synchrony constrains a domain’s frequency by design, the opposite of the ‘443 patent’s independent, domain-local control with no cross-domain relationship. EX1004, ¶[0036]; EX1001, 6:32–36; 6:71–7:3; POR, 51–53. Second, Shenai seeks synchronous operation among loads, whereas the ‘443 Patent discloses asynchrony

across domains (while permitting synchrony within a domain). EX1004, ¶[0036]; EX1001, 4:29–33; 4:38–40; 6:71–7:3; POR, 35, 51–53, 59; Reply, 21–22.

VI. Antedation — Actual reduction to practice

The Reply raises several arguments directed at Patent Owner’s antedation showing. Among them, Petitioner asserts that the invention—particularly the claimed microprocessor—requires a fabricated physical embodiment to demonstrate actual reduction to practice, and that the inventors’ use of simulation and modeling is legally insufficient. See Reply, 14–15. Petitioner also argues that the reduction to practice evidence fails to show that the invention “worked for its intended purpose.” *Id.* And with respect to method claims, Petitioner questions whether the claimed steps were actually performed in the absence of working silicon.

These arguments misframe both the legal standard for actual reduction to practice and the technological context of the invention. As shown below, the Federal Circuit has held that actual reduction to practice does not require commercial use or hardware fabrication where laboratory or simulation results demonstrate that the invention operates for its intended purpose. That principle is particularly relevant for complex systems such as microprocessors, where pre-silicon validation through simulation, timing analysis, and performance modeling is the accepted standard in the art.

A. Apparatus claims: field appropriate lab proof (including simulation) satisfies “works for intended purpose”

Actual Reduction to Practice (RTP) requires (1) an embodiment meeting the limitations and (2) proof the invention works for its intended purpose. The Federal Circuit is explicit that the second prong “need not” be proven under “conditions of actual, commercial use.” *Taskett v. Dentlinger*, 344 F.3d 1337, 1341 (Fed. Cir. 2003) (citing *Scott v. Finney*, 34 F.3d 1058, 1063 (Fed. Cir. 1994)). The court approved laboratory proof tailored to the technology, including simulation for a satellite attitude control system— “Thus, it found the simulation adequate.” *Taskett*, 344 F.3d at 1341 (discussing *Williams v. Adm’r of NASA*, 463 F.2d 1391, 1399 (CCPA 1972)). Requiring more would impermissibly demand a “viable commercial embodiment”— “[t]his the law does not require.” *Taskett*, 344 F.3d at 1342.

That principle also applies here. Cycle accurate/timing simulation and formal verification are the accepted pre-silicon means to demonstrate operability of complex microprocessors under practical use conditions (separate domain clocks/voltages; asynchronous crossings with CDC penalties; workload driven V/f schedules). (See EX2013, ¶¶42-45, also EX1033, 84 :11-87-4). The ‘443 Patent shows exactly this: (i) a flow that computes per domain V/f schedules and feeds them back to run dynamically—i.e., the claimed per domain DVFS in operation; (ii) a simulator that models independent clocks with independent jitter to capture the asynchronous crossings and timing costs; and (iii) a specification calling this an

“existence proof” that the apparatus achieves the intended energy/performance benefits. EX1001, 12:27–35; 12:41–44; 12:51–63; 12:64–71. The Reply’s refrain that “a microprocessor is a physical item” demanding fabricated silicon misstates the second RTP prong. Reply, 14–15. *Taskett/Scott/Williams* foreclose any “commercial conditions” requirement where faithful, field standard simulation demonstrates operability.

As Dr. Kaeli explained during his Deposition:

“Q. If one of your Ph.D students were to say, “I have this great branch prediction algorithm that improves performance by 50 percent,” what would you like to see the prove to you.

A. I’d like to see where they gen – where they got 50 percent. Generally, if they’re a student 10 working in my group, they would be using an execution-driven or cycle-based simulator that would, basically, be evaluating the particular they’d have a baseline design and then they’d 14 simulate their design and evaluate how much a performance changed. And I -- that would be completely acceptable and publishable and accepted by both academia and industry.

Q. So if they showed it to you, the simulation results, you would think it worked?

A. I would -- again, they would have had to demonstrate it across a representative set of workloads that, basically, are deemed, you know, appropriate for evaluating branch prediction, and -- and -- and, you know, that their results are repeatable and demonstrateable, yep.”

