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.

IPR No. IPR2024-01225

Patent No. 7,089,443 B2
Filing Date: January 23, 2004
Issue Date: August 8, 2026

DECLARATION OF DR. DAVID KAELI IN SUPPORT OF PATENT OWNER'S RESPONSE

Table of Contents

I. Introduction And Qualifications 3
II. Summary Of Materials Reviewed And Considered 6
III. Legal Standards 8
IV. Level Of Skill In The Art And Perspective Applied In This Declaration 10
V. The '443 Patent 13
1. Sarmenta (Ex. 2004) 15
2. Pratt (Ex. 2011) 18
3. Simplescalar Tool 21
VI. Claim Construction 23
VII. The Claims Of The '443 Patent Would Not Have Been Obvious 29
A. Cited Prior At 30
1. Overview of Shenai 30
2. Overview of Halepete 32
3. Overview of Sherburne 34
4. Overview of Traylor 36
5. Overview of Georgiou 36
6. Overview of Kranich 43
B. Ground 1: Claims 1, 3-5, 8, 9 and 12 Would Not Have Been Obvious Over Shenai 45
C. Ground 2: Claims 4 and 9 Would Not Have Been Obvious Over Shenai in view of Halepete 46
D. Ground 3: Claims 5-7 and 10-11 Would Not Have Been Obvious Over Shenai in view of Sherburne 48
E. Ground 4: Claims 7 Would Not Have Been Obvious Over Shenai in view of Traylor 49
F. Ground 5: Claims 1 and 3 Would Not Have Been Obvious Over Georgiou 50

G. Ground 6: Claims 4 and 9 Would Not Have Been Obvious Over Georgiou in view of Halepete 55

H. Ground 7: Claims 5 and 10 Would Not Have Been Obvious Over Georgiou in view of Kranich 56

VIII. Conclusion 58

I. Introduction And Qualifications

1. I, David Kaeli, submit this declaration on behalf of Intellectual Ventures I LLC ("Patent Owner") to provide my opinions concerning the validity of claims 1 and 3-12 ("Challenged Claims") of U.S. Patent No. 7,098,443 B2 (the "'443 patent," Ex. 1001), in support of Patent Owner's Response to the Petitioner for *Inter Partes* review of the '443 patent, which I understand was filed by Lenovo Group Ltd. ("Petitioner") and accompanied by a Declaration of Dr. Jacob Baker (the "Baker Declaration"), among other exhibits.

2. I am more than eighteen years of age, and I am a citizen of the United States, currently residing in the Commonwealth of Massachusetts.

3. In formulating my opinions, I have relied upon my knowledge, training, and experience in the relevant art. My qualifications, along with a more detailed summary of my background and publications, are stated more fully in my *curriculum vitae*, which has been provided as Exhibit 2002. Here, I provide a brief summary of my qualifications.

4. I am currently the Director of Northeastern University's Computer Architecture Research Laboratory and a Distinguished Professor of Electrical and Computer Engineering. I have been a faculty member at Northeastern since September 1993, in the Electrical and Computer Engineering Department.

5. For over thirty years, I have pursued research across a broad range of computer architectures and technologies, including memory systems, storage systems, virtualization, and secure computing systems. I have published over 400 refereed journal and conference papers in the areas of computer architecture and computer security. I also have supervised more than fifty Ph.D. theses in these areas.

6. I am a Fellow of the Institute of Electrical and Electronic Engineers (IEEE) and a Fellow of the Association of Computing Machinery (ACM). I have served two terms as the Chair of the Executive Committee of the IEEE Technical Committee on Computer Architecture. I am serving in my second term as the Editor-in-Chief of the ACM Transactions on Architecture and Code Optimization. I am also an Associate Editor for Elsevier's Journal of Future Generation Computer Systems.

7. Before I joined Northeastern's faculty in 1993, I spent twelve years working for IBM. From 1981-1986, I worked on the design and delivery of the IBM System/370 3090 mainframe system. From 1986-1993, I worked at the IBM T.J. Watson Research Center, designing, and analyzing IBM Personal Systems 2 computers and associated software. During that time, I also worked on the design of a new Application System/400 multi-processor system.

8. I am an active educator and researcher and have extensive experience in the broader field of computer architecture. I have conducted academic research on, and published multiple papers relating to dynamic voltage and frequency scaling in the context of cybersecurity attacks, as well as energy-efficient microarchitectures. *See, e.g.*, Sabbagh et al., *Overdrive Fault Attacks on GPU*, IEEE/IFIP International Conference on Dependable Systems (2021); Sabbagh et al., *A Novel GPU Overdrive Fault Attack*, ACM/IEEE Design Automation Conference (2020); Tavana et al., *Exploiting Adaptive Data Compression to Improve Performance and Energy-efficiency of Compute Workloads in Multi-GPU Systems*, 33rd IEEE International Parallel and Distributed Processing Symposium (2019); and Ning et al., Power Aware External Bus Arbitration for System-on-a-Chip Embedded Systems, *High Performance Embedded Architectures and Compilers* (2005).

9. I have also delivered many tutorials and conducted research in cycle-based simulation methods and technology. See e.g., Sun et al., MGPUSim: Enabling Multi-GPU Performance Modeling and Options, *ACM/IEEE 46th International Symposium on Computer Architecture* (2019), Ubal et al., Multi2Sim: A Simulation Framework for CPU-GPU Computing, *ACM 21st International Conference on Parallel Architectures and Compilation Techniques* (2012), and Kaeli, Issues in Trace-driven Simulation, *IFIP International*

Symposium on Computer Performance Modeling, Measurement and Evaluation (1993).

10. I earned a Bachelor of Science degree in Electrical Engineering from Rutgers University in 1981. I earned a Master of Science degree in Computer Engineering from Syracuse University in 1985. I earned a Ph.D. in Electrical Engineering from Rutgers University in 1992.

II. Summary Of Materials Reviewed And Considered

11. The opinions contained in this Declaration are based on the documents I reviewed and my knowledge and professional judgment. In forming the opinions expressed in this Declaration, while drawing on my experience, I reviewed the following documents:

	Description
EX1001	United States Patent No. 7,089,443 B2 ("the '443 patent")
EX1002	Declaration of Dr. Jacob Baker ("Baker Decl.")
EX1004	U.S. Patent Publication No. 2003/0030326 A1 to Shenai et al. ("Shenai") published on February 13, 2003
EX1005	International Publication No. WO 01/53921 A1 to Halepete et al. ("Halepete") published on July 26, 2001
EX1006	U.S. Patent Publication No. 2002/0169990 A1 to Sherburne ("Sherburne") published on November 14, 2002
EX1007	U.S. Patent No. 5,473,756 to Traylor ("Traylor") issued on December 5, 1995

	Description
EX1008	U.S. Patent No. 6,047,248 to Georgiou et al. ("Georgiou") issued on April 4, 2000
EX1009	International Publication No. WO 01/33352 A1 to Kranich et al. ("Kranich") published on May 10, 2001
EX2002	David Kaeli CV
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
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.
Paper 11	Institution Decision

12. My opinions are additionally guided by my appreciation of how a person of ordinary skill in the art would have understood the claims of the '443 patent at the time of the invention, which I have been asked to assume is February 2, 2002. I understand from counsel that this is the publication date for a paper titled "Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic

Voltage and Frequency Scaling" (Semeraro et al., Ex. 2005). I also understand that Column 5, Lines 22-27, of the '443 Patent states: "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." Accordingly, I will use the February 2, 2002 priority date when judging relevant prior art in this discussion.

III. Legal Standards

13. I am not an attorney. However, I have been informed by counsel of certain legal standards and have applied those standards, described below, in formulating my opinions.

14. I understand that the claims of a patent define the limits of a patentee's exclusive rights. I understand that to determine the scope of the claimed invention, courts may construe claim terms, the meaning of which the parties may dispute. I understand that claim terms should generally be given their ordinary and customary meaning as understood by a person of ordinary skill in the art at the time of filing of the patent application.

15. I understand that claims must be construed in light of and consistent with the intrinsic evidence. In this context, I understand that intrinsic evidence includes the claims themselves, the written disclosure in the specification, and the patent's

prosecution history. I understand that the specification is always highly relevant to the claim construction analysis and often is the single best guide to the meaning of a disputed term. I understand that extrinsic evidence may also be considered when construing claims and may include, for example, technical dictionaries, technical publications and books, treatises, and expert testimony, though the intrinsic record will be more persuasive as to the meaning of a particular term.

16. I understand that one or more prior art references can render a patent claim obvious to a person of ordinary skill in the art if the differences between the subject matter set forth in the patent claim and the prior art are such that the subject matter of the claim would have been obvious at the time the claimed invention was made. In analyzing obviousness, I understand that it is important to consider the scope of the claims, the level of skill of a person of ordinary skill in the relevant art, the scope and content of the prior art, the differences between the prior art and the claims, and "secondary considerations" of non-obviousness, which I also understand can include, for example, evidence of commercial success of the invention, evidence of a long-felt need that was solved by an invention, evidence that others copied an invention, or evidence that an invention achieved an unexpected result.

