UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

LENOVO (UNITED STATES) INC.
Petitioner

v.

UNIVERSITY OF ROCHESTER
Patent Owner

Inter Partes Review No.
IPR2024-01225

Patent No. 7,089,443 B2
Filing Date: January 23, 2004
Issue Date: August 8, 2006

PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 7,089,443 B2

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Cases

Statutes

Other Authorities

LISTING OF EXHIBITS

Exhibit	Description
Exhibit 1001	U.S. Patent No. 7,089,443 B2
Exhibit 1002	Declaration of R. Jacob Baker, P.E., Ph.D.
Exhibit 1003	Prosecution History of U.S. Patent No. 7,089,443 B2
Exhibit 1004	U.S. Patent Publication No. 2003/0030326 A1 to Shenai *et al.* ("*Shenai")* published on February 13, 2003
Exhibit 1005	International Publication No. WO 01/53921 A1 to Halepete *et al.* ("*Halepete*") published on July 26, 2001
Exhibit 1006	U.S. Patent Publication No. 2002/0169990 A1 to Sherburne ("*Sherburne*") published on November 14, 2002
Exhibit 1007	U.S. Patent No. 5,473,756 to Traylor ("*Traylor*") issued on December 5, 1995
Exhibit 1008	U.S. Patent No. 6,047,248 to Georgiou *et al.* (*"Georgiou*") issued on April 4, 2000
Exhibit 1009	International Publication No. WO 01/33352 A1 to Kranich *et al.* ("*Kranich*") published on May 10, 2001
Exhibit 1010	Webpage for Datamath Calculator Museum - archived by WayBackMachine on February 26, 2001 ("*Datamath*")
Exhibit 1011	Stanley Mazor, *The History of the Microcomputer – Invention and Evolution* (1995) ("*Mazor*")
Exhibit 1012	William M. Johnson, *Super-Scalar Processor Design*, (1989) ("*Johnson*")
Exhibit 1013	Altera, *Excalibur Arm-Based Hardware Reference Manual*, January 2001 ("*Excalibur*")
Exhibit 1014	Texas Instruments, *TMS320C64x Technical Overview*, January 2001 ("*TMS*")

Exhibit 1015	Kenneth C. Yeager, *The MIPS R100 Superscalar Microprocessor,* (1996) ("*Yeager*")
Exhibit 1016	Texas Instruments, *FIFO Architecture, Functions, and Applications*, November 1999 ("FIFO")
Exhibit 1017	Eric Tune *et al.*, *Dynamic Prediction of Critical Path Instructions,* (2001) ("*Tune*")
Exhibit 1018	U.S. Patent No. 6,289,465 B1 to Kuermerle ("*Kuermerle*") issued on September 11, 2001
Exhibit 1019	Anand Lal Shimpi, *AMD's K6-2+ in Notebooks – Evaluating Performance with PowerNow*, September 12, 2000 (*"POWERNOW!"*)
Exhibit 1020	Anand Lal Shimpi, *AMD's K6-2+ in Notebooks – Evaluating Performance with PowerNow*, September 12, 2000 ("*POWERNOW! Operating Modes*")
Exhibit 1021	Bruce R. Childers *et al.*, *Adapting Processor Supply Voltage to Instruction-Level Parallelism*, (2001) ("*Childers*")
Exhibit 1022	Andrew Royal *et al.*, *GALS, Globally Asynchronous, Locally Synchronous Systems*, April 27th, 2001 ("*GALS*")
Exhibit 1023	U.S. Patent No. 6,141,762 to Nicol *et al.* issued on October 31, 2000 ("*Nicol*")
Exhibit 1024	*Reserved*
Exhibit 1025	*Reserved*
Exhibit 1026	Actel Corporation, *3200DX Dual-Port Random Access Memory (RAM)*, September 1997 ("*3200DX*")
Exhibit 1027	Barr Group, *Types of Memory in Embedded Systems*, May 1, 2001
Exhibit 1028	Wavecrest corporation, *Understanding Jitter*, 2001

Exhibit 1029	U.S. Application No. 60/441,759 filed January 23, 2003
Exhibit 1030	U.S. Patent No. 7,739,537 to Albonesi *et al.* ("*Albonesi"*) issued on June 15, 2010
Exhibit 1031	Declaration of Ankit Aggarwal

I. INTRODUCTION

Lenovo (United States), Inc. ("Petitioner") requests *inter partes* review ("IPR") of claims 1 and 3-12 ("Challenged Claims"[1]) of USP 7,089,443 ("'443 Patent").

II. MANDATORY NOTICES

A. Real Party-in-Interest

Petitioner hereby names Lenovo (United States) Inc. as a real-party-in-interest and, solely because it is named as a defendant in the co-pending district court case listed below, further identifies Lenovo Group Ltd. as a real party-in-interest.

B. Related Matters

Intellectual Ventures I LLC *et al.* has asserted the '443 Patent against Lenovo Group Limited in *Intellectual Ventures I LLC et al. v. Lenovo Group Limited*, 6:23-cv-307 (WDTX) ("co-pending litigation").

The '443 Patent is also asserted in the following cases:

- *Intellectual Ventures I LLC v. TCL Electronics Holdings Ltd. et al.,* No. 6-23-cv-00309 (WDTX, filed April 26, 2023)

[1] A subset of the Challenged Claims are being asserted (and at issue) in the co-pending litigation.

- *Intellectual Ventures I LLC et al. v. TCL Electronics Holdings Ltd. f/k/a TCL Multimedia Technology Holdings, Ltd. et al.,* No. 6-23-cv-00293 (WDTX, filed April 20, 2023);
- *Intellectual Ventures II LLC v. Lenovo Group Limited*, No. 6-23-cv-00068 (WDTX, filed February 2, 2023, terminated on April 26, 2023).

C. Counsel and Service Information

Lead counsel is Dinesh Melwani (Reg. No. 60,670). Back-up counsel are Ankit Aggarwal (Reg. No. 67,882) and William Uhr (Reg. No. 71,282). Service information is: Bookoff McAndrews, PLLC, 2000 Pennsylvania Avenue NW, Suite 4001, Washington, D.C. 20006; Tel.: 202.808.3497; Fax.: 202.450.5538; email: docketing@bomcip.com, dmelwani@bomcip.com, aaggarwal@bomcip.com, and wuhr@bomcip.com. Petitioner consents to electronic service.

III. PAYMENT OF FEES

The PTO is authorized to charge any fees due during this proceeding to Deposit Account No. 50-5906.

IV. GROUNDS FOR STANDING

Petitioner certifies that the ’443 Patent is available for review and Petitioner is not barred/estopped from requesting review on these grounds. 37 C.F.R. § 42.104(a).

V. PRECISE RELIEF REQUESTED AND GROUNDS

A. Identification of Challenge[2]

Petitioner requests the Challenged Claims to be found unpatentable under the following grounds:

Ground	35 U.S.C. §	Claims	References
1	103	1, 3–5, 8-9, 12	US-PUB 2003/0030326 ("*Shenai*")
2	103	4, 9	*Shenai* in view of PCTPUB WO01/53921 ("*Halepete*")
3	103	5–7, 10, 11	*Shenai* in view of US-PUB 2002/0169990 ("*Sherburne*")
4	103	7	*Shenai* in view of *Sherburne* and USP 5,473,756 ("*Traylor*")
5	103	1, 3, 8, 12	USP 6,047,248 ("*Georgiou*")
6	103	4, 9	*Georgiou* in view of *Halepete*

[2] For each Ground, Petitioner does not rely on any reference other than those listed here. Other references are discussed to show the state of the art at the time of invention. *See Ariosa Diagnostics v. Verinata Health, Inc.*, 805 F.3d 1359, 1365 (Fed. Cir. 2015).

7	103	5, 10	*Georgiou* in view of PCTPUB WO01/33352 ("*Kranich*")

U.S. Application 10/762,550 ("'550 Application") that matured into the '443 Patent was filed 1/23/2004 and purportedly claims priority from U.S. Provisional Application 60/441,759 ("'759 Application") filed 1/23/2003. Ex.-1001, p.1; Ex.-1029. For purposes of this proceeding, Petitioner assumes the priority date of the '443 Patent is 1/23/2003.

VI. LEVEL OF ORDINARY SKILL

As of 1/23/2003, a person having ordinary skill in the art ("POSITA") would have had a bachelor's degree in electrical engineering, computer engineering, or the equivalent, and 2–3 years of experience designing computing circuits and systems. More practical experience could offset the aforementioned education and vice-versa. Ex.-1002, ¶¶29-31.

VII. SUMMARY OF THE '443 PATENT

A. The '443 Patent

The '443 patent describes systems and methods for microprocessors with multiple clock domains. Ex.-1001, 1:19-21; Ex.-1002, ¶57. Specifically, the '443 patent describes a multiple clock domain ("MCD") microarchitecture using a globally-asynchronous, locally-synchronous clocking style. Ex.-1001, p.1, Abstract;

Ex.-1002, ¶57. The ’443 patent asserts that the techniques described are useful in increasing microprocessor performance by providing an approach that “allows for aggressive future frequency increases [and] maintains a synchronous design methodology[.]” Ex.-1001, 4:9-12; Ex.-1002, ¶57. Figure 1 below illustrates an MCD processor block diagram.

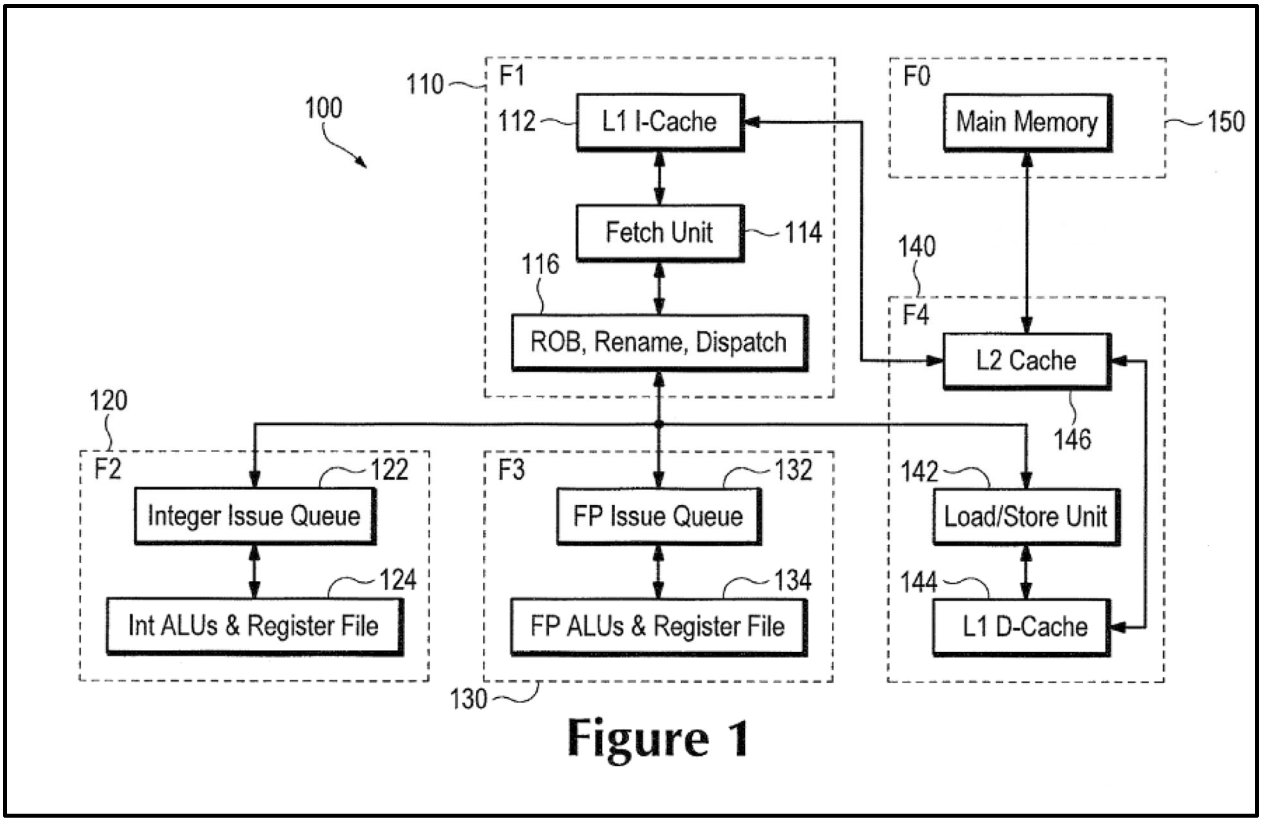

Figure 1 of USP 7,739,537[3]

[3] Figure-1 of the ’443 patent is illegible. Figure-1 above is from USP 7,739,537, which claims priority to ’550 application that matured into the ’443 patent and includes the same content as Figure-1 of the ’443 patent.

Ex. 1030, p.3; Ex.-1002, ¶57.

Figure 1 includes four domains: front end 110, integer issue/execute 120 (including integer issue queue 122), floating point issue/execute 130 (including floating point issue queue 132), and load/store issue/execute 140 (including load/store queue within load/store unit 142). Ex. 1030, 6:1-8, 6:46; Ex.-1002, ¶¶58-60. The '443 patent discloses "scaling frequency and voltage in different domains dynamically and independently." Ex.-1001, 5:8-10; Ex.-1002, ¶¶61, 66-67.

The '443 patent states that the main disadvantage of an MCD processor is "performance overhead due to inter-domain synchronization." Ex.-1001, 6:27-28; Ex.-1002, ¶58. Queues in the domains of Figure 1 are chosen as inter-domain synchronization points to help hide the synchronization cost "whenever the queue is neither full nor empty." Ex.-1001, 6:45-53; Ex.-1002, ¶¶58-60.

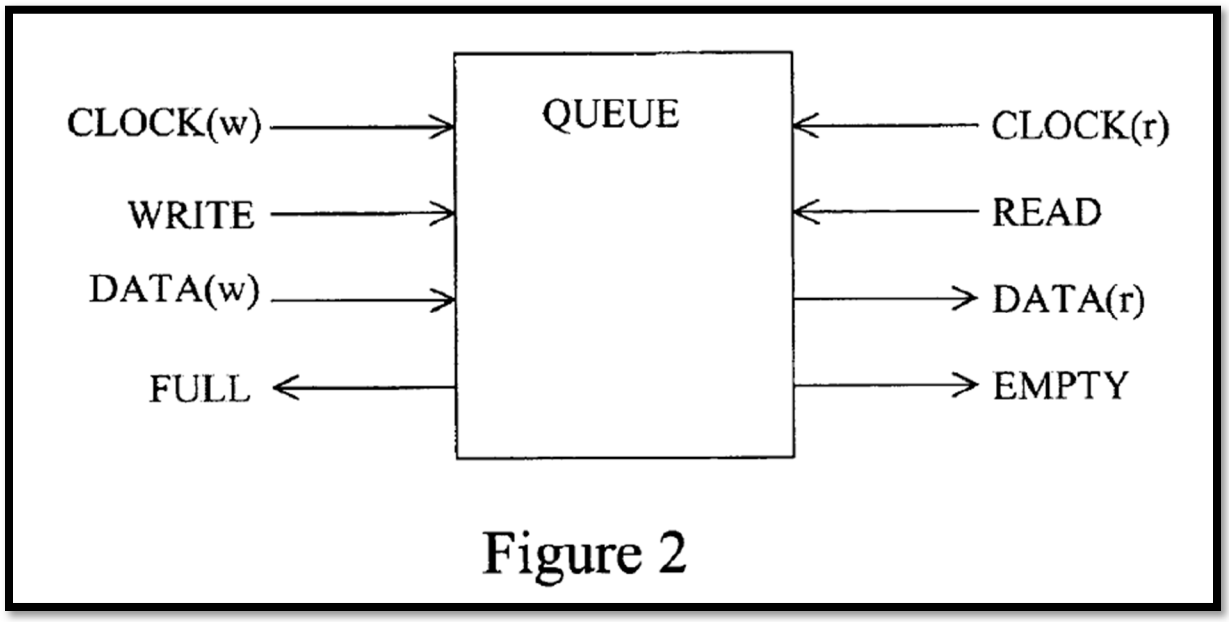

Figure 2 of '443 Patent

Ex.-1001, p.2; Ex.-1002, ¶58. Figure 2 above shows a general queue structure used for inter-domain communication. Ex.-1001, 6:54-55; Ex.-1002, ¶58. "The assertion of the Full flag indicates to the producer that it can no longer write to the queue until the flag is deasserted ($\overline{\text{Full}}$), while the Empty flag when asserted indicates that there is no valid data for the consumer to read from the queue. The consumer waits until Empty is deasserted before reading again." Ex.-1001, 6:55-60; Ex.-1002, ¶¶58-59.