EX1033, 109:22-110-:25

As a result, Patent Owner’s antedation showing satisfies both RTP prongs. An embodiment meeting the claim-1 limitations was implemented in the inventors’ field-standard, pre-silicon environment: the flow computes per-domain frequency/voltage schedules and feeds them back to run dynamically (i.e., per-domain DVFS in operation), and the simulator models independent clocks with independent jitter to capture asynchronous crossings and their timing cost—evidence the apparatus works for its intended purpose, which the specification itself characterizes as an “existence proof.” EX1001, 12:27–35; 12:41–44; 12:51–63; 12:64–71. The same record supports the additional limitations of claims 2–7 (including domain-local clock/frequency control and per-domain voltage control exercised in the dynamic runs, and the modeled inter-domain behavior addressed in the CDC analysis), and it accords with how the field evaluates operability pre-silicon. As Dr. Kaeli makes clear, that cycle-based simulation demonstrating repeatable performance across representative workloads is “completely acceptable and publishable and accepted by both academia and industry.” EX1033, 110:15-17.

Petitioner’s refrain that “a microprocessor is a physical item” (Reply, 14–15) cannot raise the bar: the Federal Circuit confirms the “works for intended purpose” prong “need not” be proven under commercial conditions and expressly approves laboratory proof, including simulation; demanding a tape-out would improperly require a “viable commercial embodiment.” *Taskett v. Dentlinger*, 344 F.3d 1337,

1341–42 (Fed. Cir. 2003). On this record, the apparatus of claims 1–7 was reduced to practice prior to the critical date.

B. Method claims: steps were actually performed

For the method claims (e.g., claim 8), the antedation inquiry is whether the claimed steps were performed—not whether a chip was fabricated. The record shows they were. First, the inventors implemented the offline analysis that computes per-domain frequency/voltage schedules and feeds them back for a second, dynamic run, i.e., executing the method’s per-domain DVFS control in operation. EX1001, 12:27–35. Second, the simulation models independent clocks with independent jitter to capture asynchronous inter-domain crossings and their timing/CDC costs, matching the claim requirement that domains operate synchronously within and asynchronously across domains. EX1001, 12:51–63; 17:63–18:1. Third, the specification describes the queue/flag boundary behavior used at domain interfaces and characterizes the overall methodology as an “existence proof” of the invention’s intended energy/performance benefits—i.e., the method, as performed, worked for its intended purpose. EX1001, 12:41–44; 17:55–18:8. See also POR, 22–23 (actual RTP on a host; “worked for intended purpose”).

Element-by-element, claim 8’s asynchrony requirement appears at EX1001, 17:63–18:1 and its dynamic, per-domain control at 17:66–18:8; the offline-then-dynamic flow and the cycle-accurate simulation with independent clocks/jitter are

the performance of those steps, with concrete outputs recorded and analyzed. EX1001, 12:27–35; 12:51–63. Petitioner’s suggestion that simulation “would have had limited value ... absent a physical prototype” (Reply, 14–15) applies the wrong test to method claims: the statute asks whether the method was carried out and shown to work for its intended purpose, which the contemporaneous implementation and the specification’s existence-proof description supply. EX1001, 12:41–44; 17:55–18:8; POR, 22–23.

VII. Shenai’s provisional entitlement fails (written description)

Even apart from RTP, Petitioner still bears the §119(e) burden. Their Reply argues the 2001 provisional (Ex. 2006) suffices because a POSITA would understand a controller “capable of determining ... current ... targets” and a microprocessor context. Reply, 16–18. But priority turns on what the provisional actually describes, not on ipse dixit capability or POSITA supplementation. *Dynamic Drinkware* and *Riggs* require written description in the provisional for the specific portions relied upon—here, a controller determining current targets and a microprocessor. POR, 39–41.