17. I understand that claimed subject matter can be obvious when a person of ordinary skill in the art would have been motivated to combine pre-existing

elements, and that combination yields no more than what a person of ordinary skill in the art would have expected. I also understand that, in assessing whether a claim is obvious, one must consider whether the claimed improvement is more than the predictable use of prior art elements according to their established functions. I understand that there need not be a precise teaching in the prior art directed to the specific subject matter of a claim because one can take account of the inferences and creative steps that a person of ordinary skill in the art would employ. However, I further understand that obviousness cannot be based on the hindsight combination of components selectively culled from the prior art.

18. I understand that for a reference to qualify as a relevant prior art reference in an obvious analysis, the reference must be "analogous"—meaning it is from the same field of endeavor as the patented technology, and if not, then the reference must be reasonable pertinent to the particular problem facing the inventors of the patent at issue. I also understand that a reference is reasonably pertinent if a person of ordinary skill in the art would have reasonably consulted those references and applied their teachings in seeking a solution to the problem facing the inventors.

IV. Level Of Skill In The Art And Perspective Applied In This Declaration

19. I understand that certain issues relating to validity must be judged from the perspective of a person of ordinary skill in the art ("POSITA"). I further

understand that factors that may be considered in determining the level of skill of a person of ordinary skill in the art include, but are not limited to, (1) the type of problems encountered in the art; (2) prior art solutions to those problems; (3) rapidity with which innovations are made; (4) sophistication of the technology; and (5) the educational level of active workers in the field.

20. I have been informed and understand that a person of ordinary skill in the art is a hypothetical person who is presumed to be aware of all pertinent prior art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. Importantly, this hypothetical person is of ordinary skill in the art to which the claimed subject matter pertains.

21. With respect to the '443 patent, a person of ordinary skill in the art at the time of the earliest priority date, February 2, 2002, would have been familiar with techniques that could manage voltage and frequency on a microprocessor, and underlying synchronization technologies used with such techniques.

22. A person of ordinary skill in the art would have had at least a Bachelor's degree in computer science, computer engineering, or a related field, and two years of industry experience or academic research experience on microprocessor design. In my view, more education can supplement relevant experience and vice versa.

23. At the time of the earliest priority date for the '443 patent (February 2, 2002), I met the qualifications to be a person of ordinary skill in the art, and I

believe that I also have a sufficient level of knowledge, experience, and education to provide expert opinions in the field of the '443 patent.

24. While my own level of skill exceeded that of the ordinary level of skill as of February 2002, I am well-acquainted with the actual performance and capabilities of a person of ordinary skill in the art as defined above. For example, during the relevant timeframe, I was supervising and evaluating the performance of such individuals, such as my graduate students who were performing research in this area. Furthermore, my participation in professional organizations and in conferences and other events exposes me to the state of the industry and the level of experience of individuals who, like I did, had at least the ordinary level of skill.

25. I have applied the above-defined perspective of a person of ordinary skill in the art, as of February 2002, in forming and expressing my opinions. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art at that relevant time.

26. In the Baker Declaration, Dr. Baker sets out the following articulation of the level of ordinary skill:

> In my opinion, based on the materials and information I have reviewed and based on my experience in the technical areas relevant to the '443 Patent, a POSITA at the time of the alleged invention of the '443 Patent would have had a bachelor's degree in electrical engineering, computer engineering, or the equivalent, and two to three years of experience designing computing circuits and systems. More

> practical experience could offset the aforementioned education while a higher level of education could offset a lesser amount of practical experience.

Ex. 1002 (Baker Decl.), ¶22. At the time of the earliest priority date for the '443 patent (February 2002), I met Dr. Baker's proposed qualifications to be a person of ordinary skill in the art, and my opinions in this Declaration would not change if the Board were to adopt and apply that proposed level of ordinary skill.

V. The '443 Patent

27. The '443 patent is titled "Multiple Clock Domain Multiprocessor" and relates to methods for designing a globally synchronous, locally asynchronous (GALS) microprocessor, '443 patent, Cover.

28. The '443 patent explains that, around the time of its priority date (February 2002), microprocessor vendors were enjoying the ability to scale frequency, but this scaling had recently been impacted by issues related to wire delays '443 patent, 1:29-37. The '443 patent described problems due to clock skew, which raised new questions about whether designers should consider the potential for asynchrony in their designs. Ex. 1001 '443 1:37:45.

29. The '443 patent discusses problems that occur when leveraging asynchronous clocks in a single design. For example, the '443 patent discusses "Two major reasons for this reluctance are the immaturity of asynchronous design tools relative to those in the synchronous domain, and the cost and risk of moving

away from the mature design infrastructures that have been successfully used to create many generations of microprocessor products." Ex. 1001, 1:50-55. The '443 patent also explains ways the industry had tried to deal with such issues. For example, the '443 patent discusses how "several multiprocessor systems run the memory bus off of a different clock than the processor core," where "the logic in each of the two clock domains is designed using conventional synchronous design methodologies." Ex. 1001, 1:56-1:62.

30. The '443 patent distinguishes between earlier art that introduced GALS designs, and a multiple clock domain (MCD), where "each functional block operates with a separately generated clock, and synchronizing circuits ensure reliable inter-domain communication." Ex. 1001, 4:17-19. Further, the '443 patent explains "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." Ex. 1001, 4:21-24. The '443 patent discusses how the use of a local phase lock loop (PLL) can be used to support independence between the clocks used for each domain. This avoids issues of clock skew that have been a growing issue as design frequencies and chip sizes have increased and allows designers to utilize higher frequency clocks. Further, "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, than can be achieved with single clock, single-core-voltage systems." Ex. 1001, 4:46-50. 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.

31. The '443 patent also recognizes that prior methods for supporting a multiple clock domain (MCD) globally-asynchronous locally-synchronous (GALS) clocking did not address issues related to supporting synchronization and communication between domains. Ex. 1001, 1:62-65. The '443 patent teaches that "[s]ome synchronization schemes restrict the phase relationship and relative frequencies of the clocks, thereby eliminating the need for hardware arbitration [27]." (Ex. 1001 6:33-35). The '443 patent includes a citation [27] "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." Ex. 1001 1:54-56. (Ex. 2005) (hereinafter, "Sarmenta"), which provides an example of a GALS design with one exemplar of dependent clock control.

1. Sarmenta (Ex. 2004)

32. 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." Ex. 2005 Sarmenta Abstract. 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.

33. As an example, as described by Sarmenta, if a first clock operates at a frequency (f_m) and a second clock at (f_n), their frequencies satisfy the relationship ($f_m/f_n = M/N$). When one clock is a multiple of the other clock, then periodic alignment is possible, described by Sarmenta as a "coincidence cycle" Ex. 20005 (Sarmenta) 271 Section 2. This duration of a coincidence cycle is the least common multiple (LCM) of the two clock periods, and during this interval, system *m* completes exactly MM cycles while system *n* completes NN cycles. The significance of the coincidence cycle is that within it, the relative timing of all active edges of both clocks is fixed and can therefore be analyzed and used to precompute which moments are safe for data transfers between the two systems. The result is that the clock edges of the two clocks will eventually align periodically. This predictable behavior is key to managing data transfers between different domains, without requiring further synchronization operations. For

example, if M=5M = 5 and N=6N = 6, the coincidence cycle would span 30 ticks of the global clock which corresponds to 5 clock cycles of system m and 6 clock cycles of system n, and the safe and unsafe transfer opportunities can be determined once and reused in every repetition of the cycle. Thus, the coincidence cycle serves as the basic scheduling unit that enables reliable and high-throughput communication between rationally-clocked systems.

34. However, as the '443 patent discusses, when introducing the Sarmenta example, "these schemes impose significant restrictions on the possible choices of frequencies." Ex. 1001 6:35-37. To remedy the limitations of systems such as Sarmenta, the '443 patent explains that "[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 speeds, latency, and the exploitation of application parallelism via complex hardware structures." Ex. 1001, 2:34-39. While Sarmenta introduces a synchronization approach leveraging rational frequency relationships, Sarmenta's scheme imposes significant restrictions on frequency selection, as noted in the '443 patent. Ex. 1001 6:33-44. For example, Sarmenta cannot support a situation where a coincidence point does not exist (e.g., multiple clock domains running at that same clock frequency but starting at different times) or the coincidence point is unknown in advance (e.g., multiple clock domains that can each independently change its

respective clock frequency without constraint). 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.

2. Pratt (Ex. 2011)

35. The publication by Sarmenta et al., cited in the record as Ex. 2004, states that "a U.S. patent application has been filed on behalf of MIT covering portions of this work." Ex. 2004, p. 1. That referenced application is U.S. Patent Application No. 07/929,879, which matured into U.S. Patent No. 5,634,041, titled "Rationally Clocked Communication Interface" (hereinafter, "Pratt"), issued to Massachusetts Institute of Technology on behalf of Gill Pratt and Stephen Ward—both of whom are co-authors of the Sarmenta publication.

36. Pratt elaborates upon the concept briefly introduced in Sarmenta, namely the use of rationally related clocks—clock signals whose frequencies are in a ratio of two integers—as a means of achieving deterministic synchronization between processing units. While Sarmenta refers to this recurring synchronization interval as a "coincidence cycle," Pratt refers to it as the "beating period." Both terms describe the same core idea: a predictable interval over which two clocks realign in phase, enabling reliable data communication without asynchronous handshaking. Ex. 2011, 2:12-23.