A purported advantage of the MCD design is it simplifies the global clock distribution network, "requiring only the distribution of the externally generated clock to the local Phase Lock Loop (PLL) in each domain." Ex.-1001, 4:25-28; Ex.-1002, ¶¶61-63.

A microprocessor's execution flow may include an event with multiple incoming dependencies with different overall speeds of execution. Ex.-1001, 11:25-32; Ex.-1002, ¶¶64-65. Thus, one arc, the path having the slowest execution speed, becomes the critical path while the other arcs, which execute faster than the critical path, include slack indicating "the previous operation completed earlier than necessary." Ex.-1001, 11:25-32; Ex.-1002, ¶¶64-65.

B. The '443 Patent's File History

The '443 patent was filed as the '550 Application on 1/23/2004 and claims priority to the '759 Application, filed 1/23/2003. The '550 Application was filed with three claims. Ex.-1003, p.98; Ex.-1002, ¶¶68-69. A preliminary amendment canceled the original three claims and added 12 new claims. Ex.-1003, pp.39-42; Ex.-1002, ¶¶68-69.

On 05/19/2005, all pending claims were allowed in a Notice of Allowance. Ex.-1003, pp.24-29; Ex.-1002, ¶¶68-69. The Notice stated the "cited prior art describe systems that change the voltage and frequency of components in data processing systems [but] … does not describe the domains in a microprocessor and does not describe dynamically controlling the clock frequency and voltage in each domain independently of the clock frequencies and voltages in each of the other domains." Ex.-1003, p.28; Ex.-1002, ¶¶68-69. Additionally, the prior art "does not

describe a plurality of domains in a microprocessor or using their system in a microprocessor having a plurality of domains." Ex.-1003, p.28; Ex.-1002, ¶¶68-69.

VIII. SUMMARY OF THE PRIOR ART

1. *Shenai*

Shenai was filed 08/12/2002, published 02/13/2003, and is prior art at least under Section 102(e). Ex.-1002, ¶70.

Shenai discloses a "power distribution management apparatus for supplying power to two or more loads [that] includes a power and clock distribution controller capable of determining voltage, current and clock signal frequency targets for the loads." Ex.-1004, p.1, Abstract; Ex.-1002, ¶71. The power distribution management apparatus includes a plurality of power sources and a plurality of clock signal sources coupled to respective loads to selectively provide target voltages and frequencies to each load. *Id*.

In *Shenai,* circuit 400 balances power consumption and performance for a plurality of loads. Ex.-1004, [0025]; Ex.-1002, ¶71. The loads can be "components or subsystems of a microprocessor or microcontroller such as cache memory, register files, arithmetic logic units (ALU), integer units, instruction pipeline circuitry, hardware multipliers, floating point circuits, non-volatile memory units, or circuitry clusters[.]" Ex.-1004, [0025]; Ex.-1002, ¶71. This can be seen in Figure 4 below:

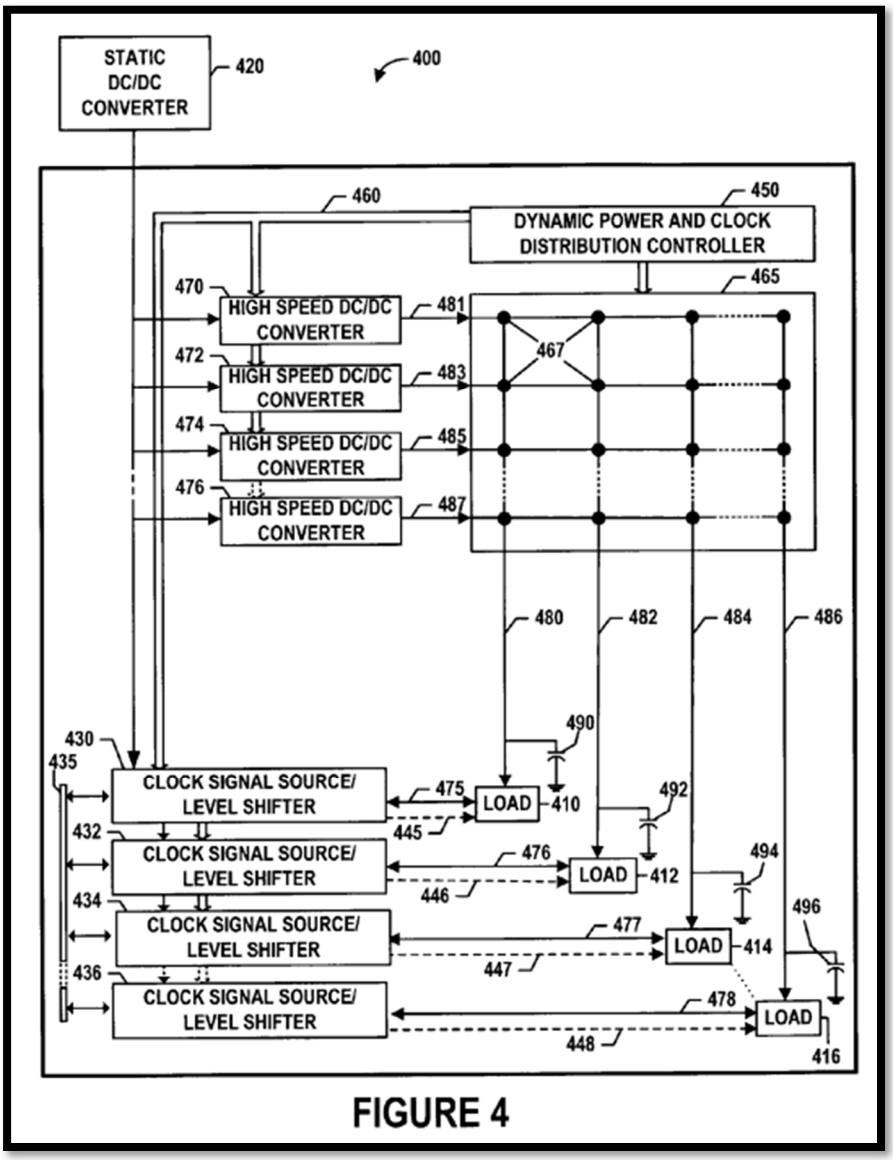

FIGURE 4

Figure 4 of *Shenai*

Ex.-1004, p.5; Ex.-1002, ¶¶71-72.

Controller 450 uses clock/shifters 430, 432, 434, 436 to generate target frequencies to respective loads, and DC/DC converters 470, 472, 474, 476 to generate target voltages to respective loads. Ex.-1004, [0035] and [0037]; Ex.-1002, ¶73. The clock/shifters are "coupled with loads 410-416 so as to provide clock

signals via clock signal buses 445, 446, 447 and 448 to the loads at the target frequencies, which may be determined by controller 450[.]" Ex.-1004, [0035]; Ex.-1002, ¶73. Similarly, controller 450 selectively couples "DC/DC converters 470-476 with loads 410-416 via connection network 465's nodes 467 and voltage supply lines [480-487]." Ex.-1004, [0037]; Ex.-1002, ¶73.

As "certain processing activities of software running on a microprocessor may utilize various subsystems more heavily," controller 450 allows each load to receive variable voltage and frequency values depending on the needs of the specific load. Ex.-1004, [0030]; Ex.-1002, ¶¶74-75.

2. *Halepete*

Halepete published 07/26/2001 and is prior art at least under Section 102(b). Ex.-1002, ¶76.

Halepete describes a system for determining the necessary operating frequency and voltage for a processor and dynamically changing the current operating frequency and voltage to match the necessary values determined. Ex.-1005, p.1, Abstract; Ex.-1002, ¶77. Figure 1 of *Halepete* shows processor 10, clock generator 11, programmable voltage generator 12, and other components. Ex.-1005, 4:21-23; Ex.-1002, ¶78.

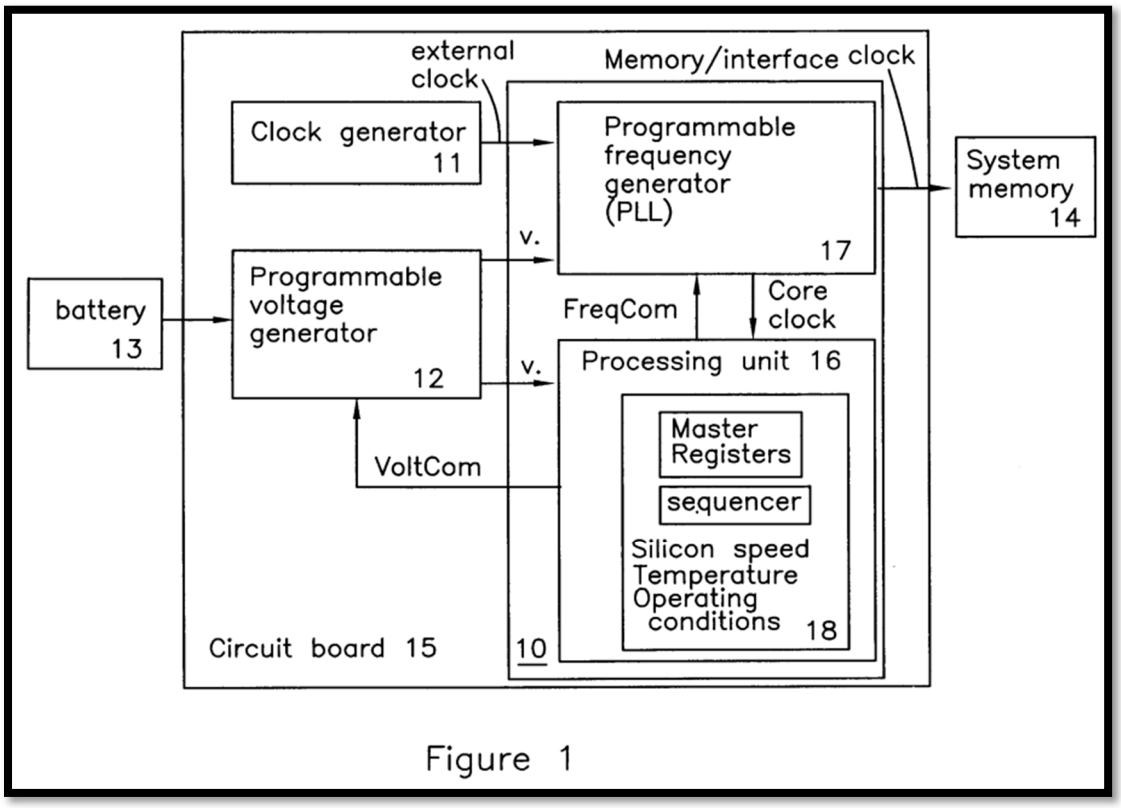

Figure 1 of *Halepete*

Ex.-1005, p.24; Ex.-1002, ¶78. As shown, clock generator 11 supplies an external clock to processor 10, and to phase-locked loop (PLL) 17. Ex.-1002, ¶78.

PLL 17 "receives an external frequency often referred to as a 'slow clock' from the external clock generator 11." Ex.-1005, 6:3-5. *Halepete* states that processor 10 "typically includes a number of other components which are known to those skilled in the art but are not pertinent to the present invention and are therefore not illustrated." *Id.*, 5:7-10; Ex.-1002, ¶78. "Processing unit 16 includes a number

of logical components including a master control unit 18 which is the central portion for accomplishing clock and voltage control." Ex.-1005, 5:10-12; Ex.-1002, ¶79.

For more adaptability, separate clocks are provided for each component within processor 10. Ex.-1005, 6:5-10; Ex.-1002, ¶79. *Halepete* states that the frequency generator can "provide individual frequencies selectable for each of these components." Ex.-1005, 6:11-13; Ex.-1002, ¶79. Control software "detects operating characteristics and determines whether those characteristics indicate that the frequency and the voltage of operation should be changed." Ex.-1002, 11:14-16; Ex.-1002, ¶79.

3. *Sherburne*

Sherburne was filed 03/21/2001, published 11/14/2002, and is prior art at least under Sections 102(a) and (e). Ex.-1002, ¶80.

Sherburne discloses "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 [by] varying the clock frequency of each processing unit to optimize speed and processing power for a task; and a high-density memory array core coupled to the processing units." Ex.-1006, p.1, Abstract; Ex.-1002, ¶81. *Sherburne* teaches a FIFO buffer that allows processing units to communicate, as shown in Figure 3 below:

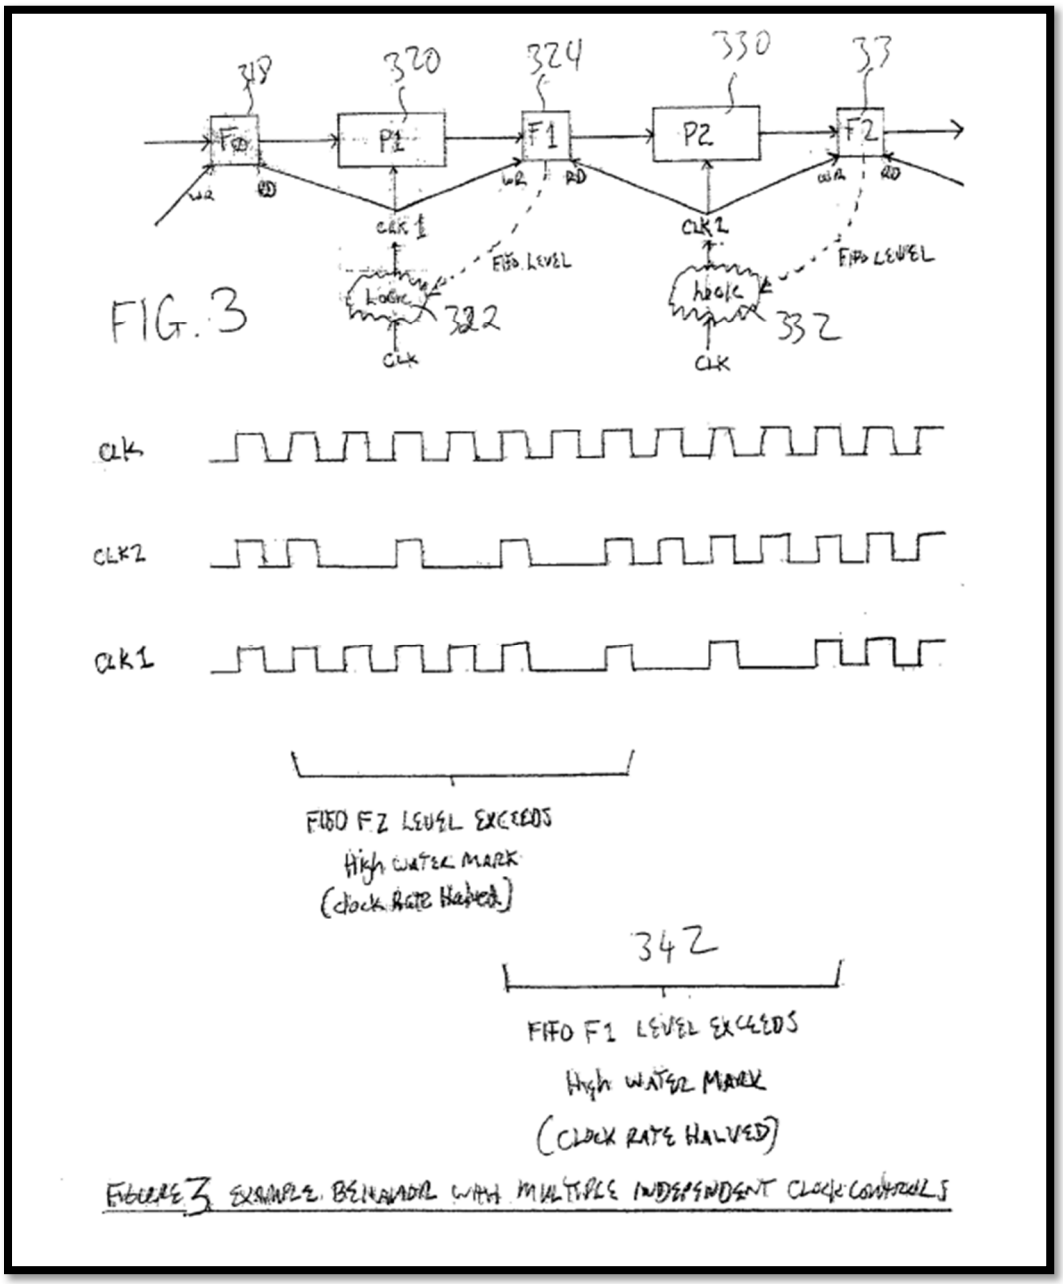

Figure 3 of *Sherburne*

Ex.-1006, p.4; Ex.-1002, ¶81. A data processing unit, such as processor 330, receives the output of a FIFO buffer, such as buffer 324. Ex.-1006, [0029]; Ex.-1002, ¶82.