Shenai’s utility recites a “controller capable of determining voltage, current and clock signal frequency targets.” EX1004, ¶[0026]; see also ¶[0031]. The provisional passages the Reply cites (dynamic controller setting voltage and frequency; cross bridge control) do not disclose current targets—Petitioner argues

only capability by inference, and the first appearance of “current targets” is in the utility, not the provisional. Reply, 16–18; POR, 39–41 (collecting the exact Ex. 2006 quotes). Under *Dynamic Drinkware/Riggs*, capability and POSITA fill ins are not written description of the specific subject matter relied upon. POR, 39–41. The Reply likewise tries to fill the absence of the word “microprocessor” in Ex. 2006 by inference. Reply, 18–19. But §119(e) requires possession in the provisional of the claimed subject matter (including the microprocessor context the Petition uses), not merely an IC arrangement that could be interpreted as such. POR, 41–43.

As a result, Petitioner’s §119(e) theory rests on ipse dixit “capability” and POSITA inference; the provisional describes neither “current targets” nor even a microprocessor as claimed—so Shenai does not antedate February 2002 on this record. Reply, 16–19; EX1004, ¶[0026], ¶[0031]; POR, 39–43.

VIII. Conclusion

Under the intrinsic record, “separately generating” and “independently changeable” require domain-local clocks without rational-ratio tethering, and GALS operation that is synchronous within a domain and asynchronous across domains. Georgiou is a selector among master-derived, predetermined fractions that aims to maintain synchronization; Shenai teaches integer/half/quarter multiples expressly to enable synchronous inter-load communication. Neither reference—alone or in combination—discloses the claimed domain-local clocks or independent

per-domain frequency control the ‘443 patent describes and claims. EX1001, 4:29–56; 6:32–36; 6:71–7:3; 17:19–27; 17:63–18:1; EX1008, 8:37–47; 8:61–67; 10:48–51; EX1004, ¶[0036]; Petition, 68–70; Reply, 11, 21–27; POR, 51–57, 59–60.

On the merits, Petitioner’s combinations do not supply the required asynchronous-across-domains operation (method) or convert multiplexer selection into a clock that separately generates (apparatus). Nor do they meet the dependent claim limitations directed to queue/flag domain-boundary behavior in a GALS setting. EX1001, 17:19–27; 17:63–18:1; 5:7–19; 12:51–63; Petition, 68–70; Reply, 21–27; POR, 54–57.

On antedation, the record shows actual reduction to practice consistent with the law and the field: a cycle-accurate/timing implementation that (i) computes and executes per-domain V/f schedules, (ii) models independent clocks with independent jitter for asynchronous crossings, and (iii) documents an “existence proof” of the intended energy/performance benefits. For apparatus claims, Federal Circuit precedent confirms that “works for intended purpose” need not be proven under commercial conditions and endorses laboratory/simulation proof as sufficient; for method claims (e.g., claim 8), the question is whether the steps were performed—and the record shows they were. EX1001, 12:27–35; 12:41–44; 12:51–63; 12:64–71; 17:55–18:8; Reply, 14–15; POR, 22–23.

Finally, Petitioner’s §119(e) rescue fails: the provisional passages it cites do not describe a controller determining current targets or even a microprocessor, and written-description law does not allow POSITA inference to supply what the provisional lacks. Reply, 16–19; EX1004, ¶[0026], ¶[0031]; POR, 39–43.

As a result, Petitioner has not shown, by a preponderance of the evidence, that the challenged limitations are taught or would have been obvious under the proper construction; Patent Owner’s antedation is established; and Shenai lacks provisional entitlement. The Board should therefore deny the requested relief and confirm the patentability of the challenged claims.

Respectfully submitted,

Dated: September 19, 2025

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's Sur-Reply complies with the type-volume limits of 37 C.F.R. § 42.24(b)(1) because it contains 4,043 words, excluding the parts of this Patent Owner's Sur-Reply that are exempted by 37 C.F.R. § 42.24(a), according to the word processing system used to prepare this Patent Owner's Sur-Reply.

Respectfully submitted,

Dated: September 19, 2025

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 SUR-REPLY** was served via email (as consented to by counsel) on September 19, 2025 to lead and backup counsel of record for Petitioners as follows:

Dinesh Melwani
dmelwani@bomcip.com

Ankit Aggarwal
aaggarwal@bomcip.com

William Uhr
wuhr@bomcip.com

Bookoff McAndrews, PLLC
2000 Pennsylvania Avenue NW
Suite 4001
Washington, D.C. 20006
Service email: docketing@bomcip.com

Dated: September 19, 2025

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