37. As explained in Pratt, when two clocks operate at frequencies Fm and Fn such that Fm / Fn = M / N, where M and N are small integers, their least common multiple defines a recurring interval—the beating period—in which their respective clock edges periodically coincide. Pratt expresses this relationship mathematically as M × Tm = N × Tn, where Tm and Tn are the periods of the respective clocks. This beating period defines a deterministic cycle over which the relative phase of the two clocks repeats. Ex. 2011, 3:65-4:29

38. Pratt further discloses that this periodicity is exploited to precompute a schedule of valid and invalid communication cycles between the transmitting and receiving systems. During each beating period, only certain transmitter clock cycles are designated as "valid" for data transmission, and only certain receiver clock cycles are valid for sampling. This schedule is determined based on timing constraints such as data propagation delay and setup/hold time requirements. Pratt provides Figure 2 (reproduced below) to illustrate the "beating period." Ex. 2005, 4:22-28.

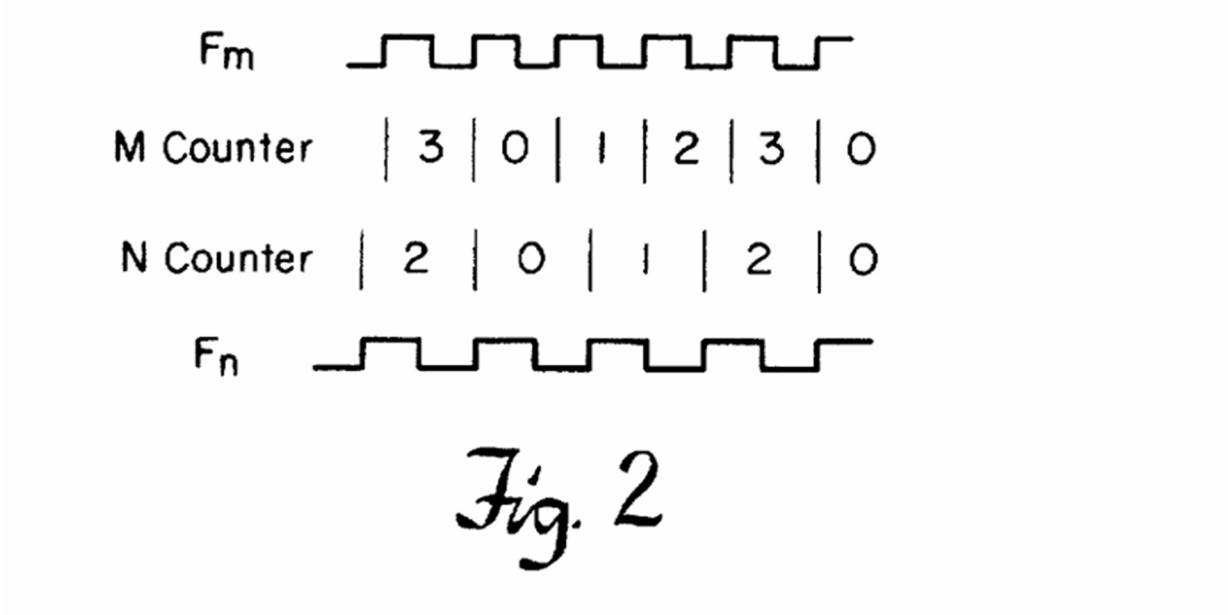


Fig. 2

Ex. 2011, Figure 2

39. By using the beating period to constrain data transfers to known-safe intervals, Pratt avoids the need for asynchronous handshaking protocols or metastability-hardening techniques. The system relies solely on the pre-characterized clock relationship to ensure timing correctness. As a result, the beating period functions as the fundamental temporal framework for the system's synchronization strategy—coordinating the exchange of data between independently clocked domains with precision and determinism. Ex. 2005, 1-5-19.

40. In this way, Pratt provides the detailed implementation and theoretical foundation that underpins the concept of the "coincidence cycle" referenced in Sarmenta. The "coincidence cycle" and "beating period" are functionally synonymous terms for the recurring alignment window of two rationally related

clocks, and both serve as the basis for the synchronization scheme described across both documents. Ex. 2011, 4:5-15.

41. This predictable timing structure yields a powerful architectural advantage: it enables the system to precompute a schedule of valid and invalid transmission and reception intervals. As a result, data exchanges between independently clocked systems can be precisely timed to occur only during safe windows—thereby avoiding metastability and eliminating the need for asynchronous handshaking protocols or dynamic arbitration logic. Communication proceeds deterministically and with high throughput, even across clock domains operating at different speeds. In this context, the coincidence cycle—or "beating period"—serves as the fundamental mechanism that transforms what would otherwise be an asynchronous communication challenge into a synchronously solvable problem, achieved through rational clocking and deterministic scheduling.

3. Simplescalar Tool

42. Counsel has asked me to discuss the Simplescalar Tool that is mentioned in the Semeraro paper. In the same timeframe as the Semeraro paper (2002), the computer architecture community, including my own research group, was using the Simplescalar Tool in our research.

43. 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. See Ex. 2010

44. Claim 1 of the '443 Patent recites "a plurality of domains' each of which includes "a clock for separately generating a clock signal at a frequency for that domain ..." and "a voltage input for receiving a voltage which is dynamically changeable ...". When 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. In addition, 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.

45. Industry relies heavily on simulators (such as SimpleScalar) to demonstrate the operation of architectural designs. Similarly, in academia, research papers in the field of computer architecture often include simulation results from simulators, enabling the research community to effectively quantify the benefits a design. In both research and academia, results on a validated simulator (such as SimpleScalar) are typically accepted as evidence that an architectural design has been sufficiently tested to demonstrate that it will work for its intended purpose. In addition, in the context of the Paper, given the details and thoroughness of the simulation as described herein, it is my opinion that others 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. At that point, skilled artisans in the field could have implemented the architectural design in other forms, such as a stand-alone silicon chip, through application of ordinary skills.

VI. Claim Construction

46. In my analysis, I have applied the plain and ordinary meanings of the claim terms in the '443 patent, as a person of ordinary skill in the art would have understood them in the context of that patent in approximately February 2002. And it is my opinion that the prior art set forth in the Petition fails to disclose or render

obvious the subject matter of the challenged claims under any reasonable construction and interpretation of the challenged claims' terms, including their plain meanings, as well as the constructions Petitioner's expert, Dr. Baker, proposes.

47. Based on the '443 patent's claims and the patent's specification, a person of ordinary skill in the art would understand that the claim language in claims 1 and 8 describe a multiple-clock domain microprocessor, where the multiple clocks are independent of each other, and where the input voltages are independent of each other. Further, the frequency of each clock in a domain can be changed independently of other clocks, and the voltage value for each domain can be changed independently of the voltage values of other domains.

48. Skilled artisans would also understand these limitations because they would immediately recognize that ability to design a multi-domain microprocessor with independent clock domains and independent voltage domains provides a new degree of flexibility not present in microprocessors in February 2002.

49. My understanding of the plain and ordinary meaning of the terms utilized in the '443 Patent was developed after reading the patent and the references incorporated in the Patent. From this perspective, the plain and ordinary meaning of the claims in the '443 Patent reflects a system in which each clock domain generates its own clock signal and is capable of dynamically changing its

frequency independently of the other domains. That is, each domain's frequency selection is not constrained by or dependent on a predetermined small integer relationship with the frequency of any other domain.

50. This understanding is confirmed by the specification, which repeatedly emphasizes the design goal of independence across clock domains. At column 6, lines 32–36, the specification contrasts the disclosed invention with synchronization schemes that "restrict the ... relative frequencies of the clocks," specifically criticizing such systems for imposing "significant restrictions on the possible choices of frequencies." These statements clearly inform a POSITA that the invention is designed to avoid architectures where clock domains must maintain a small-integer relationship to each other.

51. The references cited in the specification—particularly Sarmenta and by extension Pratt—reinforce this distinction. As a POSITA, I recognize that Sarmenta and Pratt describe systems that rely on rational clocking—that is, systems where the clock frequencies of different domains must maintain a ratio of small integers (e.g., 3:2, 5:4). These systems exploit the predictability of phase alignment over a fixed interval (referred to as a "coincidence cycle" or "beating period") to schedule communication without requiring run-time arbitration.

52. From a technical perspective, these rational clocking systems necessarily impose frequency constraints on each domain. In order to maintain the

fixed alignment interval that enables deterministic communication, each domain's clock frequency must be selected from a narrow set of permitted values that satisfy the required rational ratio. A POSITA reading Sarmenta or Pratt would understand that such constraints fundamentally limit the system's ability to independently vary domain frequencies.

53. In contrast, the '443 Patent advocates for an approach that frees each domain from such constraints. It states that "designers of each domain [are] no longer constrained by the speeds of critical paths in other domains" and emphasizes that each domain may "optimize the tradeoffs among clock speed, latency, and ... parallelism" independently. Ex. 1001 at 6:36–39. A POSITA would recognize that such independence is inconsistent with any scheme requiring fixed or known relationships among clock domains.

54. Although not specifically cited in the '443 Patent, "Efficient Self-Timed Interfaces for Crossing Clock Domains" by Chakraborty (hereinafter, "Chakraborty") confirms that rational clocking systems—defined by known integer frequency ratios—are fundamentally distinct from systems with arbitrary clock frequencies.