Processor 330 then generates data and provides the data to buffer 334, which provides the data to the next processor. Ex.-1006, [0029]; Ex.-1002, ¶82. Buffers can generate feedback signals such as "BUFFER_FULL, BUFFER_EMPTY[.]" Ex.-1006, [0029]; Ex.-1002, ¶82. Controller 332, in response, varies the clock signal to the processor either reading or writing on the buffer to adjust the fill rate. Ex.-1006, [0029]; Ex.-1002, ¶82*; see also* Ex.-1006, [0013]. The clock rate is reduced in response to the signal provided. For example, the BUFFER_FULL signal tells the system to prevent any further writes and the BUFFER_EMPTY signal tells the system to prevent any reads. Ex.-1006, [0029]; Ex.-1002, ¶¶82-83.

4. *Traylor*

Traylor published 12/05/1995 and is prior art at least under Section 102(b). Ex.-1002, ¶84.

Traylor teaches "generating control signals for a high speed First In First Out (FIFO) buffer" that "limits the instances where signal glitches may occur." Ex.-1007, p.1, Abstract; Ex.-1002, ¶85. FIFO buffer logic allows for a flag that generates FULL and EMPTY control signals. Ex.-1007, 5:38-42; Ex.-1002, ¶¶85-86. *Traylor* also teaches that "[g]enerally, a FIFO is comprised of a memory device, e.g. a Static Random Access Memory (SRAM), and FIFO controller." Ex.-1007, 1:19-21; Ex-1002, ¶188.

"EMPTY and FULL signals are used to avoid reading from an empty FIFO or writing to a full FIFO." Ex.-1007, 5:44-46; Ex.-1002, ¶86.

5. *Georgiou*

Georgiou published 04/04/2000 and is prior art at least under Section 102(b). Ex.-1002, ¶87.

Georgiou describes a system and method for a scalar microprocessor (e.g., a multi-functional unit processor) using "thermal feedback to cooperatively vary a voltage and frequency of a circuit to control heating while maintaining synchronization." Ex.-1008, p.1, Abstract, 3:50-52; Ex.-1002, ¶88-89.

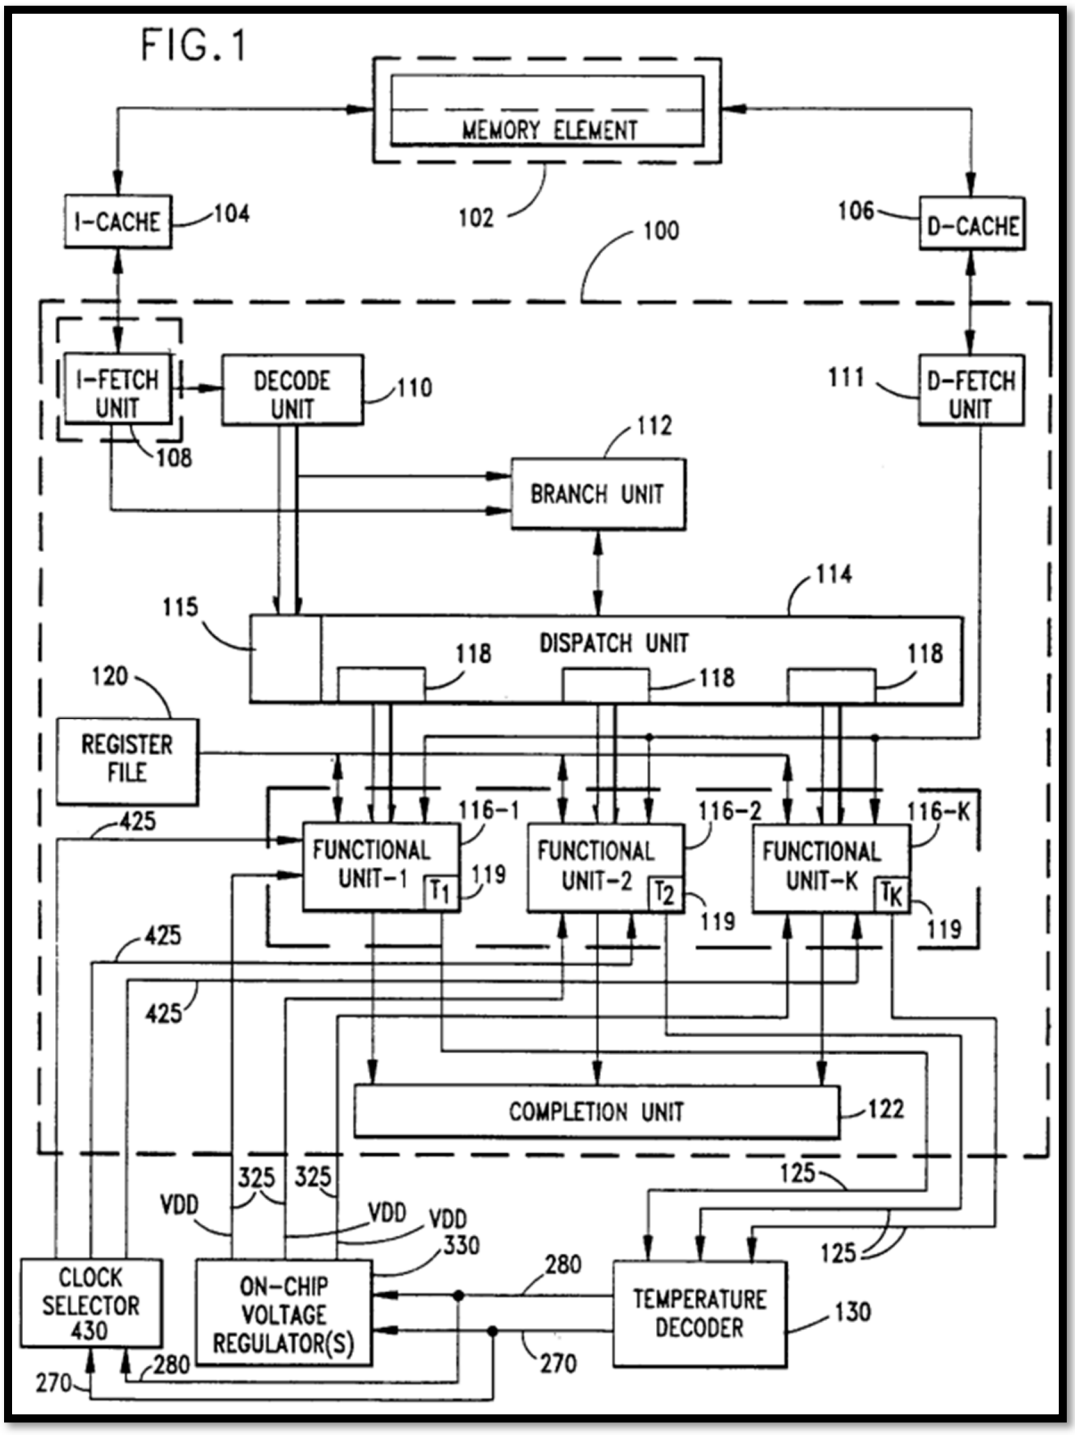

Figure 1 of *Georgiou*

Ex.-1008, p.3; Ex.-1002, ¶88.

Processing unit 100 includes functional units 116-1, 116-2…116-K, dispatch unit 114, branch unit 112, and other components. Ex.-1008, p.3; Ex.-1002, ¶¶88-89. The functional units may include "fixed-point units, floating point units, load/store units, and branch units found in typical general purpose and superscalar microprocessors." Ex.-1008, 3:59-61; Ex.-1002, ¶¶88-89; *see also* Ex.-1008, 4:52-54 ("As is typical, the functional units may include a floating point unit, multiple load/store and fixed point units, and a branch execution unit.").

Clock selector 430 provides a clock frequency 425 (e.g., a clock) to each functional unit via multiplexer 470. Ex.-1008, p. 3; Ex.-1008, 4:27-51; Ex.-1002, ¶¶90-91. Each clock and voltage provided is dynamically changeable independently of the clocks and voltages of the other functional units. Ex.-1002, ¶¶90-91; *see* Ex.-1008, 5:5-6 ("[E]ach functional unit may be operating under a different voltage 325 and frequency 425[.]").

Temperature decoder 130 generates control signal 270 to compare values related to temperature change to a threshold value. Ex.-1008, 4:28-32; Ex.-1002, ¶¶90-91. Signal 270 is then supplied to voltage regulators 330 to, if necessary, reduce the supply voltage 325. Ex.-1008, 4:34-36; Ex.-1002, ¶¶90-91. Signal 270 is provided to clock selector 430 to "modulate the clock frequency sufficiently to maintain system synchronization at the reduced supply voltage." Ex.-1008, 4:36-39; Ex.-1002, ¶¶90-91. This clock frequency 425 is applied to a clock input of the

associated functional unit. Ex.-1008, 4:42-44. Clock selector 430 "may be implemented on-chip and associated with each independently controllable functional unit[.]" Ex.-1008, 5:2-4; Ex.-1002, ¶¶90-91.

6. *Kranich*

Kranich published 05/10/2001 and is prior art at least under Section 102(b). Ex.-1002, ¶92.

Kranich teaches a "processor architecture containing multiple closely coupled processors in a form of symmetric multiprocessing system" utilizing FIFO buffers for interconnection between processors. Ex.-1009, p.1, Abstract; Ex.-1002, ¶93. Figure 3A illustrates the relationship between the processors and buffers:

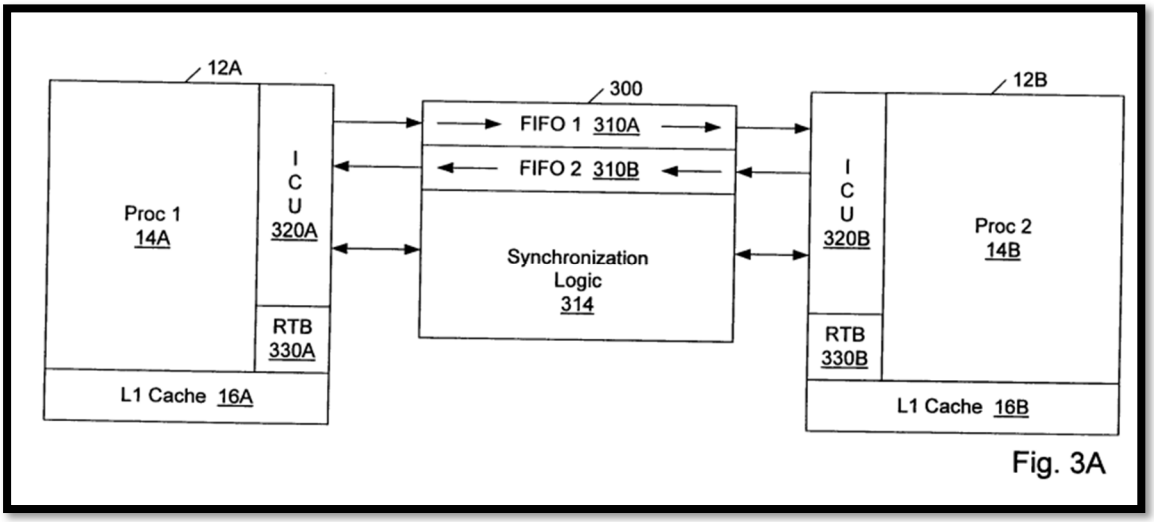

Figure 3A of *Kranich*

Ex.-1009*,* p.26; Ex.-1002, ¶93. Processors 12A-12B communicate with one another through FIFO 310A and 310B. Ex.-1002, ¶¶93-94; Ex.-1009, 10:8-10 ("Each processor 12A-12B includes a processing core 14A-14B, an L1 cache 16A-16B, and an interprocessor communication unit (hereinafter ICU) 320A-320B."). ICUs 320A and 320B are coupled to thread control device 300 "which facilitates communication between processors 12A and 12B" and includes FIFO buffers 310A-310B. Ex.-1002, ¶¶93-94; Ex.-1009, 10:12-13.

IX. CLAIM CONSTRUCTION

The claims of the '443 Patent should be construed under the *Phillips* standard. 37 C.F.R. § 42.100(b); *see generally Phillips v. AWH Corp*., 415 F.3d 1303 (Fed. Cir. 2005). Under *Phillips*, claim terms are typically given their ordinary and customary meanings, as would have been understood by a POSITA, at the time of the invention, having taken into consideration the language of the claims, specification, and prosecution history. *Id.* at 1313; *see also id.* at 1312-16. The Board, however, only construes the claims when necessary to resolve the underlying controversy. *Toyota Motor Corp. v. Cellport Sys., Inc*., IPR2015-00633, Paper 11 at 16 (P.T.A.B. Aug. 14, 2015). Petitioner believes that no construction of the claims

is necessary to assess whether the prior art reads on the Challenged Claims.[4] Ex.-1002, ¶32.

X. THE CHALLENGED CLAIMS ARE UNPATENTABLE

A. Ground 1: Claims 1, 3–5, 8, 9, and 12 are obvious in view of *Shenai*

1. *Claim 1*

a. (Element 1(a)) A multiple clock domain microprocessor comprising:

If the preamble is limiting, *Shenai* teaches *a multiple clock domain microprocessor*. Ex.-1002, ¶¶104-107. *Shenai* is directed to the architecture of a microprocessor where "certain processing activities of software running on a microprocessor may utilize various subsystems more heavily." Ex.-1004, [0030]; Ex.-1002, ¶¶104, 107. *Shenai* describes circuit 400 (the claimed "multi clock domain microprocessor") including loads 410-416, which may be "individual components, functional blocks of a single component[,] or independent systems." Ex.-1004, [0025]; Ex.-1002, ¶¶104-107.

Figure 4 of *Shenai* (shown below) illustrates "plural loads 410, 412, 414, and 416[.]" Ex.-1004, [0025], p.5. *Shenai*'s loads are "components or subsystems of a

[4] Petitioner reserves all rights to raise claim construction and other arguments in other proceedings.

microprocessor … [such as] arithmetic logic units (ALU), integer units," "non-volatile memory units," a "floating point unit," or even an "integer unit." *Id*., [0025], [0028]; Ex.-1002, ¶¶104-107. Similarly, the '443 Patent describes an "integer issue/execute" domain, a "floating point issue/execute" domain, and a "load/store issue/execute" domain. Ex.-1001, 6:3-8; Ex.-1002, ¶¶104-107.

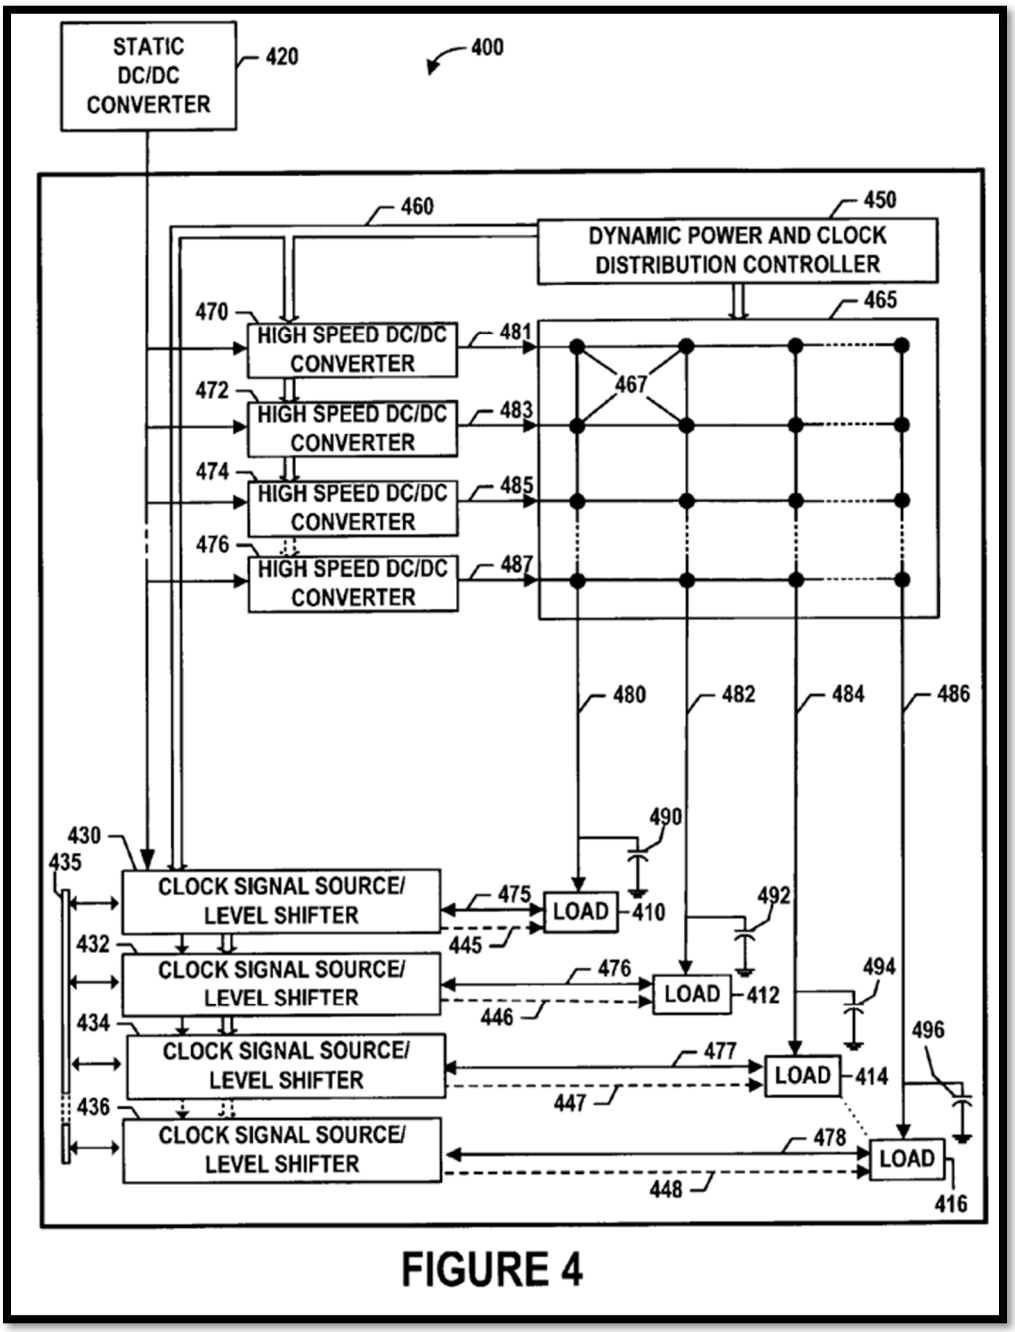

FIGURE 4

Figure 4 of *Shenai*

Ex-1004, p.5; Ex.-1002, ¶104. *Shenai's* controller 450 determines clock targets for each load while clock/shifters 430-436 provide the **clock signals to corresponding loads** 410-416. Ex-1004, [0026], [0035]; Ex.-1002, ¶¶104-107.[5]

b. (Element 1(b)) a plurality of domains;

Shenai teaches *a plurality of domains* (loads 410, 412, 414, 416) overlapping the domains discussed in the '443 Patent. *See supra* Section X(A)(1)(a); Ex.-1002, ¶¶104-107, 108-109. *Shenai's* Figure 4 follows.

[5] All emphasis and annotations added by Petitioner unless noted otherwise.

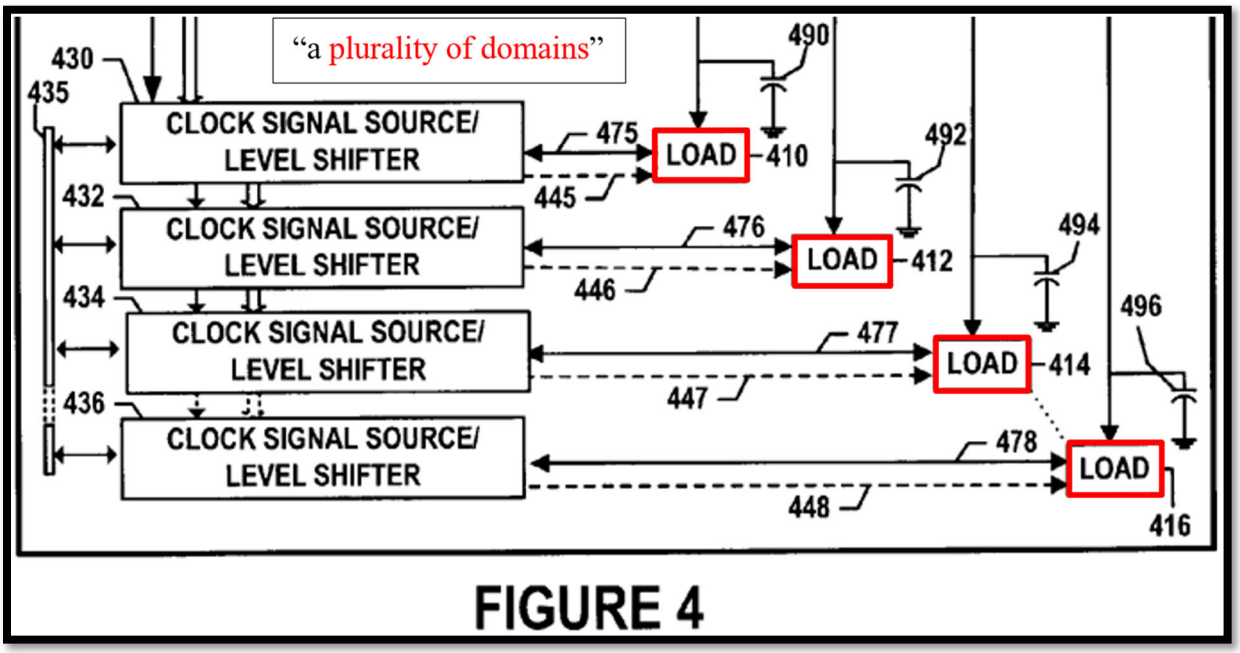

Figure 4 of *Shenai* (Excerpted)[6]

Ex-1004, p.5 (excerpted); Ex.-1002, ¶108.

c. (Element 1(c)) for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains; and

Shenai's "[c]lock/shifters 430-436 … [are] coupled with loads 410-416 so as to provide clock signals via clock signal buses 445, 446, 447 and 448 to the loads at the target frequencies" Ex.-1004, [0035]; Ex.-1002, ¶110. Clock/shifters 430-436

[6] Quoted claim elements provided in text boxes are added in annotated figures by Petitioner, unless indicated otherwise.

(each the recite "clock") include circuits generating clock signals for loads 410-416 (each a claimed "domain") by, for example, operating "at an integer multiplication, a half-integer multiplication or a quarter-integer multiplication of a reference clock signal." Ex.-1004, [0036]; Ex.-1002, ¶¶110-111. *Shenai's* Figure 4 follows.

"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"

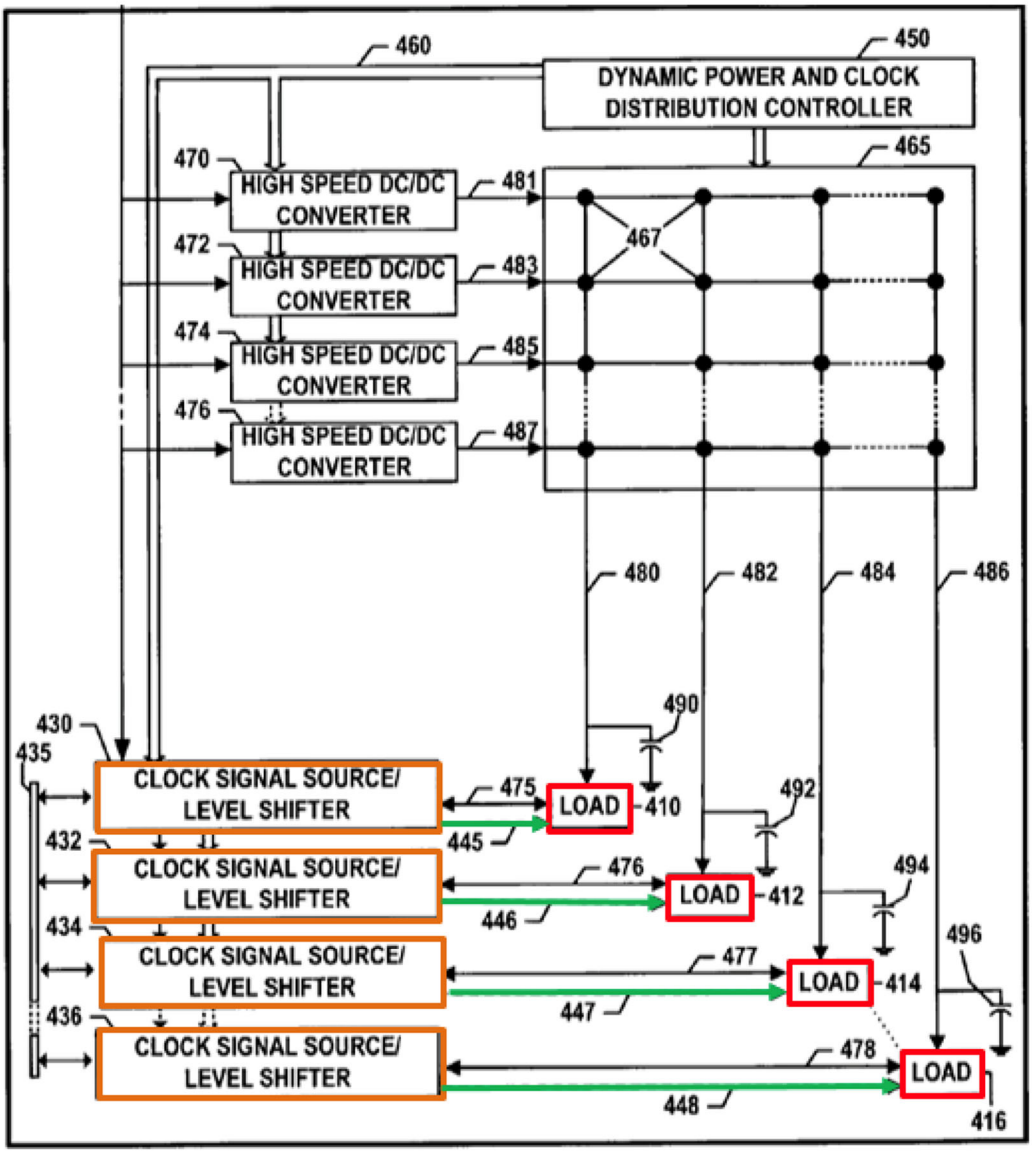

FIGURE 4

Figure 4 of *Shenai* (Excerpted)

Ex.-1004, p.5 (excerpted); Ex.-1002, ¶110.

"[C]lock/shifters 430-436 may comprise, phase-locked-loop circuits, delay-locked-loop circuits, or combinational logic circuits" that generate clock signals for loads 410-416 by, for example, operating at "at an integer multiplication, a half-integer multiplication or a quarter-integer multiplication of a reference clock signal." Ex.-1004, [0036]; Ex.-1002, ¶¶110-111. The loads are "identified as requiring higher (or lower) clocking frequency[.]" Ex.-1004, [0030]; Ex.-1002, ¶¶110-111. As seen above, clock/shifter 430 provides a clock signal via bus 445 to load 410, clock/shifter 432 provides a clock signal via bus 446 to load 412, and so forth. *Id.* The controller communicates with clock/shifters 430-436 "associated with each [respective load]" to "**establish appropriate** operating voltage and **frequency values**" as it "makes **dynamic determinations** of actual performance versus required performance, and **further adapts** voltage and **frequency parameters** of the various subsystems." Ex.-1004, [0028]-[0027]; Ex.-1002, ¶¶110-111. Such dynamic determination and corresponding adaptation of "frequency parameters of the various subsystems," based on required performance for each respective load, corresponds to the claimed "the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains." Ex.-1002, ¶¶110-113.

The loads are dynamically controlled by varying the generated clock signals based on operating conditions. Ex.-1004, [0043]; Ex.-1002, ¶¶111-113. The loads

can be partitioned to function at a voltage and frequency provided for each load. Ex.-1004, [0042]; Ex.-1002, ¶¶111-113. "[S]ome non-critical components may be provided with clock and voltage settings that allow them to operate [at a lower performance level] while only those load clusters that must operate at a higher level of performance [do so]." Ex.-1004, [0042]; Ex.-1002, ¶¶112-113. As a clock signal is dynamically generated for a given load at a frequency based on that load's respective operations, the clock signal for each load is generated independent of the clock signal for each other load. Ex.-1002, ¶¶110-113.

d. (Element 1(d)) for each of the plurality of domains, a voltage input for receiving a voltage which is dynamically changeable independently of the voltages applied to said others of the plurality of domains.

Shenai teaches receiving an "appropriate operating voltage" for each load. Ex.-1004, [0028]; Ex.-1002, ¶¶114-119. DC/DC converters are "responsive to controller 450, such that they are selectively coupled with loads 410-416 to provide the **respective target voltages** and currents **to the loads**." Ex.-1004, [0031]; Ex.-1002, ¶117. The controller communicates with the converters to "establish appropriate operating voltage[s]" as it "makes **dynamic determinations** of actual performance versus required performance, and **further adapts voltage** and frequency **parameters** of the various subsystems." Ex.-1004, [0028], [0027]; Ex.-1002, ¶¶116-118. This can be seen in *Shenai's* Figure 4 below.

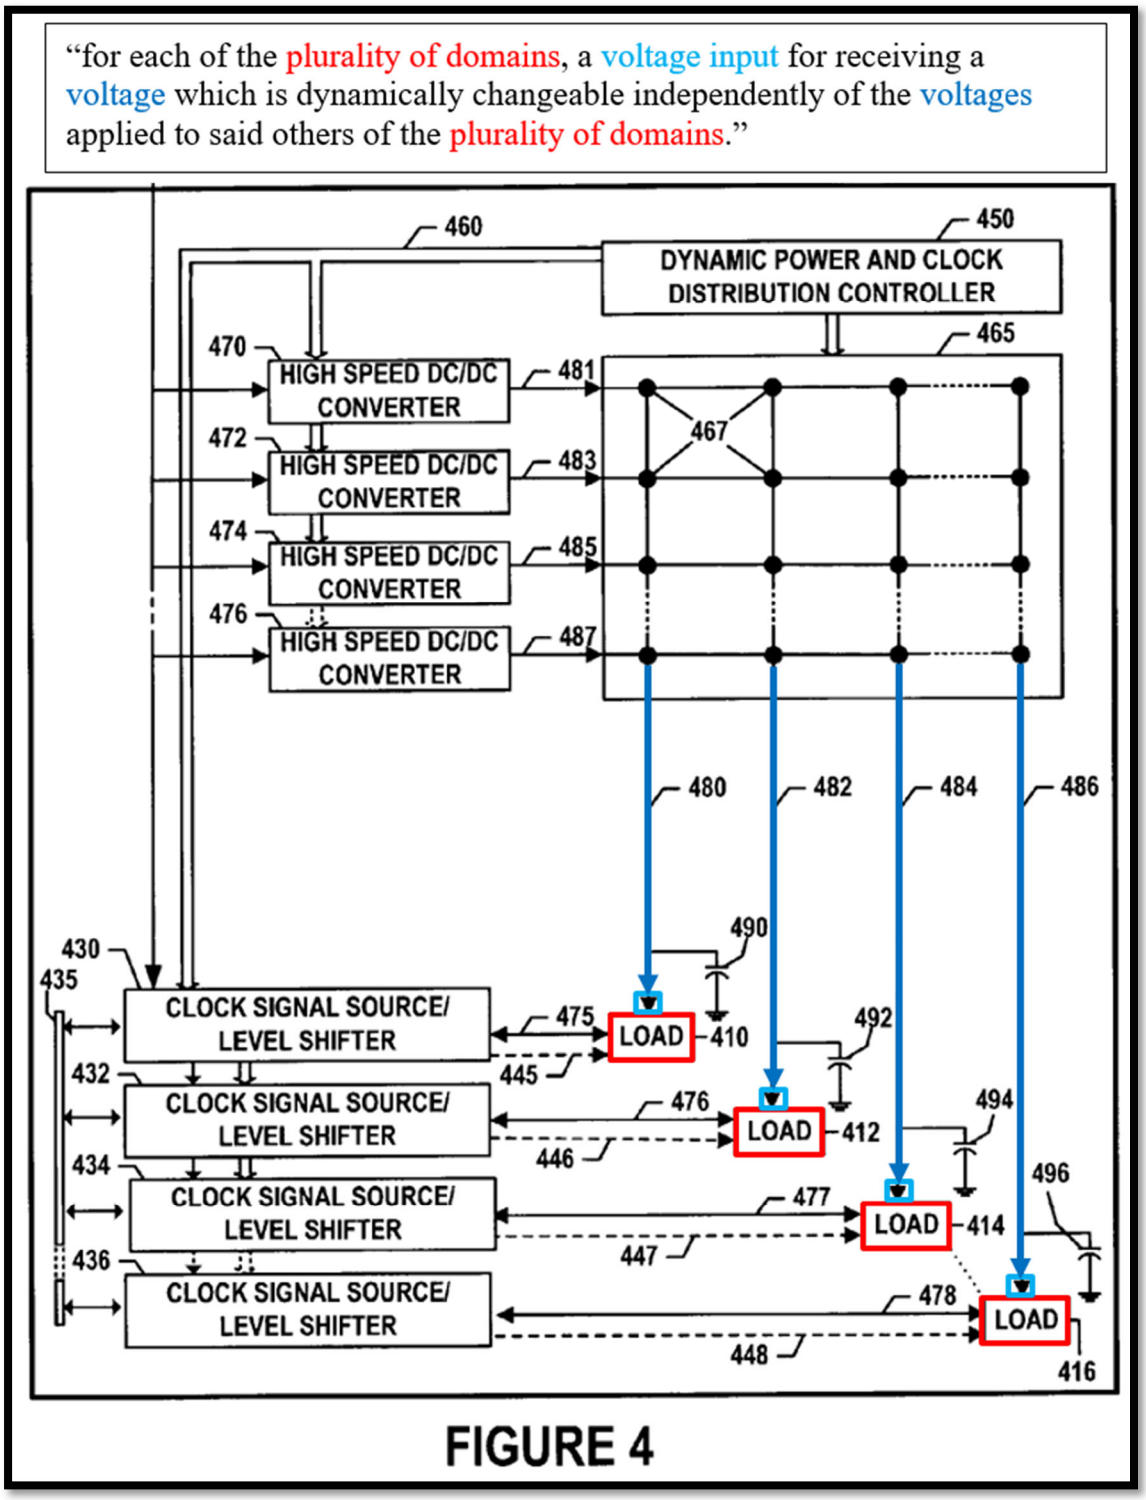

FIGURE 4

Figure 4 of *Shenai* (Excerpted)

Ex.-1004, p.5 (excerpted); Ex.-1002, ¶114.