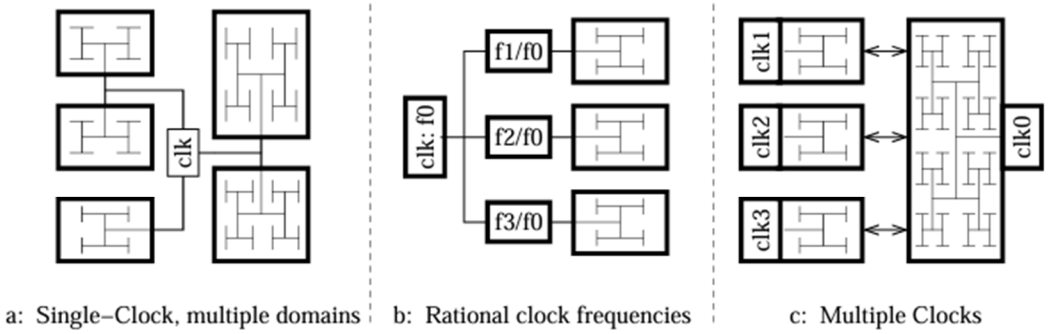


Figure 1. Chips with Multiple Clock Domains

55. Specifically, Figure 1(b) (reproduced above) of Chakraborty illustrates a chip architecture in which different regions operate at distinct clock frequencies that are rationally related—that is, their frequencies are integer multiples or ratios of one another—and all clocks are derived from a common source. Chakraborty explains that this configuration is typical in system-on-chip (SoC) and multi-rate digital signal processing designs, where different IP blocks are optimized to run at different rates based on performance or power requirements. Ex. 2012, 2.

56. Chakraborty further states:

> "Although different domains operate at different frequencies, these frequencies are all rational multiples of each other, and these ratios are typically known in advance or are determined by pre-designed operating modes of the design. In this case, we can take advantage of knowing the exact relationship between the various clocks to implement an interface that operates with low latency and without synchronization." Ex. 2012, 2.

57. Chakraborty also makes clear that systems utilizing rational clock frequencies are different from systems that operate at arbitrary frequencies. Specifically, Chakraborty explains:

> "As described above, we present designs when the sender and receiver operate at exactly the same frequency; at frequencies that are rational multiples of each other; at closely matched frequencies; and at arbitrary, relatively stable frequencies." Ex. 2012, 2.

58. To emphasize this point, Chakraborty dedicates entirely separate sections of the paper to "Rational Clock Frequency Multiples" (see Ex. 2012, 6) and "Arbitrary Clock Frequencies" (see Ex. 2012, 9), reinforcing that these are treated as distinct classes of systems.

59. Accordingly, Chakraborty makes clear that rational clocking systems—where clock frequencies are related by a known ratio of small integers—are fundamentally different from systems that operate with arbitrary or unrelated frequencies. In rational clocking, the known relationship between clocks enables the design of interfaces that operate with low latency and without synchronization, leveraging the predictability of clock phase alignment. This stands in contrast to arbitrary clocking systems, which do not rely on any predefined frequency relationship and therefore must account for unknown or drifting phase differences, typically requiring more complex synchronization mechanisms. Chakraborty expressly distinguishes between these two classes of systems, devoting separate

sections of the paper to "Rational Clock Frequency Multiples" and "Arbitrary Clock Frequencies," thereby underscoring that rational clocking is not simply a special case within a broader asynchronous framework, but rather a distinct and recognized approach that depends on a known, rational frequency relationship between clock domains.

60. Based on the specification and the incorporated references, it is my opinion that a person of ordinary skill in the art would understand the claim language—particularly the phrases "independently of the frequencies" (claims 1 and 12)—to require that each domain's frequency be selectable and dynamically changeable without being tethered to, derived from, or constrained by the frequencies of other domains. A system in which domain frequencies are required to bear a fixed rational relationship (as in Sarmenta or Pratt) would not satisfy the plain and ordinary meaning of these limitations as understood by a POSITA.

VII. The Claims Of The '443 Patent Would Not Have Been Obvious

61. Petitioner challenges claims 1 and 3-11 of the '443 patent based on seven grounds of purported obviousness, each of which depends on prior art or combinations of prior art.

62. In my opinion, none of the claims of the '443 patent would have been obvious over Petitioner's prior art, for reasons I explain below. I will first review the Petitioner's prior art. Then I will address each of the seven grounds.

A. Cited Prior At

1. Overview of Shenai

63. Shenai is titled "Distributed Power and Supply Architecture" and discloses a power management system that is able to provide power to multiple loads, and includes a clock distribution controller that is able to adjust voltages, current and clock frequencies for the loads. Ex.1004, Title, Abstract, 1:0005. Figure 4 illustrates the primary components of the system:

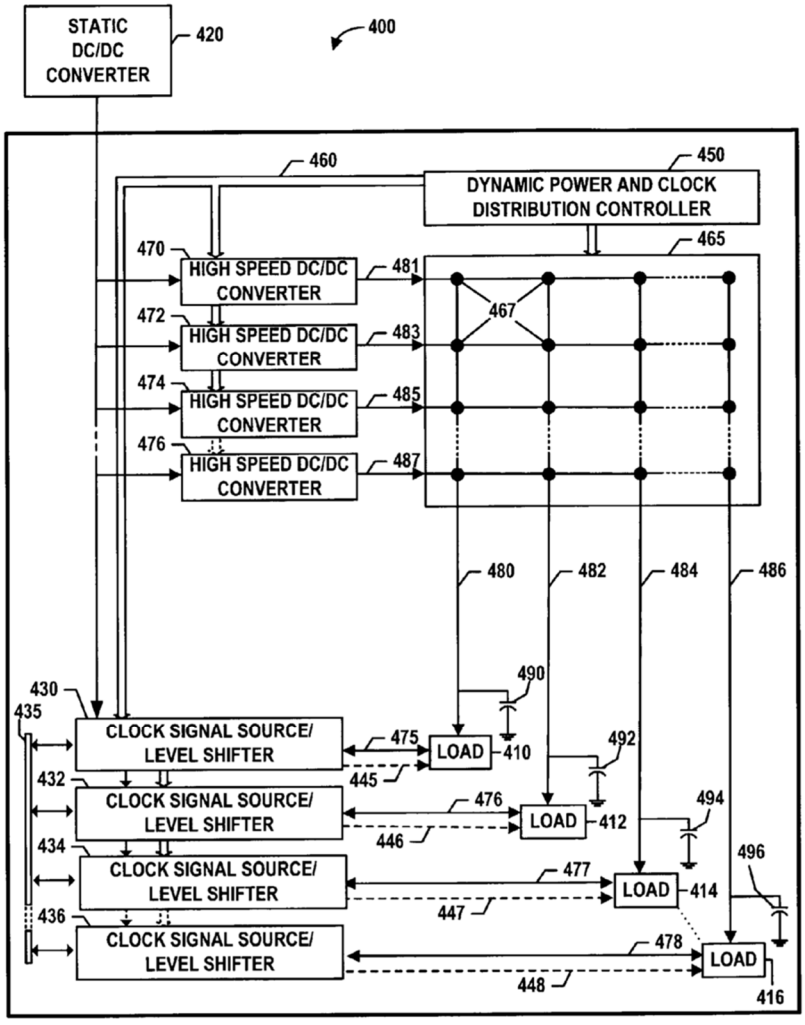


FIGURE 4

64. As seen in Figure 4, Shenai discloses embodiments of a clock and voltage distribution network Ex. 1004, 0010. As described by Shenai, "These circuits may operate at an integer multiplication, a half-integer multiplication or a quarter-integer multiplication of a reference clock signal." Ex. 1004, 0036. Further, referring to these multiple clock frequencies, all rational multiples of a common clock, Shenai adds "Such configurations may allow for synchronous communication between loads 410-416, as opposed to asynchronous communication, which may be undesirable." Ex. 1004, 0036.

65. In terms of distributing voltages, again referencing Figure 4, Shenai states "Circuit 400 may also include plural high-speed DC/DC converters 470, 472, 474 and 476….. DC/DC converters 470-476 may be responsive to control 450, such that they are selectively coupled with loads 410-416 to provide the respective target voltages and current to the loads." Ex. 1004, 0032. As shown in Figure 4, the high-speed DC/DC converters 470-476 are fed off a common voltage source 420.

66. It should be noted that Shenai is focused on supplying a clock and voltage source to an analog to digital converter, in contrast to the microprocessor targeted in the '443 patent Ex. 1004, 3:0025. Further, Shenai's characterization of systems requiring asynchronous communications as "undesirable" would not

motivate a POSITA to implement a globally-asynchronous architecture as required by the ‘443 Patent..

2. Overview of Halepete

67. Halepete is titled “Adaptive Power Control” and presents a method for managing the power supplied to a computer, monitoring the operating characteristics of central processor, and adjusting “the frequency and voltage steps”. Ex. 1005, Cover, Abstract.

68. More specifically, Halepete describes “…methods for varying the amount of power used by such systems during use of the systems” Ex. 1005, 1:0008. Haleprete describes methods to “extending the operating life of a portable computer” Ex. 1005, 2:20-21. Figure 1 from Halepete is shown below.

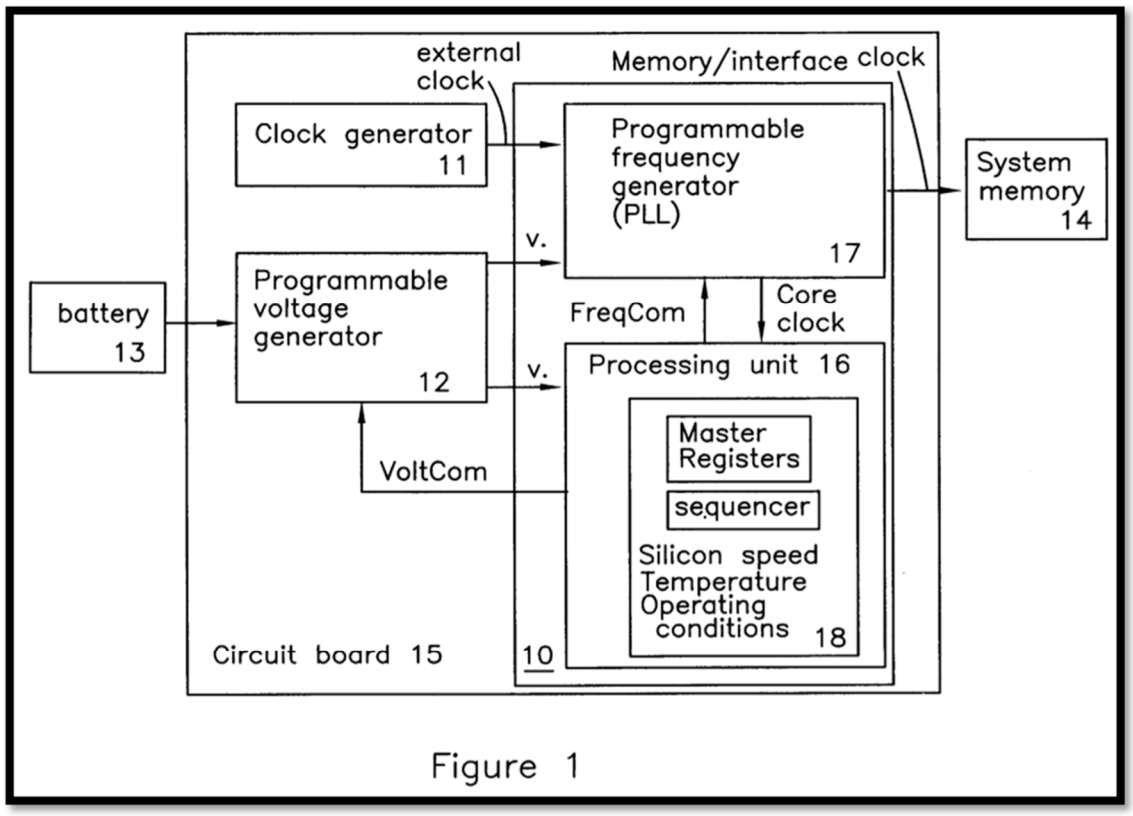


Figure 1 of *Halepete*

69. Figure 1 includes a clock generator (11) that feeds a programmable frequency generator (17), providing a core clock to a processing unit (16). The external clock is referred to as the "slow clock" Ex. 1005, 6:4, which is a common clock used to generate the core clock and the memory system clock. "The core frequency for the processing unit 16 is generated by multiplying the slow clock by a factor." Ex. 1005, 7:1-2. Further, "The frequencies at which the other components

of the system operate are determined from the core frequency determined by multiplying the slow clock by the core processor factor." Ex. 1005, 7:8-10

3. Overview of Sherburne

70. Sherburne is titled "Low Power Clocking Systems and Methods" and is generally directed toward enabling "one or more processing units, each unit having a clock input that controls the performance of the unit; one or more clock controllers having clock inputs coupled to the clock inputs of the processing units…". Ex. 1006, Cover, Abstract.

71. More specifically, Sherburne describes "A low power a reconfigurable processor core includes one or more processing units, each unit having a clock input that controls the performance of the unit; one or more clock controllers having clock outputs coupled to the clock inputs of the processing units, the controller operating varying the clock frequency of each processing unit to optimize speed and processing power for a task…" Ex. 1006, 1:0008. Sherburne describes first-in first-out buffers that can communicate between processing units. Figure 3 from Sherburne is shown below.

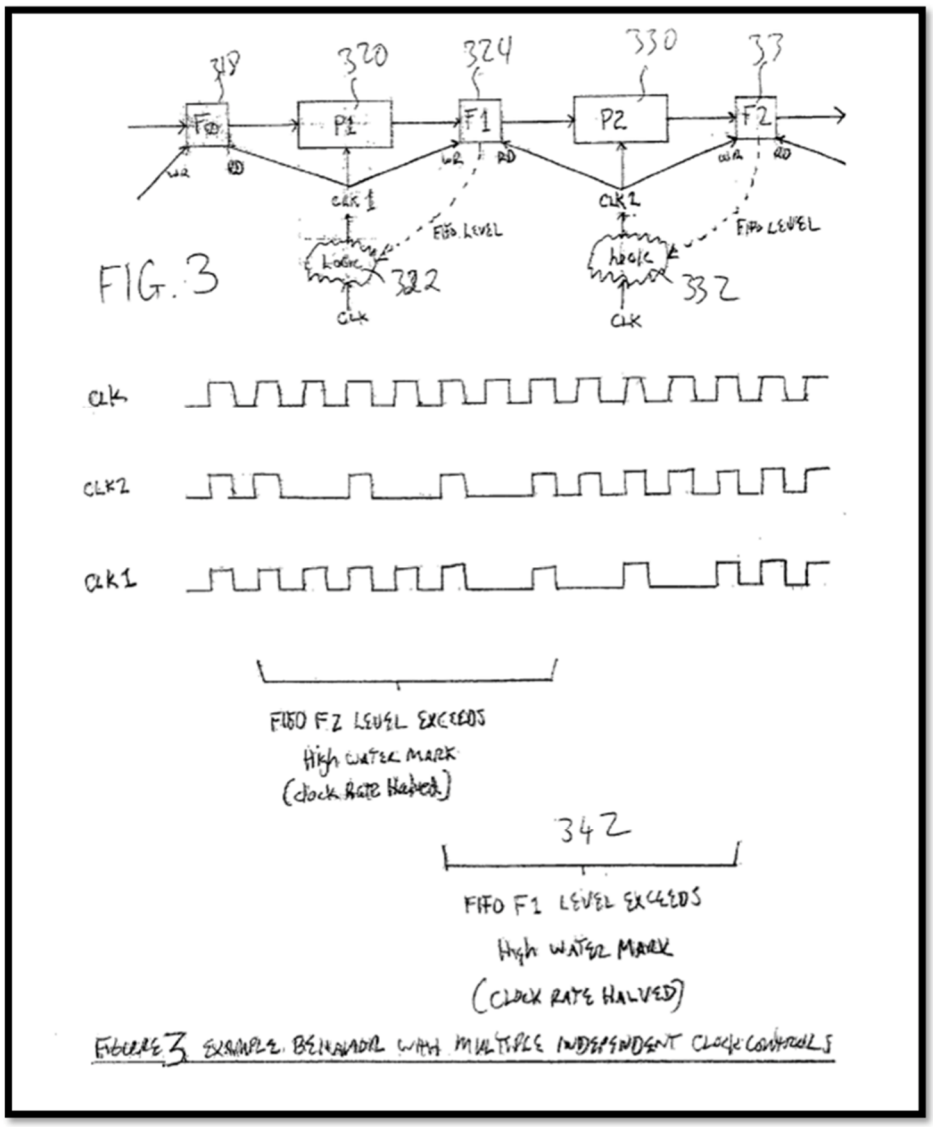


Figure 3 of *Sherburne*

72. As shown in Figure 3, Sherburne describes a common clock that feeds logic to produce processing unit clocks. Sherburne also includes FIFO buffers between processing units, providing a path for communication between adjacent processing units. Sherburne also describes a high-water mark in Figure 3, indicating the clock frequency needs to be halved Ex. 1006, 3:0029.

4. Overview of Traylor

73. Traylor is titled "FIFO Buffer with Full/Empty Detection by Comparing Respective Registers in Read and Write Circular Shift Registers." Traylor describes methods for avoiding signal glitches that may occur in a First In First Out (FIFO) buffer, Ex. 1007, Cover, Abstract. The methods include a mechanism for tracking the status of the FIFO, indicating when the FIFO is EMPTY or FULL. Additionally, signals are generated that indicated when the FIFO is ALMOST EMPTY and ALMOST Full. Ex. 1007, Cover, Abstract.

74. Traylor describes methods for "generating control and flag signals for a First In First Out (FIFO) buffer directly on a read or write cycle" Ex. 1007, 2:19-20. There is no mention of clocking or voltage in this patent.

5. Overview of Georgiou

75. Georgiou is titled "Performance-Temperature Optimization by Cooperatively Varying the Voltage and Frequency of a Circuit." Georgiou discusses a system and method "using thermal feedback to cooperatively vary a voltage and frequency of circuit to control heating while maintaining synchronization." Ex. 1008, Cover, Abstract. Georgiou is focused on changing voltage-frequency pairs to adapt a circuit based on thermal sensor feedback. Georgiou explains that "Reducing the voltage swing at the outputs of the circuit advantageously reduces power dissipation by a factor of the voltage squared. The

circuit's clock frequency is correspondingly varied with the voltage to maintain critical timing paths within specifications." Ex. 1008, 2:26-30.