"Network 465 may be responsive to controller 450 so as to selectively couple DC/DC converters 470-476 with loads 410-416 via connection network 465's nodes 467 and voltage supply lines 480[–]487." Ex.-1004, [0037]. As seen in Figure 4, voltage supply line 480 supplies voltage to the voltage input of load 410, voltage supply line 482 supplies voltage to the voltage input of load 412, and so forth. Ex.-1002, ¶115. A POSITA would have understood, for each load to receive the voltage from the voltage supply lines, loads must include a voltage input. Without an input, loads would have been unable to receive the voltage supplied. Ex.-1002, ¶115. A POSITA would have been motivated to apply known techniques of using a voltage input (e.g., an input pin) to known devices such as *Shenai's* microprocessor to yield the predictable result of receiving a voltage from a respective voltage supply line with a reasonable expectation of success. *Id*; *see KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 415-421 (2007).

Controller 450 determines "**voltage**, current and clock signal frequency **targets for the plural loads**" by communicating with DC/DC converters (e.g., DC/DC converters 470-476) to determine the appropriate operating voltage. Ex.-1004, [0026], [0028]; Ex.-1002, ¶¶116-118. By "sensing high levels of activity of certain strategically identified nodes, the requirements for those subsystems may be identified as requiring higher (or lower) clocking frequency, and hence higher (or

lower) **operating voltages**." *Id*. Based on this, DC/DC converters 470-476 are selectively coupled to loads 410-416 to provide "respective target voltages[.]" Ex.-1004, [0031]; Ex.-1002, ¶¶117-118.

Loads are dynamically controlled by varying the generated voltages based on operating conditions. Ex.-1004, [0043]; Ex.-1002, ¶¶118-119. Additionally, loads can be partitioned to function at a voltage and frequency provided for each load. Ex.-1004, [0042]; Ex.-1002, ¶¶118-119. "That is, some non-critical components may be provided with clock and voltage settings that allow them to operate [at a lower level of performance] while only those load clusters that must operate at a higher level of performance [do so]." *Id*. As voltage is dynamically changed for a given load based on that load's respective operations, the voltage for each load is changed independent of the voltage for each other load. Ex.-1002, ¶¶114-119.

2. *Claim 3 – The microprocessor of claim 1, wherein there are at least four of said domains.*

Shenai renders obvious the microprocessor of claim 1. *See supra* Section X(A)(1). *Shenai* includes loads 410, 412, 414, 416. Ex.-1004, [0025], p.5; Ex.-1002, ¶¶104-107, 120-121. These loads can be seen in Figure 4 of *Shenai* below. As discussed above, *Shenai's* loads correspond to the claimed "domains." *See supra* Section X(A)(1)(a); Ex.-1002, ¶¶104-107, 120-121.

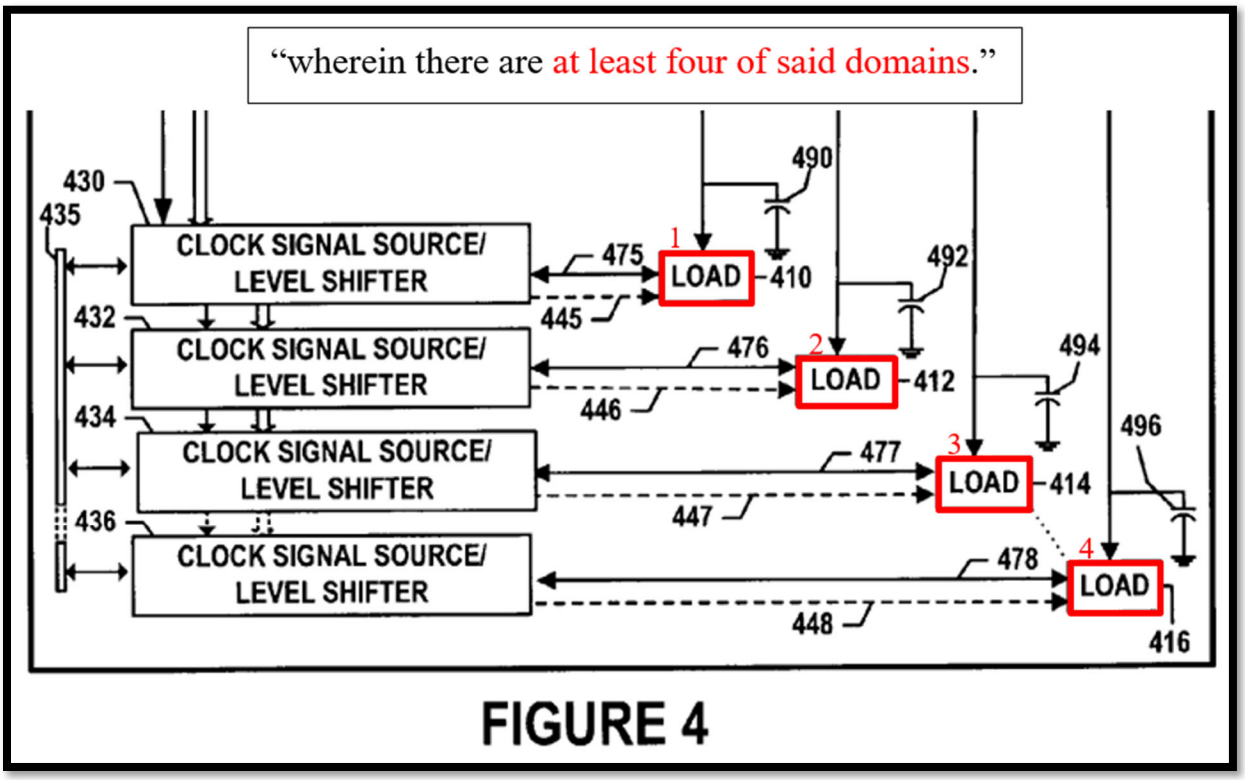

Figure 4 of *Shenai* (Excerpted)

Ex-1004, p.5 (excerpted); Ex.-1002, ¶120.

3. *Claim 4 – The microprocessor of claim 1, wherein the microprocessor is programmed to determine a slack in processing in one of the domains and to reduce the clock frequency and the voltage in said one of the domains to reduce the slack.*

Shenai renders obvious the microprocessor of claim 1. *See supra* Section X(A)(1). The '443 Patent explains slack "indicates…[a] previous operation completed earlier than necessary." Ex.-1001, 11:28-29; Ex.-1002, ¶122. Functional units performing different functions (e.g., at respective speeds) was known in the

art. Ex.-1002, ¶38; *see* Ex.-1014, p.8-9. A POSITA would have understood slack in the context of the '443 Patent as the difference in computing time between the slowest dependency, known as the critical path, and a different dependency within the same operation. Ex.-1002, ¶¶122-123, 45-52; *see* Ex.-1017, p.1; *see also* Ex.-1019, p.1; *see also* Ex.-1020, p.1. Thus, slack is a way to measure what dependencies are completing faster than needed to actively adjust the provided power. When power is reduced to dependencies with slack, their slack is reduced since the completion time will increase. *Id*.

Shenai teaches determining slack. Ex.-1002, ¶¶122-126; Ex.-1004, [0026], [0030]. For example, "controller 450 may comprise software instructions for determining the voltage, current and frequency targets based, on information provided by circuitry included with controller 450, which may **indicate an amount of circuit activity for each of the loads**." Ex.-1004, [0026]; Ex.-1002, ¶¶123-124. Controller 450 can sense "high levels of activity of certain strategically identified nodes, [and] the requirements for those subsystems may be identified as **requiring higher (or lower) clocking frequency**, **and hence higher (or lower) operating voltages**." Ex.-1004, [0030]; Ex.-1002, ¶¶123-124. Further, "controller 450 may comprise sensing circuits connected to **critical-nodes (nodes that tend to limit the rate at which a given operation may be completed) and responsively adjust**

(lower) frequencies of other non-critical signal paths, or adjust (increase) the frequency of the circuit elements containing the critical node paths." *Id*.

Sensing circuits "provide circuit activity information for loads 410-416" and "may determine a number of logic transitions in a predetermined time period." Ex.-1004, [0029], [0030]; Ex.-1002, ¶¶123-125. Calculating processing needs to adjust voltages and frequencies to save energy was also known in the art. Ex.-1002, ¶¶123-126, 45-52. A POSITA would have looked to use such known techniques to determine the processing needs of *Shenai's* loads to identify the critical and non-critical paths in the same way as was known in the art with a reasonable expectation of success. Ex.-1002, ¶¶123-126, 45-52; *see KSR*, 550 U.S. at 415-421.

Based at least on the above, determining "slack" associated with non-critical paths and lowering the frequencies of such non-critical paths while increasing the frequency of critical paths, to reduce "slack" would have been obvious to a POSITA based on *Shenai.* Ex.-1002, ¶¶122-126.

4. *Claim 5 – The microprocessor of claim 1, further comprising a queue for communication between at least two of the domains.*

Shenai renders obvious the microprocessor of claim 1. *See supra* Section X(A)(1). *Shenai* discloses intercommunication between loads. Ex.-1004, [0029], [0035]; Ex.-1002, ¶¶127-128. For example, *Shenai* states "**clock/shifters 430-436 may** also **include circuitry that allows** electrical signal **communication between**

loads operating at varying target voltages via signal buses 475, 476, 477 and 478. Such an approach could have been implemented using a queue and would have improved the noise tolerance of such electrical signal communication and/or allow respective loads to be ready to receive communication before such communication is initiated. Ex.-1004, [0035]; Ex.-1002, ¶¶127-128.

To the extent *Shenai* is found not to specify how communication between loads is carried out, using a queue for communicating between loads would have been obvious to a POSITA. Ex.-1002, ¶¶127-128. Using queues for communication between loads was known in the art. Ex.-1002, ¶128, 45-52; *see* Ex.-1015, p.1 (teaching FIFOs are queues that may hold instructions between communicating processors, processing units, etc., until the instructions are ready to be written or read. Since out-of-order superscalar processors are complex, microprocessors were designed to include an instruction queue.) A POSITA would have recognized that *Shenai's* circuitry allowing communication between loads would have been improved by implementing a queue in the same predictable way as known in the art. Ex.-1002, ¶128. A POSITA would have been motivated to combine prior art elements (e.g., the communication between loads taught in *Shenai* with queue structures known in the art) according to known methods to yield predictable results, and would have a reasonable expectation of success. Ex.-1002, ¶¶127-128; *see KSR*, 550 U.S. at 415-421.

5. *Claim 8*

a. (Element 8(a)) A method of operating a microprocessor, the method comprising:

To the extent the preamble is limiting, *a method of operating a microprocessor* would have been obvious in view of *Shenai*. *See supra* Section X(A)(1)(a); Ex.-1002, ¶¶104-107; *see also* Ex.-1010, p. 1. *Shenai* is directed to the "organization and **method of operation**" of a microprocessor or microcontroller. Ex.-1004, [0006]; Ex.-1002, ¶129. A POSITA would have applied known techniques (e.g., *Shenai's* method of operation) to a known device (e.g., *Shenai's* microprocessor) to yield predictable results (e.g., performing a method of operating a microprocessor) with an expectation of success. Ex.-1002, ¶129; *see KSR*, 550 U.S. at 415-421.

b. (Element 8(b)) (a) providing a plurality of domains in the microprocessor;

As discussed for element 1(b), *Shenai* provides a plurality of functional units in its microprocessor. See *supra* Section X(A)(1)(b); Ex.-1002, ¶130.

c. (Element 8(c)) (b) clocking each of the domains separately at a clock frequency;

As discussed for claim element 1(c), each of *Shenai's* clock/shifters provide clock frequencies to each respective load. *See supra* Section X(A)(1)(c); Ex.-1002, ¶131.

d. (Element 8(d)) (c) applying a voltage to each of the domains separately;

As discussed for element 1(d), each of *Shenai's* DC/DC converters apply independent voltages to each respective load. *See supra* Section X(A)(1)(d); Ex.-1002, ¶132.

e. (Element 8(e)) (d) operating the microprocessor such that each domain operates synchronously, while the domains operate asynchronously relative to one another; and

The circuitry of *Shenai* allows for a microprocessor to operate a plurality of loads at a plurality of independent frequencies and voltages which would cause the loads of the microprocessor to operate asynchronous relative to one another. Ex.-1004, [0035], Ex.-1002, ¶¶133-141. Clock/shifters 430-436 are coupled with loads 410-416 to "provide clock signals via clock signal buses 445, 446, 447 and 448 to the loads at the target frequencies" and include "circuitry that allows electrical signal communication between loads operating at varying target voltages via signal buses 475, 476, 477 and 478." Ex.-1004, [0035]; Ex.-1002, ¶¶133-134. *Shenai* explains its loads may be partitioned such that "each load of a partitioned (or dynamically partitioned) complex electronic system may function at voltage and frequency values that allow the load to realize [] power savings." Ex.-1004*,* [0042]; Ex.-1002, ¶¶133-134. Accordingly, *Shenai's* loads operate asynchronously relative to one another

based on each given load's respective voltage and frequency. *See supra* Sections X(A)(1)(c)-(d); Ex.-1002, ¶¶110-119, 138.

Shenai loads are intra-load synchronized. Ex.-1002, ¶¶134-137, 139, 141. *Shenai* explains that loads may be "independent systems." Ex.-1004, [0025]; Ex.-1002, ¶136. Each load is provided a respective frequency and voltage, based on its operating needs (e.g., for each load to realize power savings). Ex.-1002, ¶¶110-119, 138; *see supra* Sections X(A)(1)(c)-(d), *see also* Ex.-1004, [0042]. Accordingly, components of each independent system encompassing a load operate using that load's respective frequency and voltage, and each load (e.g., components in an independent system) operates intra-load synchronously using its respective frequency and voltage. Ex.-1002, ¶134, 136. Intra-load synchronization was also known in the art. Ex.-1002, ¶¶138-139; Ex.-1022, p.1 ("Leaving a synchronous block into the asynchronous domain is a reasonably trivial matter."). Globally asynchronous, locally synchronous (GALS) systems, as also described in the '443 Patent, were known in the art prior to the '443 Patent. Ex.-1002, ¶¶137, 139, 53-56. Such an architecture provided a system where loads operated synchronously (e.g., locally synchronous) while the system operated asynchronously (e.g., the domains operate asynchronously relative to one another). *Id.*, *see* Ex.-1022, p.1 (teaching that GALS systems have "a locally generated clock for every module" that is "locally distributed.") It would have been obvious to a POSITA to apply a known technique

(e.g., globally asynchronous, locally synchronous communication) to known devices (e.g., *Shenai's* loads that may be independent systems) ready for improvement to yield predictable results (e.g., synchronized intra-load communication). Ex.-1002, ¶¶137, 139, 141; *see KSR*, 550 U.S. at 415-421.