76. Georgiou discloses a system in which the operating voltage and clock frequency of individual functional units are dynamically adjusted based on thermal feedback. A critical aspect of this system is the use of pre-characterized, safe voltage-frequency pairs when performing such adjustments. This requirement ensures that any changes in operating conditions do not compromise the timing integrity of the circuit.

77. It is well understood in the field of digital microprocessor design that reducing the supply voltage of a circuit—while beneficial for lowering power consumption and reducing heat—increases the propagation delay through its logic elements. In synchronous systems, where computation is coordinated by a clock signal, this increase in delay must be compensated by reducing the clock frequency, thereby providing more time per cycle for signal transitions to settle properly.

78. However, indiscriminately reducing the clock frequency or voltage can lead to timing violations if the new combination does not satisfy the circuit's critical path requirements. Critical paths are the longest paths through the logic, and they dictate the minimum clock period required for correct operation.

79. To avoid such errors, Georgiou specifies that the system must select a new clock frequency from a set of voltage-frequency pairs that have been pre-characterized—meaning that, for each functional unit, specific voltage and frequency combinations have been validated in advance to ensure that all timing paths meet specification under worst-case conditions, including manufacturing, temperature, and voltage variations.

80. These voltage-frequency pairs are therefore considered safe in the sense that they maintain system synchronization and prevent logic hazards, even as environmental or workload-induced thermal conditions fluctuate.

81. This design constraint is central to the effectiveness and reliability of Georgio's system. It ensures that each functional unit, even when operating at reduced voltage and frequency due to localized heating, will continue to function correctly without violating synchronization boundaries or causing system-wide timing errors.

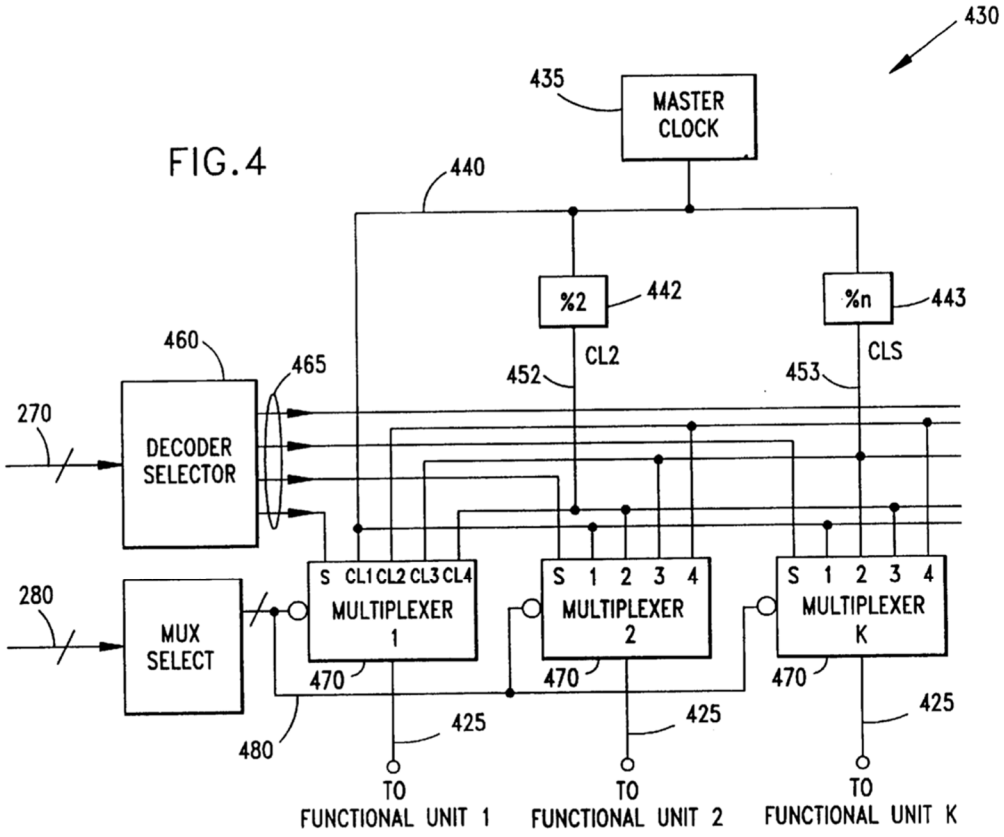
82. In short, the use of pre-characterized, safe voltage-frequency pairs allows the system disclosed in Georgiou to dynamically throttle power and manage heat without compromising system correctness, which is a fundamental requirement in complex, high-performance microprocessors.

83. Georgiou discloses a system that dynamically adjusts the operating voltage and clock frequency of individual functional units within a processor in

response to real-time temperature feedback. This allows the system to manage heat dissipation without compromising timing integrity or synchronization across the chip.

84. One component integral to this voltage-frequency scaling mechanism is a clock divider, which is a well-understood digital circuit element that produces a lower-frequency output signal by dividing the frequency of a higher-frequency input signal. In the context of Georgiou, the clock divider receives a high-frequency master clock signal and generates reduced-frequency clock signals such as f/2, f/4, or f/N, where f is the frequency of the master clock and N is an integer.

85. As illustrated in Figure 4 (reproduced below) of Georgiou and described in the corresponding disclosure (see col. 9, lines 1–15), the master clock (element 435) drives a set of clock dividers (e.g., 442, 443), each configured to output a specific submultiple of the master clock frequency. These outputs are then provided as inputs to a series of multiplexers (element 470), each corresponding to a functional unit within the processor.



86. Georgiou describes a design where all of the clocks in the system are dependent versions of the master clock. Georgiou teaches "A method of using thermal feedback to cooperatively vary a voltage and a frequency applied to a synchronous circuit to control heating while maintaining synchronization." Ex. 1008, 12:16-19. The system is run synchronously, where synchrony is maintained by restricting clock frequencies that are fixed fractions of the master clock (435).

87. The output frequency for each functional unit is selected by a clock select signal (element 465), which is derived from a decoded temperature feedback signal (element 270) received from a local thermal sensor (element 119) and interpreted

by the temperature decoder (element 130). The selection process is guided by a pre-characterized mapping between safe operating voltages and corresponding frequencies that satisfy the worst-case propagation delays for critical timing paths in the unit.

88. The purpose of employing clock dividers in this manner is to enable discrete, hardware-efficient scaling of the clock frequency for a functional unit experiencing thermal stress. By selecting from a set of known, reliable frequency levels derived from the master clock, the system avoids the need for dynamically synthesizing new clocks, thereby reducing complexity and ensuring stability.

89. For example, if the master clock operates at 1 GHz, the system may utilize a divide-by-2 clock divider to output 500 MHz, or a divide-by-4 divider to output 250 MHz. The selected lower-frequency clock is then routed to the overheating functional unit to reduce power dissipation while maintaining synchronization with the rest of the system.

90. In summary, the clock dividers described in Georgiou are essential to implementing the disclosed thermal-feedback-driven performance control system. They provide a practical means for realizing safe, pre-determined reductions in operating frequency in response to real-time temperature conditions, and they do so in a way that preserves system synchronization and avoids the timing hazards that might otherwise arise from arbitrary or asynchronous clock scaling.

91. In addition to the use of clock dividers, Georgiou also discloses the use of multiplexers (element 470 in FIG. 4) to facilitate selectable clock routing for each functional unit based on thermal feedback. The multiplexers are a critical part of the clock selection subsystem, enabling the processor to route the appropriate reduced-frequency clock signal to the overheating unit, thereby implementing the desired voltage-frequency pair dynamically and per-unit.

92. Each multiplexer 470 receives as inputs the outputs of several clock dividers, which are themselves fed by a master clock signal (element 435). The clock dividers generate a range of clock frequencies, such as f, f/2, f/4, ..., f/N, where f is the frequency of the master clock. The multiplexer selects one of these divided clock signals as its output based on a select control signal (element 465).

93. The select control signal is derived from thermal data obtained from a thermal sensor (element 119) associated with a particular functional unit. This sensor produces a temperature-dependent signal (element 125), which is processed by a temperature decoder (element 130) and ultimately used to generate a current/clock control signal (element 270). The decoder 460 (within the clock selector 430) interprets this signal and determines which clock divider output should be selected to match the temperature condition of the corresponding unit.