Shenai's loads are also inter-load synchronized. Ex.-1002, ¶¶134-135, 139-141. "[C]lock/shifters 430-436 may also include circuitry that allows electrical signal **communication between loads operating at varying target voltages** via signal buses 475, 476, 477 and 478." Ex.-1004, [0035]; Ex.-1002, ¶134. *Shenai* contemplates circuit operation configurations that "allow for synchronous communication between loads 410-416, as opposed to asynchronous communication, which may be undesirable." Ex.-1004, [0036]; Ex.-1002, ¶135. Moreover, *Shenai* describes a prior art microprocessor including a clock/shifter supplying a clock signal with circuitry to "modify the voltage level of electrical signals, such as the clock signals and signals communicated between load 310 and 315 via signal lines 375 and 376 when they are operating at different voltage levels." Ex.-1004, [0022]; Ex.-1002, ¶140. This is similar to the '443 Patent: "The preferred embodiment uses four domains, one of which includes the L2 cache, so that domains may vary somewhat in size and **still be covered by a single clock**." Ex.-1001, 5:7-9; Ex.-1002, ¶140. *Shenai* explains "level shifting techniques are known, and may

reduce the effect of noise on the communication of electrical signals between loads 310 and 315." Ex.-1004, [0022]; Ex.-1002, ¶140.

Inter-load synchronous communication was also known in the art. Ex.-1002, ¶¶134-135, 139-141, 53-56; *see* Ex.-1022*,* p.1 ("GALS systems [do] away with a global clock and instead [have] a locally generated clock for every module. … Obviously, modules cannot be [synchronized] if they all use different clocks, so we use an **asynchronous protocol to communicate between them**."). To implement inter-load synchronous communication, "GALS systems stretch the clock [by preventing] the clock from rising if asynchronous data is arriving" *Id.* It was known that "if the frequencies at which [] individual [processing elements 'PEs'] operate differ from one another and from other elements within the system where the multiprocessor chip is employed … a **synchronization schema** must be implemented when there is a need to communicate data between PEs (or with other system elements) that operate at different frequencies." Ex.-1023, 6:14-21; Ex.-1002, ¶¶137-138. "To effect such synchronization, each PE…is connection [sic] to an arrangement comprising elements 150 and 160. Level converter 150 converts the variable voltage swings of the PEs to a fixed level swing, and network 160 resolves the issue of different clock domains" *Id*., 6:29-31. Ex.-1002, ¶¶138-140, 53-56.

A POSITA would have applied known synchronization techniques (e.g., level shifting techniques as disclosed in *Shenai* and known the art such as clock stretching,

asynchronous protocols, or synchronization schemas) to *Shenai's* inter-load communication to yield predictable results for improved synchronized communication such by improving "the noise tolerance of such electrical signal communication." Ex.-1004, [0035]; Ex.-1002, ¶¶133-141, 53-56; *see KSR*, 550 U.S. at 415-421.

f. (Element 8(f)) (e) dynamically controlling the clock frequency and the voltage in each of the plurality of domains independently of the clock frequencies and the voltages in others of the plurality of domains.

Shenai dynamically controls the clock frequency and voltage for each given load independent of each other load, based on each given load's operating conditions. *See supra* Sections X(A)(1)(c)-(d); Ex.-1002, ¶¶110-119, 142.

6. *Claim 9*

a. (Element 9(a)) The method of claim 8, wherein step (e) comprises:

Shenai renders obvious the method of claim 8. *See supra* Section X(A)(5). *See also supra* Section X(A)(5)(f); Ex.-1002, ¶143.

b. (Element 9(b)) (i) determining a slack in processing in one of the domains; and

As discussed for claim 4, *Shenai* determines loads having non-critical paths. *See supra* Section X(A)(3); Ex.-1002, ¶144.

c. (Element 9(c)) (ii) reducing the clock frequency and the voltage in said one of the domains to reduce the slack.

As discussed for claim 4, *Shenai* reduces the clock frequency and voltage of loads identified as having non-critical paths. *See supra* Section X(A)(3); Ex.-1002, ¶145.

7. *Claim 12 – The method of claim 8, wherein there are at least four of said domains.*

As discussed for claim 3, *Shenai* discloses at least four loads. *See supra* Section X(A)(2); Ex.-1002, ¶146.

B. Ground 2: Claims 4 and 9 are obvious over *Shenai* in view of *Halepete*

1. *Claim 4*

To the extent it is argued claim 4 is not obvious in view of *Shenai* (*see supra* Section X(A)(3)), which Petitioner does not concede, a POSITA would have found claim 4 obvious over *Shenai* in view of *Halepete*. Ex.-1002, ¶¶147-155.

Halepete discloses a processor with a processing unit (e.g., a load) and "typically includes a number of other components which are known to those skilled in the art[.]" Ex.-1005, 5:5-9; Ex.-1002, ¶147. Processing unit 16 includes master control unit 18 accomplishing "clock and voltage control" by "monitoring the operating characteristics of the processor." *Id*., 5:10-14; Ex.-1002, ¶147. Again, the '443 Patent explains slack in processing "indicates that [a] previous operation completed earlier than necessary." Ex.-1001, 11:28-29; Ex.-1002, ¶148. *Halepete* discusses a system determining "when the **operating characteristics of the central**

processor are significantly different than required by the operations being conducted (Figure 2 increase and decrease frequency and voltage conditionals), **and chang[es] the operating characteristics of the central processor to a level commensurate with the operations being conducted** (Figure 2 increase and decrease frequency and voltage steps)." Ex.-1005, p.1, Abstract; Ex.-1002, ¶148.

Thus, a POSITA would have understood that *Halepete* teaches a processor programmed to determine slack and reduce frequency and voltage supplied to the functional unit with the slack. Ex.-1002, ¶¶147-155, 45-52. *Halepete* teaches a slow clock multiplied by a factor to generate the core frequency for processing unit 16, where the factor is "computed by the control software of the present invention which monitors the operation of the processor to determine from the characteristics of the processor just what frequency should be selected." *Id*., 7:1-5; Ex.-1002, ¶¶149-150. Thus, the system determines a certain frequency required for system operations. *Id*., 7:8-25; Ex.-1002, ¶¶149-150. The clock frequency provided by *Halepete's* PLL is divided by a "plurality of different values determined by the control software, [and] the operating frequencies for the different components of the system may be individually controlled and furnished to other components of the processor[.]" *Id*., 8:11-15; Ex.-1002, ¶¶149-150.

"The **detection of such operating characteristics** therefore may indicate that the **frequency and voltage of operation should be reduced**." *Id*., 11:5-7; Ex.-1002,

¶¶149-150. *Halepete* teaches control software that "detects operating characteristics and determines whether those characteristics indicate that the frequency and voltage of operation should be changed." *Id*., 11:14-16; Ex.-1002, ¶¶149-151. *Halepete* further discloses "the control software causes **the voltage to be lowered to the calculated value**…. [a]t that point, **the control software begins again to monitor the various conditions controlling the frequency and voltage of operation**." *Id*., 15:6-12; Ex.-1002, ¶¶149-151. It was also known in the art to dynamically change both clock frequency and voltage in a processing unit based on the unit's operations. Ex.-1002, ¶¶148, 152-155, 45-52; *see* Ex.-1012, p1; *see also* Ex.-1021, p.1, 3 (describing "a technique for adjusting supply voltage and frequency at run-time to save energy. Our technique monitors a program's instruction-level parallelism (ILP) and adjusts processor voltage and speed in response to the amount of observed ILP." "The processor measures ILP over some specified interval of time in terms of a [Million Instructions Per Second ('MIPS')] rate. If the MIPS rate changes appreciably either up or down, the processor adjusts its supply voltage accordingly." "Using this MIPS rate, a new frequency (and, hence, voltage level) is calculated.")

Halepete explains that "the programmable frequency generator is able to provide individual frequencies selectable for each [] component[]." Ex.-1005, 6:11-13; Ex.-1002, ¶¶147-150. *Shenai* explains that "controller 450 may comprise software instructions for determining the voltage, current and frequency targets[.]"

Ex.-1004, [0026]; Ex.-1002, ¶¶147. A POSITA would have recognized the benefits from using software as suggested by *Halepete* to control the allotted power to each load in a multi-load microprocessor of *Shenai*. Ex.-1002, ¶¶147-155, 45-52. For example, a POSITA would have understood dynamically controlling frequency and voltage would reduce thermal output, increase battery life, and provide overall system improvements to the microprocessor. *Id.* Such a combination would have been nothing more than a combination of prior art elements (e.g., using *Halepete's* control software with *Shenai's* multi-load microprocessor) according to known methods (e.g., changing a load's frequency and voltage based on detecting operating characteristics such as MIPS rates) to yield predictable results with an expectation of success. *Id; see KSR*, 550 U.S. at 415-421.

2. *Claim 9*

d. Element 9(a)

Shenai renders obvious the method of claim 8. *See supra* Section X(A)(5); Ex.-1002, ¶156.

e. Element 9(b)

As discussed for claim 4, *Shenai* in view of *Halepete* discloses monitoring operations of a load to identify which operating characteristics can be reduced based on the operations. *See supra* Section X(B)(1); Ex.-1002, ¶157.

f. Element 9(c)

As discussed for claim 4, *Shenai* in view of *Halepete* discloses a clock frequency and voltage of a load based on operating needs of the load. *See supra* Section X(B)(1); Ex.-1002, ¶158.

C. Ground 3: Claims 5-7 and 10-11 are obvious over *Shenai* in view of *Sherburne*

1. *Claim 5*

A microprocessor that includes communication between loads would have been obvious in view of *Shenai*. *See supra* Section X(A)(4). To the extent it is argued claim 5 is not obvious in view of *Shenai*, a POSITA would have found claim 5 obvious over *Shenai* in view of *Sherburne*. Ex.-1002, ¶¶159-166.

Sherburne describes FIFO buffer elements (the claimed "queue for communication between at least two of the domains"). *Id.*; Ex.-1006, [0004]. *Sherburne* teaches a system where "a data processing unit [] is fed data through a data buffering element (such as a first-in-first-out (FIFO)) and which outputs data to another data buffering element." Ex.-1002, ¶¶159-166; Ex.-1006, [0004]. The FIFO unit is between functional units (the claimed "domains"), as shown in *Sherburne's* Figure 3:

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

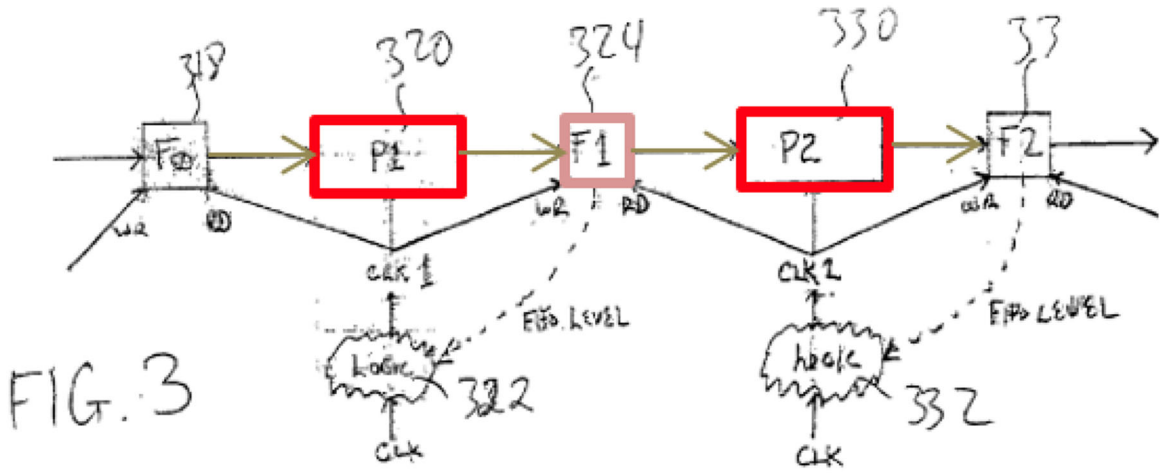

Figure 3 of *Sherburne* (Excerpted)

Ex.-1006, p.4 (excerpted); Ex.-1002, ¶¶162-164. Figure 3 shows processor 320 (e.g., first domain) receiving data from FIFO 318. Ex.-1006, [0031]; Ex.-1002, ¶¶162-164. FIFO 324 "accepts data generated by the processor [320]." *Id.* FIFO 324 "provides a FIFO fill level feedback signal to the controller 322 such as FIFO_FULL, FIFO_EMPTY" to assist in communicating if a read or write is available. *Id.* The output of FIFO 324 is received by processor 330 (e.g., second domain), which then supplies data to another buffer. Ex.-1006, [0029]; Ex.-1002, ¶¶162-164. "**Based on the FIFO fill level feedback signal**, the controller 322 can vary the clock signal to

the processor 320 and the FIFOs 318 and 324 to increase or decrease the processing speed of the processor 320, **thereby adjusting the read rate from the FIFO 318 and the fill rate** [(write rate)] **to the FIFO 324**." *Id.* Accordingly, *Sherburne* discloses a FIFO buffer (the "queue") adjusting read and write rates for communicating between a plurality of functional units. Ex.-1002, ¶¶162-164

Shenai describes "power supply systems and, more particularly, to frequency and voltage tuning of distributed power supply systems" and "approaches for balancing power consumption reduction and performance." Ex.-1004, [0002], [0004]; Ex.-1002, ¶¶165-166. Similarly, *Sherburne* describes "power clocking systems and methods" where "power can be saved by varying the clock frequency, the core voltage or a combination thereof, if necessary, to reduce heat or to reduce battery power consumption." Ex.-1006, p.1, Abstract, [0021]; Ex.-1002, ¶¶165-166. As *Shenai* and *Sherburne* are in the same field of endeavor, a POSITA would have been motivated to modify *Shenai´s* multi-load microprocessor to realize inter-load communication improvements based on *Sherburne*. Ex.-1002, ¶¶165-166

A POSITA would have found it obvious, and been motivated, to combine the explicit teachings of a buffer to help facilitate communication between functional units as *Sherburne* describes with a multi-load processor that is capable of inter-load communication as *Shenai* describes. Ex.-1002, ¶¶159-166; *see KSR*, 550 U.S. at 415-421. Using queues for communication between units as was also known in the art.

Id., ¶¶160, 41-44; *see* Ex.-1015, p.1 (teaching that a microprocessor instruction pipeline "writes the renamed instructions into the queues and reads the busy-bit table to determine if the operands are initially busy. Instructions wait in the queues until all their operands are ready.")

Shenai's loads 410-416 would have benefited from having a queue (e.g., *Sherburne's* FIFO buffer or other queues as generally known in the art) to facilitate communication between one another the same way *Sherburne's* functional units benefit from FIFO queues. *Id.*, ¶¶159-166. A POSITA would have understood the advantages of a queue (e.g., better communication, avoiding overflow, allowing for decoupling, increasing load flexibility, etc.) were well known in the art at the time of the '443 Patent. *Id*., ¶¶159-166, 41-44. A POSITA would have been motivated to combine prior art elements (e.g., *Shenai's* inter-load communication with *Sherburne's* FIFO buffer elements or other queues known in the art) according to known methods (e.g., FIFO fill level feedback signals, busy-bit tables, etc.) to yield predictable results with an expectation of success. *Id*; *see KSR*, 550 U.S. at 415-421.

2. *Claim 6*

a. (Element 6(a)) The microprocessor of claim 5,

Shenai in view of *Sherburne* renders obvious the microprocessor of claim 5. *See supra* Section X(C)(1); Ex.-1002, ¶167.

b. (Element 6(b)) wherein the queue has a Full flag and an Empty flag, and

Sherburne teaches this element. Ex.-1002, ¶¶168-170; Ex.-1006, [0040], [0028]-[0031]. The '443 Patent states that a Full flag "indicates to the producer that it can no longer write to the queue" and an Empty flag "indicates that there is no valid data for the consumer to read from the queue." Ex.-1001, 6:54-59. *Sherburne's* full indicators "are for the write (destination) FIFO… and empty [indicators] are for the read (source) FIFO." Ex.-1006, [0040]. *Sherburne* discloses "BUFFER_FULL" and "BUFFER_EMPTY" signals provided by each FIFO buffer element. Ex.-1002, ¶¶168-169; Ex.-1006, [0028]; *see also* Ex.-1006, [0029] ("[a buffer can provide a] buffer fill level feedback signal to the controller 332 such as BUFFER_FULL, BUFFER_EMPTY[.]"). *Sherburne* also describes each buffer element including high low water marks. Ex.-1006, [0039]; Ex.-1002, ¶¶168-169.

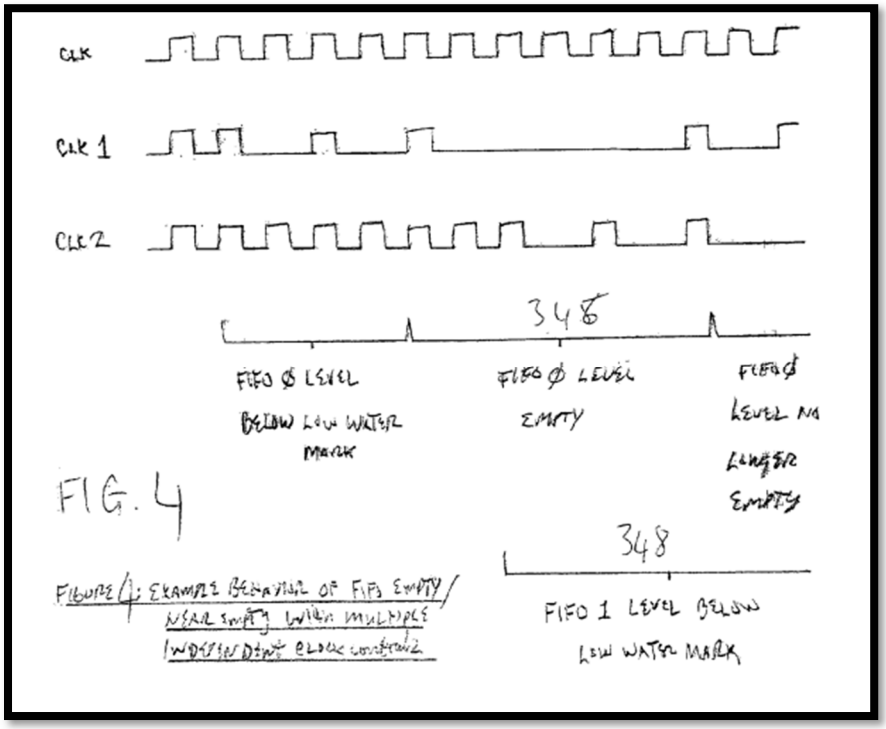

Figure 4 of *Sherburne* (Excerpted)

Ex.-1006, p.5 (excerpted); Ex.-1002, ¶168. "FIG. 4 shows an example of clock behavior including clock rate reduction during FIFO empty condition." Ex.-1006, [0039]. Instead of "BUFFER_FULL" or "BUFFER_EMPTY" signals, the FIFO can provide a "FIFO fill level feedback signal to the controller 312 such as FIFO_FULL, FIFO_EMPTY[.]" Ex.-1006, [0022]; Ex.-1002, ¶169. Flags have also been used in computing logic and programming well before the time of the '443 Patent. Ex.-1002, ¶¶41-44; *see also* Ex.-1016, p.19 ("The **flags** of the TI synchronous FIFOs bear the following designations: $\overline{\textbf{EMPTY}}$ = OR (OUTPUT READY) $\overline{\textbf{FULL}}$ = IR (INPUT READY)"). POSITA would have understood *Sherburne's* BUFFER_FULL and BUFFER_EMPTY or FIFO_FULL and FIFO_EMPTY signals to be equivalent to "Full flags" and "Empty flags," respectively. Ex.-1002, ¶¶169-170. To the extent it is argued that *Sherburne's* BUFFER_FULL and BUFFER_EMPTY or FIFO_FULL and FIFO_EMPTY signals are not "flags," as these signals indicate when read and write operations can be performed, a POSITA would have known that using these signals as flags would have been a simple substitution of one known element (*Sherburne's* signals) for another (flags, as known in the art) to obtain predictable results (e.g., to perform read and write operations). *Id*., ¶¶168-170, 41-44; *see KSR*, 550 U.S. at 415-421.

Based on *Sherburne*, a POSITA would have been motivated to improve similar communication between functional units in a microprocessor in the same way by providing the multi-load microprocessor of *Shenai* with a queue that has a BUFFER_FULL and BUFFER_EMPTY or FIFO_FULL and FIFO_EMPTY signals. Ex.-1002, ¶¶168-170, 41-44. Since queue structures and flag values were well known and conventional prior to the time of the '443 Patent, there would have been a reasonable expectation of success implementing the queue structure with flag values in *Shenai*'s microprocessor. *Id.* Implementing *Sherburne*'s buffer with BUFFER_FULL and BUFFER_EMPTY or FIFO_FULL and FIFO_EMPTY signals in *Shenai*'s microprocessor structure would have been nothing more than combining prior art elements (e.g., *Shenai*'s microprocessor structure with *Sherburne's* buffer with BUFFER_FULL and BUFFER_EMPTY or FIFO_FULL and FIFO_EMPTY signals) according to known methods (e.g., to yield predictable results) with a reasonable expectation of success. *Id*; *see KSR*, 550 U.S. at 415-421.

c. (Element 6(c)) wherein the microprocessor is programmed to prevent a write to the queue when the Full flag is asserted, until the Full flag is deasserted; and

d. (Element 6(d)) to prevent a read from the queue when the Empty flag is asserted, until the Empty flag is deasserted.

Sherburne discloses elements 6(c)-6(d). Ex.-1006, [0013], [0035]; Ex.-1002, ¶¶171-176. For example, *Sherburne* explains the clock is slowed down or stopped

"if the second buffer is full" and "if the FIFO is still empty, the processor clock may be reduced to a low rate or zero." Ex.-1006, [0013], [0035].

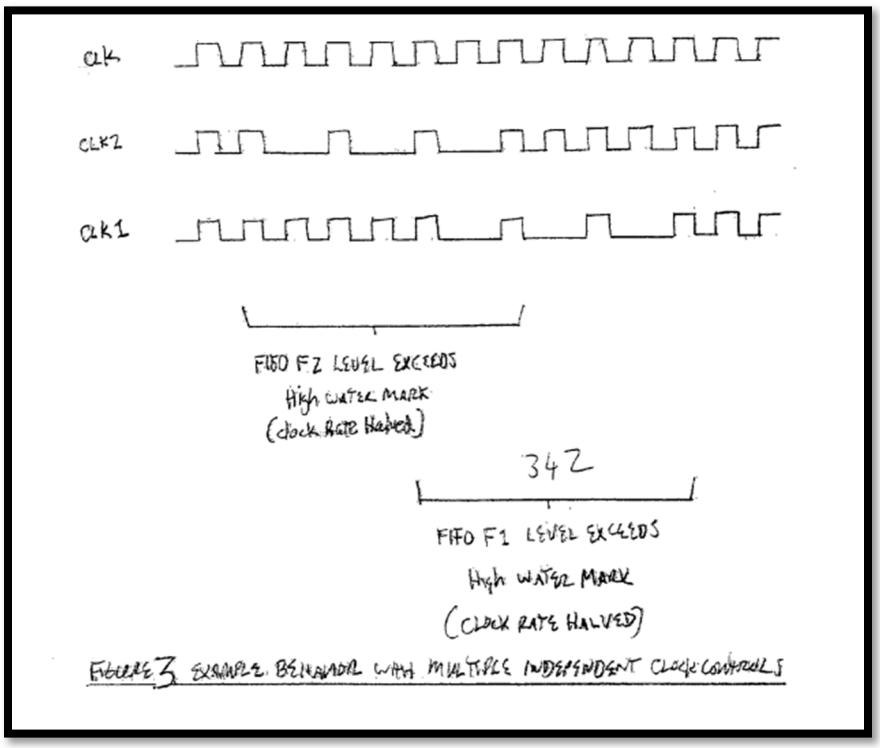

Figure 4 of *Sherburne* (Excerpted)

Ex.-1006, p.5 (excerpted); Ex.-1002, ¶172. A FIFOs accepts data generated by the previous functional unit. Ex.-1006, [0031]; Ex.-1002, ¶172. That data "is gated by a write signal from the controller 322, and the FIFO provides a FIFO fill level feedback signal to the controller 322 such as FIFO_FULL, FIFO_EMPTY[.]" Ex.-1006, [0031]; Ex.-1002, ¶¶172, 47; *see* Ex.-1018, p.1 (explaining that less critical units are gated "without delaying completion of processing in the critical units.") Based on the signal, "the controller 322 can vary the clock signal to the processor

320 and the FIFOs 318 and 324 to increase or decrease the processing speed of the processor 320, **thereby adjusting the read rate from the FIFO 318 and the fill rate** [(write rate)] **to the FIFO 324**." Ex.-1006, [0031]; Ex.-1002, ¶172.

Sherburne's system prevents further writes to a buffer that has provided a full signal to "avoid overflowing." Ex.-1006, [0023]; Ex.-1002, ¶173. If the buffer nears full, "as defined by some 'high water' mark," or if the buffer is full, "the clock rate to [the] processor … may be reduced to … zero." Ex.-1006, [0030]; Ex.-1002, ¶173. The published application of *Sherburne* includes claim 12 that states "slowing down or **stopping the clock** if the second buffer is above its high water mark or **if the second buffer is full**." Ex.-1006, p.16; Ex.-1002, ¶173.

The system of *Sherburne* also prevents further reads to a buffer that has provided an empty flag to conserve power. Ex.-1006, [0035]; Ex.-1002, ¶174. During "periods of FIFO empty condition, the units it feeds may have little or no computation to perform. Ex.-1006, [0035]; Ex.-1002, ¶174. Therefore a given system design may be amenable to reduced clock rate in presence of low or zero data input rate." *Id*., Ex.-1006, [0035]; Ex.-1002, ¶174. *Sherburne* continues, disclosing: "Once the source FIFO is empty, **assume the last data read** will be subject to processing and later written out to the next FIFO after the processor." Ex.-1006, [0035]; Ex.-1002, ¶174. *Sherburne* also includes claim 13 that states "slowing down

or **stopping the clock** if the first buffer is below its low water mark or **if the first buffer is empty**." Ex.-1006, p.16; Ex.-1002, ¶174.

A POSITA would have found it obvious to implement *Sherburne*'s queue structure in *Shenai*'s multi-load microprocessor structure as *Shenai*'s microprocessor structure already includes computing logic to facilitate communication between domains. Ex.-1004, [0035]; Ex.-1002, ¶¶171-175. As explained in reference to claim 5 and element 6(b), a POSITA would have been motivated to use teachings of *Sherburne* for *Shenai's* microprocessor architecture. *See supra* Sections X(C)(1), X(C)(2)(d); Ex.-1002, ¶¶159-166, 168-170. This would have been nothing more than a combination of prior art elements according to known methods (e.g., *Shenai's* microprocessor structure including computing logic to facilitate communication between domains with *Sherburne's* queue structure) to yield predictable results (e.g., to prevent a write when a full flag is asserted until it is deasserted and prevent a read when an empty flag is asserted until it is deasserted) with an expectation of success. Ex.-1002, ¶¶171-175; *see KSR*, 550 U.S. at 415-421.

3. *Claim 7 – The microprocessor of claim 6, wherein the queue is implemented as a dual-ported SRAM.*

Shenai in view of *Sherburne* renders obvious the microprocessor of claim 6. *See supra* Section X(C)(2). *Sherburne* discloses "the buffer 314 is a first-in, first-out memory circuit (hereinafter FIFO memory) using a dual port RAM[.]" Ex.-1006,

[0026]; Ex.-1002, ¶¶176-181. A POSITA would have understood that the claimed "dual-ported SRAM" is a type of the dual port RAM taught by *Sherburne* and was known in the art for use with microprocessors. Ex.-1002, ¶¶175-181, 34-40; *see* Ex.-1026, 5-67 (titled "3200DX **Dual-Port Random Access Memory (RAM)**" which references its 3200DX product as "dual-port SRAM" which "offers dual-port SRAM capable of supporting a system speed of 100 MHz.") It was also known in the art the "RAM family includes two important memory devices: static RAM (SRAM) and dynamic RAM (DRAM)." Ex.-1027, p.2; Ex.-1002, ¶¶177, 40. Based on *Sherburne's* use of a dual port RAM, a POSITA would have recognized that selecting a dual-ported SRAM would have been a simple design choice from a limited number of memory options in the RAM family. Ex.-1002, ¶¶177-181, 40; *see KSR*, 550 U.S. at 415-421; *see also Uber Techs. v. X One, Inc.*, 957 F.3d 1334, 1340 (Fed. Cir. 2020) ("Because a person of ordinary skill 'has good reasons to pursue the known options within his or her technical grasp'; § 103 bars the patentability of such obvious variations." (citations omitted))

As discussed above, *Shenai* and *Sherburne* are in the same field of endeavor. *See supra* Section X(C)(1); Ex.-1002, ¶179. A POSITA would have been motivated to modify *Shenai´s* multi-load microprocessor to realize inter-load communication improvements based on *Sherburne*, and *Shenai*'s microprocessor would have benefited from the implementation of a dual-port SRAM buffer for the same reasons

as *Sherburne*. Ex.-1002, ¶¶177-181. Therefore, a POSITA would have been motivated to use the known technique of implementing a dual-ported SRAM queue, as obvious in view of *Sherburne*, to optimize communication between domains of the microprocessor of *Shenai* in the same way. *Id.* Since queue structures and dual-ported SRAM were well known and conventional prior to the '443 Patent, there would have been a reasonable expectation of success implementing the dual-ported SRAM queue of *Sherburne* (obvious in view of the dual port RAM explicitly disclosed in *Sherburne*) in *Shenai*'s microprocessor. *Id.* Implementing such a dual-ported SRAM queue in *Shenai*'s microprocessor structure would have been nothing more than combining prior art elements (e.g., *Shenai's* microprocessor with a dual-port SRAM as obvious based on *Sherburne*) according to known methods (e.g., for use with *Sherburne's* buffers) to yield predictable results. *Id*; *see KSR*, 550 U.S. at 415-421.

4. *Claim 10 – The method of claim 8, wherein step (d) comprises providing a queue for communication between at least one of the domains.*

Shenai renders obvious the method of claim 8. *See supra* Section X(A)(5). Additionally, as discussed for claim 5, *Shenai* in view of *Sherburne* discloses providing buffers for communication between two loads. *See supra* Section X(C)(1); Ex.-1002, ¶182.

5. *Claim 11*

a. (Element 11(a)) The method of claim 10,

Shenai in view of *Sherburne* renders obvious the method of claim 10. *See supra* Section X(C)(4); Ex.-1002, ¶183.

b. (Element 11(b)) wherein the queue has a Full flag and an Empty flag, and wherein step (d) further comprises:

As discussed for element 6(b), *Shenai* in view of *Sherburne* discloses a buffer that has BUFFER_FULL, BUFFER_EMPTY, FIFO_FULL, and FIFO_EMPTY signals. *See supra* Section X(C)(2)(b)); Ex.-1002, ¶184.

c. (Element 11(c)) preventing a write to the queue when the Full flag is asserted, until the Full flag is deasserted; and

d. (Element 11(d)) preventing a read from the queue when the Empty flag is asserted, until the Empty flag is deasserted.

As discussed for elements 6(c)-6(d) *Shenai* in view of *Sherburne* renders obvious preventing write operations based on Full signals and preventing read operations based on Empty signals. *See supra* Sections X(C)(2)(c)-X(C)(2)(d); Ex.-1002, ¶¶185-186.

D. Ground 4: Claim 7 is obvious over *Shenai* in view of *Sherburne* and in view of *Traylor*

To the extent it is argued claim 7 is not obvious over *Shenai* in view of *Sherburne* (*see supra* Section X(C)(3)), a POSITA would have found claim 7 obvious over *Shenai-Sherburne* combination in view of *Traylor*. Ex.-1002, ¶¶187-191.

Sherburne explains the buffer "is a first-in, first-out memory circuit (hereinafter FIFO memory) using a dual port RAM." Ex.-1006, [0026]; Ex.-1002, ¶187. A POSITA would have found it obvious to use an SRAM for the teachings of *Sherburne* in view of the teachings of *Traylor*. *See supra* Section X(C)(3); Ex.-1002, ¶¶187-191, 176-181.