94. The multiplexer acts as a switching mechanism that routes the appropriate frequency from among the pre-generated options to the clock input of

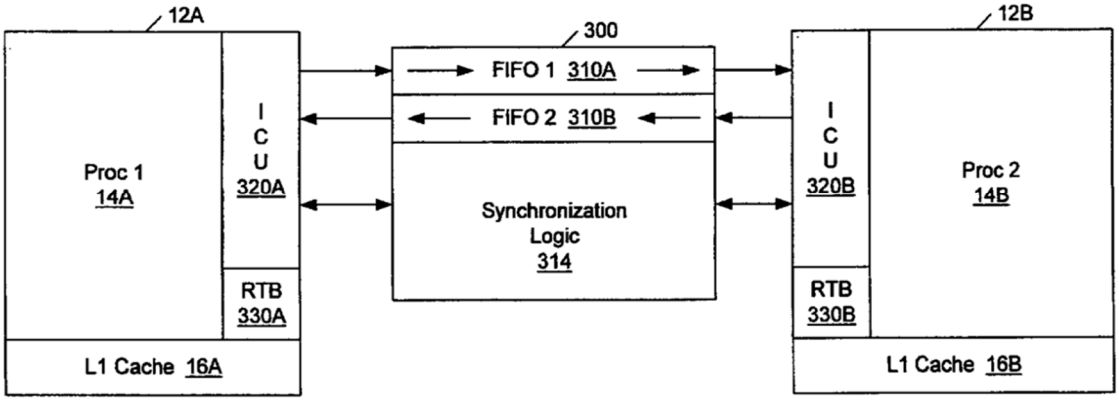
the affected functional unit. This enables the processor to reduce the operating frequency for that unit without affecting the timing of unrelated units, thereby preserving overall system synchronization while addressing local thermal issues.

95. In summary, within the thermal-feedback-based power optimization system disclosed in Georgiou, each multiplexer (element 470) serves as the final stage of the dynamic clock selection path, taking as input a set of divided clock frequencies and outputting a selected clock to the associated functional unit based on the decoded thermal state. Therefore Georgiou's architecture allows for discrete scaling from a fixed set of fractional clock outputs, but does not permit domain clocks to be generated independently or to be dynamically varied beyond the bounds of the preselected small integer ratios.

6. Overview of Kranich

96. Kranich is titled "Fast Multithreading for Closely Coupled Multiprocessors." Kranich describes a "processor architecture containing multiple closely coupled processors in a form of symmetric multiprocessing system." Ex. 1009, Cover, Abstract.

97. The figure below is from the cover of Kranich.



The figure shows synchronization logic (314) between two microprocessors (12A and 12B). The synchronization logic (314) "is included in a form of symmetrical multiprocessing system which enables the scheduling and speculative execution of multiple threads on multiple processors without the involvement and inherent overhead of the operating system." Ex. 1009, Cover, Abstract.

98. Kranich teaches "a thread control device, wherein said thread control device is configured to store and transmit instructions between said processors, wherein said instructions support the speculative execution of threads on said processors, wherein said threads may be executed in parallel." Ex. 1009, Claim 1: 7-9. As shown in the figure above from Kranich, the two microprocessors (12A and 12B) send synchronization messages through a set of First In First Out (FIFO) buffers (310A and 310B). Ex. 1009, 10:17-30.

B. Ground 1: Claims 1, 3-5, 8, 9 and 12 Would Not Have Been Obvious Over Shenai

99. The '443 patent includes two independent claims (1 and 8), as well as ten dependent claims (2-7 and 9-12). The '443 patent teaches "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" Ex. 1001, 17:11-15. Similarly, the '443 teaches "for each of the plurality of domains, a voltage input for receiving a voltage which is dynamically changeable independently of the voltages application to said others of the plurality of domains" Ex. 1001, 17:16-20. As described by Shenai, "These circuits may operate at an integer multiplication, a half-integer multiplication or a quarter-integer multiplication of a reference clock signal." Ex. 1004, 0036. The clock frequency control described in Shenai generates a clock frequency that is based on a rational multiple of a reference clock signal.

100. The '443 patent teaches how to generate independent clocks and voltage sources for a microprocessor, while Shenai is focused on the design of analog to digital converters. Given this difference in application, it is my opinion that a person of ordinary skill in the art would not have looked to Shenai when evaluating the claims of the '443 patent.

101. Based on the '443 patent's claims and the patent's specification, a person of ordinary skill in the art would understand that the claim language in claims 1 and 8 describe a multiple-clock domain microprocessor, where the multiple clocks are independent of each other, and where the input voltages are independent of each other. Further, the frequency of each clock in a domain can be changed independently of other clocks, and the voltage value for each domain can be changed independently of the voltage values of other domains. Both claims 1 and 8 in the '443 patent describe independent clock frequency generation and independent voltage sources.

102. Given that Shenai describes clock frequency generation that is derived through a multiplication/division of a reference clock signal, in my opinion, Petitioner fails to show that any of claims 1, 3-5, 8, 9 and 12 of the '443 patent would have been obvious over Shenai. Further, it is my opinion that Shenai does not disclose or suggest the limitations in the independent claims.

C. Ground 2: Claims 4 and 9 Would Not Have Been Obvious Over Shenai in view of Halepete

103. Claim 4 of the '443 patent is dependent on claim 1 and claim 9 of the '443 patent is dependent on claim 8. Shenai does not teach the independent clocks described in claim 1 and claim 8 of the '443 patent, nor does Halepete, as we will discuss next.

104. In terms of combining Shenai and Halepete, a POSITA would have recognized that these two exhibits have unrelated goals. Shenai is targeting an analog to digital converter, while Halepete is targeting a computer system. It appears to me 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 cite would, in fact, be combined in the first place. I understand that obviousness arguments based on this type of hindsight are improper.

105. Halepete describes clock and voltage control where "The frequencies at which the other components of the system operate are determined from the core frequency determined by multiplying the slow clock by the core processor factor." Ex. 1005, 7:8-10. Given that Claim 1 in the '443 patent is dependent on claim 1 and claim 9 in the '443 patent is dependent on claim 8, and that claims 1 and 8 in the '443 patent teach the ability to independently generate clocks and independently generated voltages, and Halepete does not, it is my opinion that Halepete does not disclose or make obvious claims 4 and 9 of the '443 patent. Further, it is my opinion that the combination of Shenai and Halepete do not disclose or make obvious claims 4 and 9 of the '443 patent.

D. Ground 3: Claims 5-7 and 10-11 Would Not Have Been Obvious Over Shenai in view of Sherburne

106. Claim 5 of the '443 patent is dependent on claim 1, claim 6 of the '443 patent is dependent on claim 5, and claim 7 of the '443 patent is dependent on claim 6. Claim 10 of the '443 patent is dependent upon claim 8 and claim 11 of the '443 patent is dependent on claim 10. Shenai does not teach the independent clocks described in claim 1 and claim 8 of the '443 patent, nor does Sherburne, as we will discuss next.

107. In terms of combining Shenai and Sherburne, a POSITA would have recognized that these two exhibits have unrelated goals. Shenai is targeting an analog to digital converter, while Sherburne is targeting a reconfigurable processing core. Again, it appears to me 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 cite would, in fact, be combined in the first place. I understand that obviousness arguments based on this type of hindsight are improper.

108. Sherburne describes "A low power a reconfigurable processor core includes one or more processing units, each unit having a clock input that controls the performance of the unit; one or more clock controllers having clock outputs coupled to the clock inputs of the processing units, the controller operating varying the clock frequency of each processing unit to optimize speed and processing

power for a task…" Ex. 1006, 1:0008. Further, in Figure 4, Sherburne shows a common clock (CLK) feeding logic that produces the clocks for each processor. Ex. 1006, Figure 4. A POSITA would understand that Sherburne teaches a method that uses a common clock, versus independent clocks as taught in the '443 patent. It is my opinion Shenai, in view of Sherburne, does not disclose or make obvious claims 5-7 and claims 10-11 of the '443 patent.

E. Ground 4: Claims 7 Would Not Have Been Obvious Over Shenai in view of Traylor

109. Claim 7 of the '443 patent is dependent on claim 6, which is dependent on claim 5, which is dependent on claim 1. Shenai does not teach the independent clocks described in claim 1 of the '443 patent, nor does Traylor, as we will discuss next.

110. In terms of combining Shenai and Traylor, a POSITA would have recognized that these two exhibits have unrelated goals. Shenai is targeting an analog to digital converter, while Traylor is targeting First In First Out (FIFO) buffer. Again, it appears to me 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 cite would, in fact, be combined in the first place. I understand that obviousness arguments based on this type of hindsight are improper.

111. Traylor describes methods for avoiding signal glitches that may occur in a First In First Out (FIFO) buffer. Ex. 1007, Cover, Abstract. The methods include a mechanism for tracking the status of the FIFO, indicating when the FIFO is EMPTY or FULL. There would be no reason for a POSITA to look to Traylor when considering the claims of the '443 patent, given that Traylor does not discuss clock generation or voltage source control. In my opinion, Shenai, in view of Traylor, does not disclose or make obvious claim 7 of the '443 patent.

F. Ground 5: Claims 1 and 3 Would Not Have Been Obvious Over Georgiou

112. Claim 1 of the '443 patent is an independent claim and claim 3 of the '443 patent is dependent on claim 1. Georgiou does not teach the independent clock control described in claim 1 of the '443 patent, as we will describe next.

113. Georgiou discusses a system and method "using thermal feedback to cooperatively vary a voltage and frequency of circuit to control heating while maintaining synchronization." Ex. 1008, Cover, Abstract. Georgiou describes how to utilize voltage-frequency pairs to control a circuit based on feedback from a thermal sensor.

114. In my opinion, and based on the plain and ordinary meaning as understood by a person of ordinary skill in the art, the claims of the '443 Patent require that each domain generate its clock signal independently, and that the frequency of each domain's clock be dynamically changeable without being

constrained by the frequencies of other domains. That is, each domain must be free to generate and change its clock frequency without being tethered to small integer multiples of the other domains.

115. Georgiou discloses a clock selector 430, which includes a master clock 435 and multiple clock dividers and multiplexers. The master clock 435 provides a common source signal at frequency 440, which is fed into multiple clock dividers such as divide-by-2 (element 442) and divide-by-n (element 443). Georgiou states that "clock divider 442 may be a divide-by-two clock divider whose output frequency 452 is one-half that of the master clock frequency 440, and clock divider 443 may be divide-by-n clock divider whose output frequency 453 is one-nth that of the master clock frequency 440." Ex. 1008, 8:42–47.

116. These dividers output predetermined fractional frequencies of the master clock, which are then selected via multiplexers 470. As Georgiou explains, "the master clock frequency 440 is used as the common source and provides a common clock reference for a number of clock dividers whose output is some predetermined fraction of the master clock frequency 440." Ex. 1008, 8:39–42. The multiplexer outputs the selected clock to the functional unit. Id. at 8:65–67.

117. In this architecture, each domain's clock is not generated independently. Instead, all domain clocks are constrained to be rational multiples of the master clock frequency—i.e., the output of each domain clock is determined

by applying a fixed divide ratio (e.g., 1/2, 1/n) to a shared global clock. The result is that all domain frequencies are interdependent and maintain a small integer relationship of the form M/N.

118. As a result, Georgiou's architecture necessarily creates a system in which domain clocks are not independently generated and not dynamically changeable in an unconstrained manner. Any changes in clock frequency must occur within a pre-characterized set of divide ratios. The multiplexer 470 cannot generate or select an arbitrary or unrelated frequency—it can only choose among those predetermined by the available dividers.

119. The purpose of employing clock dividers in this manner is to enable discrete, hardware-efficient scaling of the clock frequency for a functional unit experiencing thermal stress. By selecting from a set of known, reliable frequency levels derived from the master clock, the system avoids the need for dynamically synthesizing new clocks, thereby reducing complexity and ensuring stability.

120. For example, if the master clock operates at 1 GHz, the system may utilize a divide-by-2 clock divider to output 500 MHz, or a divide-by-4 divider to output 250 MHz. The selected lower-frequency clock is then routed to the overheating functional unit to reduce power dissipation while maintaining synchronization with the rest of the system. This configuration avoids the

complexity of dynamic frequency generation by selecting from a constrained set of precomputed clock values derived from the master clock

121. Further, the values of N must be relatively small because taking a large value such as 1,000 would result in the 1 GHz master clock outputting a 1 MHz signal, which is not useful in a modern computing system.

122. This reliance on rational frequency relationships is the "coincidence cycle" described in Sarmenta and the "beating period" described in Pratt, where communication between clock domains is synchronized based on a known phase alignment that recurs over a fixed interval. Georgiou's architecture imposes exactly the kind of small-integer ratio constraints that the '443 Patent specification expressly criticizes.

123. Specifically, the '443 Patent distinguishes itself from such prior art schemes by allowing each domain to operate freely, without the need for rational coordination. The specification states: "Some synchronization schemes restrict the ... relative frequencies of the clocks ... [and] impose significant restrictions on the possible choices of frequencies." Ex. 1001, 6:32–36. In contrast, the '443 invention permits each domain to operate at a frequency that is both separately generated and dynamically changeable independently of the other domains.

124. Therefore, in my opinion, Georgiou's use of a shared master clock, fixed divider ratios, and multiplexer-based selection fails to meet the plain and

ordinary meaning of the claim terms "separately generating a clock signal at a frequency for that domain" and "dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains" (claim 1), and "dynamically controlling the clock frequency … independently of the clock frequencies" (claim 12).

125. I have also reviewed Petitioner's assertion that the '443 Patent discloses a "single clock" covering multiple domains (Ex. 1001 at 5:55–62), and I find this reading technically incorrect. The cited passage, when read in context, refers to components within a single domain—not to a global clock governing multiple domains. The same passage makes clear that the preferred embodiment comprises five distinct clock domains: three for processor cores, one for the L2 cache, and one for memory. Ex. 1001, 5:52–62. There is no suggestion that these domains share a common frequency or are related by a rational factor. In contrast, Georgiou's system imposes such a rational dependency by design.

126. In conclusion, Georgiou does not disclose or suggest a system in which clock frequencies are generated independently and are dynamically controllable without reference to other domains. Instead, it teaches a globally synchronized architecture based on predetermined, rationally related frequency relationships. This structure is fundamentally incompatible with the claim

limitations of the '443 Patent and with the technical objectives expressed throughout its specification.

G. Ground 6: Claims 4 and 9 Would Not Have Been Obvious Over Georgiou in view of Halepete

127. Claim 4 of the '443 patent is dependent on claim 1 and claim 9 of the '443 patent is dependent on claim 8. Georgiou in view of Halpete does not teach the independent clocks described in claims 1 and 8 of the '443 patent, as we will describe next.

128. In terms of combining Georgiou and Halpete, a POSITA would have recognized that these two exhibits have unrelated goals. Georgiou is trying to avoid a thermal emergency on a system, while Halepete is trying to extend battery life of a computer. Again, it appears to me 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 cite would, in fact, be combined in the first place. I understand that obviousness arguments based on this type of hindsight are improper.

129. Halepete teaches how to "increase significantly the operating life of a portable computer" Ex. 1005,3:18-19. Halepete describes clock and voltage control where "The frequencies at which the other components of the system operate are determined from the core frequency determined by multiplying the slow clock by the core processor factor." Ex. 1005, 7:8-10. Given that Claim 4 in

the ‘443 patent is dependent on claim 1 in the ‘443 patent and claim 9 in the ‘443 patent is dependent on claim 8 in the ‘443 patent, and that claims 1 and 8 teach the ability to independently generate clocks and independently general voltages, and Halepete does not, it is my opinion that Halepete does not disclose or make obvious claims 4 and 9 in the ‘443 patent.

130. As discussed earlier, in Figure 4 from Georgiou, a master clock (435) is divided down (442 and 443) into slower versions of the master clock (435). These multiples of the master clock (442 and 443) are then fed to the multiplexers 1-K (470), which supply functional units 1-K via (425). A POSITA would understand the clocking structure taught by Georgiou would lack the flexibility and precision of the independent clocks that are taught in the ‘443 patent.

131. Given that neither Halepete nor Georgiou teaches the independent clocking taught in the ‘443 patent, the combination of Halepete and Georgiou do not make up for this deficiency. For at least the foregoing reasons, it is my opinion that the Petitioner has not shown that claim 4 or 9 would have been obvious.

H. Ground 7: Claims 5 and 10 Would Not Have Been Obvious Over Georgiou in view of Kranich

132. Claim 5 of the ‘443 patent is a dependent claim on claim 8 and claim 10 of the ‘443 patent is dependent on claim 1. Georgiou in view of Kranich does not teach the independent clocks described in claims 1 and 8 of the ‘443 patent, as we will describe next.

133. Kranich is titled "Fast Multithreading for Closely Coupled Multiprocessors." Kranich describes a "processor architecture containing multiple closely coupled processors in a form of symmetric multiprocessing system" Ex. 1009, Cover, Abstract.

134. In terms of combining Georgiou and Kranich, a POSITA would have recognized that these two exhibits have unrelated goals. Georgiou is trying to avoid a thermal emergency on a single system, while Kranich is providing thread synchronization between two systems. Again, it appears to me 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 cite would, in fact, be combined in the first place. I understand that obviousness arguments based on this type of hindsight are improper.

135. Kranich teaches "a thread control device, wherein said thread control device is configured to store and transmit instructions between said processors, wherein said instructions support the speculative execution of threads on said processors, wherein said threads may be executed in parallel." Ex. 1009, Claim 1: 7-9. As shown the cover figure from Kranich, the two microprocessors (12A and 12B) send synchronization messages through a set of First In First Out (FIFO) buffers (310A and 310B). Ex. 1009, 10:17-30. Kranich does not mention clocking or voltage sources.

136. As discussed earlier, in Figure 4 from Georgiou a master clock (435) is divided down (442 and 443) into slower versions of the master clock (435). These multiples of the master clock (442 and 443) are then fed to the multiplexers 1-K (470), which supply functional units 1-K via (425). Georgiou's clocking structure is dependent on a master clock, in contrast to the independent clocks taught by the '443 patent. A POSITA would understand the clocking structure taught by Georgiou would lack the flexibility and precision of the independent clocks that are taught in the '443 patent.

137. Given that neither Kranich nor Georgiou teaches the independent clocking taught in the '443 patent, the combination of Kranich and Georgiou does not make up for this deficiency. For at least the foregoing reasons, it is my opinion that the Petitioner has not shown that claim 5 or 10 would have been obvious.

VIII. Conclusion

138. For the foregoing reasons, it is my opinion that claims 1, 3-12 of the '443 patent would not have been obvious over the prior art in Petitioner's Petition for *inter partes* review.

All statements made herein of my own knowledge are true, and all statements made on information and belief are believed to be true. Further, I am aware that these statements are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001. I declare under penalty of perjury that the foregoing is true and correct.

Executed on May 9, 2025, in Boston, Massachusetts.



Dr. David Kaeli