Just like *Sherburne*, *Traylor* discloses a FIFO buffer queue with an "EMPTY signal" and a "FULL signal[.]" Ex.-1007, 1:39-50; Ex.-1002, ¶188. The read counter and the write counter are "coupled to a **two-port Static Random Access Memory (SRAM)** 330." Ex.-1007, 4:59-61; Ex.-1002, ¶¶187-188. "The SRAM 330 is the memory device [] stores data in the FIFO." Ex.-1007, 4:61-62; Ex.-1002, ¶188. *Traylor* explains prior art systems implemented FIFOs using SRAMs. Ex.-1007, 1:19-21 ("Generally, a FIFO is comprised of a memory device, e.g. **a Static Random Access Memory (SRAM)**, and FIFO controller."); Ex.-1002, ¶188. A POSITA would have been well aware of using a dual-port SRAM for a queue structure such as a FIFO. Ex.-1002, ¶¶187-191, 40.

Based on *Traylor*, it would have been obvious to a POSITA to similarly configure *Shenai*'s microprocessor architecture to include a buffer as described in *Sherburne* and *Traylor* and for that buffer to be a two-port (e.g., "dual-ported") SRAM as described in *Traylor*. *Id.*, ¶¶187-191. *Traylor* and *Sherburne* describe FIFO elements providing buffer queues between functional units and have Full and

Empty flag signals. Moreover, *Sherburne* states FIFO buffers are dual-port RAM devices, and *Traylor* provides a more specific description for the buffers as dual-port SRAM devices. *Id.*

Thus, a POSITA would have been motivated to use an SRAM to implement the FIFO described in *Sherburne* as taught by *Traylor* to realize the advantages of an SRAM (e.g., faster access times and simpler control mechanisms, making it much more suitable for high-speed buffering.) Ex.-1007, p.1, Abstract; Ex.-1002, ¶¶187-191. For example, the system of *Shenai*, being a multi-load microprocessor, would have been a high-speed buffering system. *Id.*, ¶¶190-191. Multi-domain microprocessors often use dual-ported SRAM to allow a plurality of domains to access memory simultaneously. *Id.*; *see* Ex.-1026, p.1, 2.

A POSITA would have known to use known techniques (e.g., implementing a dual-ported SRAM queue, as disclosed in *Traylor* and known in the art), to optimize communication between domains of *Shenai's* microprocessor in the same way. Ex.-1002, ¶¶187-191. Since dual-ported SRAM and queue structures were known in the art prior to the '443 Patent, there would have been a reasonable expectation of success for implementing the *Traylor*-modified dual-ported SRAM queue of *Sherburne* in *Shenai*'s microprocessor, to yield predictable results. *Id*; *see KSR*, 550 U.S. at 415-421.

E. Ground 5: Claims 1 and 3 are obvious in view of *Georgiou*

1. *Claim 1*

a. Element 1(a)

To the extent the preamble is limiting, *Georgiou* teaches *a multiple clock domain microprocessor*. Ex.-1002, ¶¶192-194. For example, *Georgiou* teaches a "system and method for cooperatively varying the voltage and operating frequency of a functional unit in a multi-functional unit processor such as a superscalar **microprocessor**" where the functional units may be "fixed-point units, floating point units, load/store units, and branch units found in typical general purpose and superscalar microprocessors." Ex.-1008, 3:49-63; Ex.-1002, ¶192. A "**clock** selector" is associated with each "particular functional unit" to provide the varying clock signal and voltage for that functional unit. Ex.-1008, 8:37-67; Ex.-1002, ¶193. Multiplexers of clock selector are associated with each "independently controllable functional unit" such that "each functional unit may be operating under a different voltage 325 and frequency 425 combination according to the thermal feedback from its associated thermal sensor 119." Ex.-1008, 5:2-6; Ex.-1002, ¶193. Accordingly, multiplexers 470 provide a plurality of independent clocks for the claimed "multiple clock" domain microprocessor. *Id.*

Georgiou teaches a plurality of functional units 116-1, 116-2,…116-K (each a claimed "domain"). *Id.*, 5:35-38 ("The dispatch unit 114 analyses dependencies

among pending instructions, and schedules execution of instructions, one on each of functional units 116-1 through 116-K."). Ex.-1002, ¶¶192-194.

The '443 Patent describes a microprocessor that includes "integer issue/execute 120" domain, "floating point issue/execute 130" domain, and "load/store issue/execute 140" domain. Ex.-1001, 6:4-8; Ex.-1002, ¶194. Similarly, *Georgiou* states its functional units can be a "floating point unit" (e.g., the floating point issue/execute 130 domain described in the '443 Patent) or "multiple load/store and fixed point units" (e.g., load/store issue/execute 140 domain described in the '443 Patent). Ex.-1008, 4:52-54; Ex.-1002, ¶¶192-194. Based on this overlap between the '443 Patent's domains and *Georgiou's* functional units, a POSITA would have recognized that *Georgiou's* functional units correspond to the claimed "domains." Ex.-1002, ¶¶192-194. To the extent it is argued that *Georgiou's* functional units are not "domains," based at least on the overlap discussed above, a POSITA would have known that applying *Georgiou's* functional units as the claimed "domains" would have been a simple substitution of one known element (*Georgiou's* functional units) for another (domains) to obtain predictable results. *Id*; *see KSR*, 550 U.S. at 415-421.

b. Element 1(b)

Georgiou teaches *a plurality of domains* (functional units 116-1 through 116-K) similar to the domains discussed in the '443 Patent. *See supra* X(E)(1)(a); Ex.-1002, ¶¶195-196. The architecture of *Georgiou* is shown in Figure 1 below:

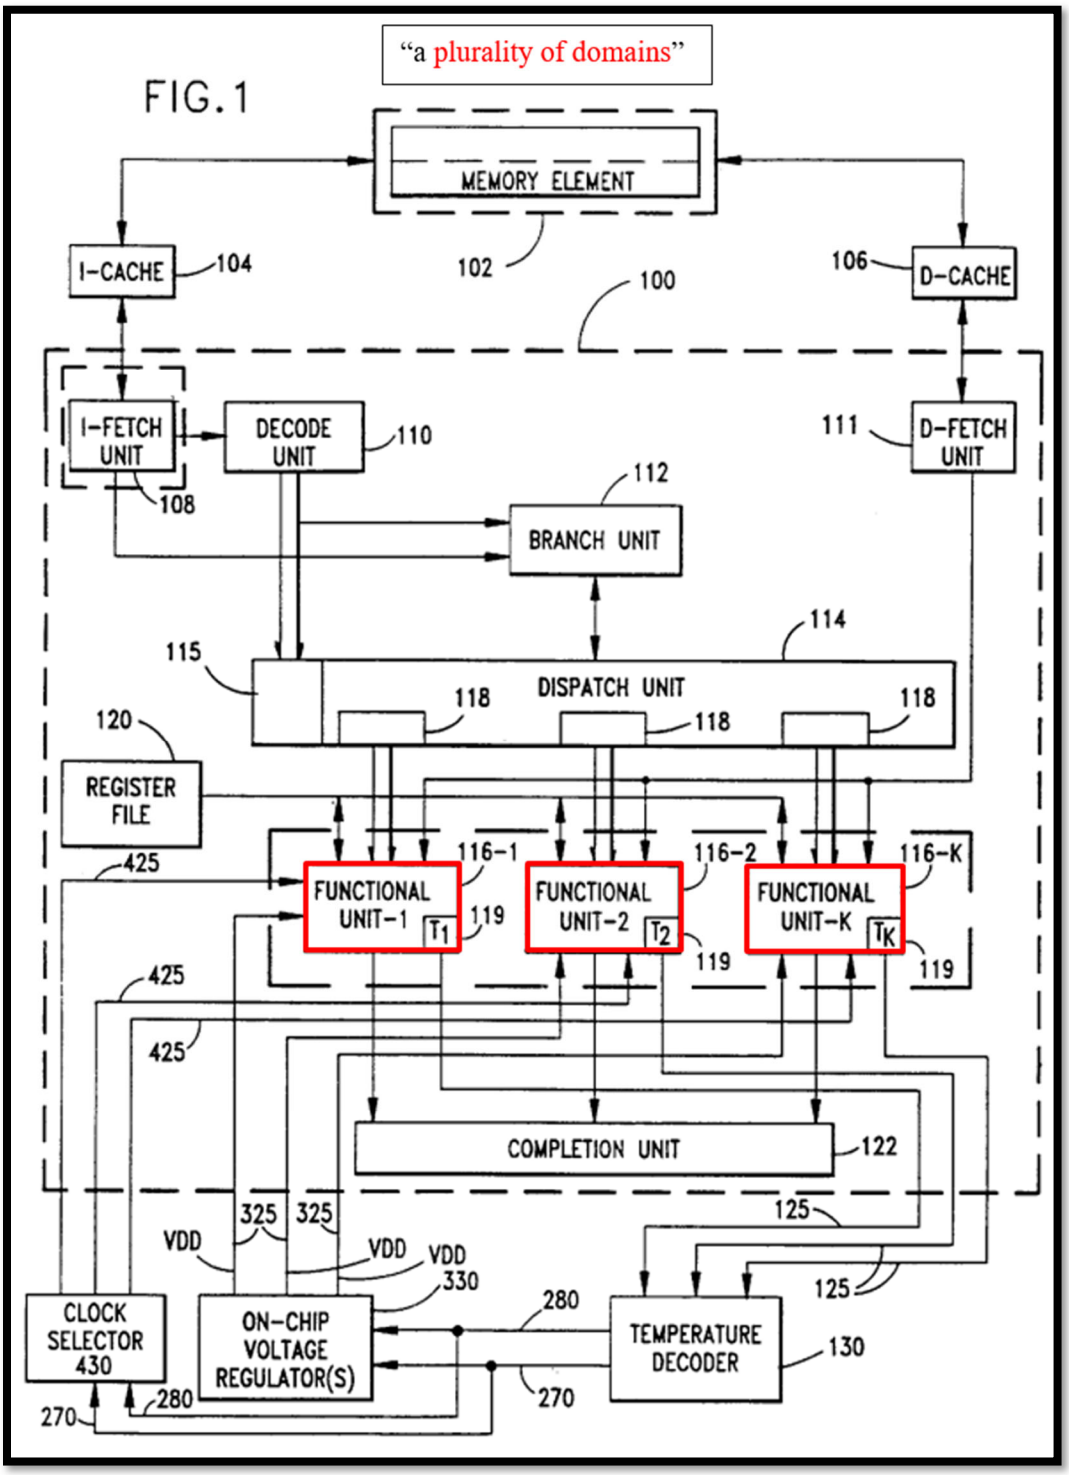

Figure 1 of *Georgiou*

Ex.-1008, p.3; Ex.-1002, ¶195.

c. Element 1(c)

Georgiou discloses a thermal feedback system for "**varying the** voltage and **operating frequency of a functional unit** in a single or multi-functional unit processor to control heating." Ex.-1008, 1:22-26; Ex.-1002, ¶¶197-201. The "modulated clock frequency 425 which corresponds to the reduced supply voltage 325 is applied to the clock inputs of the associated functional unit." Ex.-1008, 4:41-44; Ex.-1002, ¶¶197-198. Figure 1 follows:

"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"

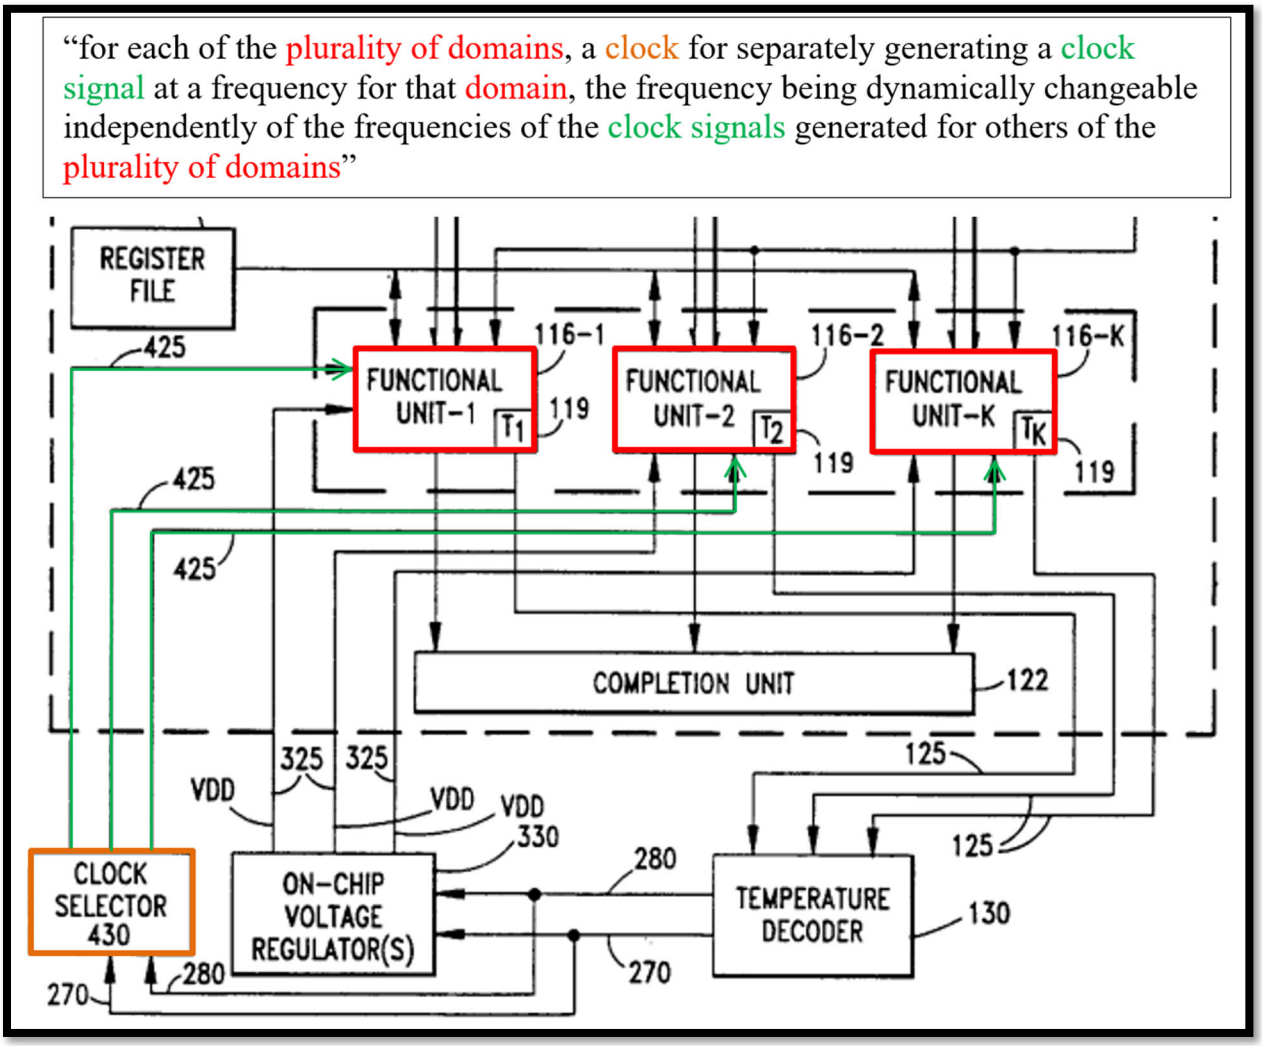

Figure 1 of *Georgiou* (Excerpted)

Ex.-1008, p.3 (excerpted); Ex.-1002, ¶197. Current/clock control signal 270 is coupled to the "clock selector 430 to cooperatively modulate the clock frequency[.]" Ex.-1008, 4:37-39; Ex.-1002, ¶¶197-198. As shown in annotated Figure 4 below, respective clock frequency 425 is generated by multiplexer 470 corresponding to a respective functional unit, and respective clock frequency 425 is then applied to each clock input of a corresponding functional unit such that the "clock frequency of the functional unit is modulated[.]" Ex.-1008, 4:61-62; Ex.-1002, ¶¶197-198. Multiplexers 470 of clock selector 430 are respectively associated with each "**independently** controllable functional unit" such that "**each functional unit may be operating under a different** voltage 325 and **frequency 425** combination according to the thermal feedback from its associated thermal sensor 119." Ex.-1008, 5:2-6; Ex.-1002, ¶¶197-198. "Thus, the **modulated** voltage and **frequency are applied to the functional unit to dynamically control heating** while maintaining synchronization and reliability." *Id.*, 10:48-51; Ex.-1002, ¶¶197-198.

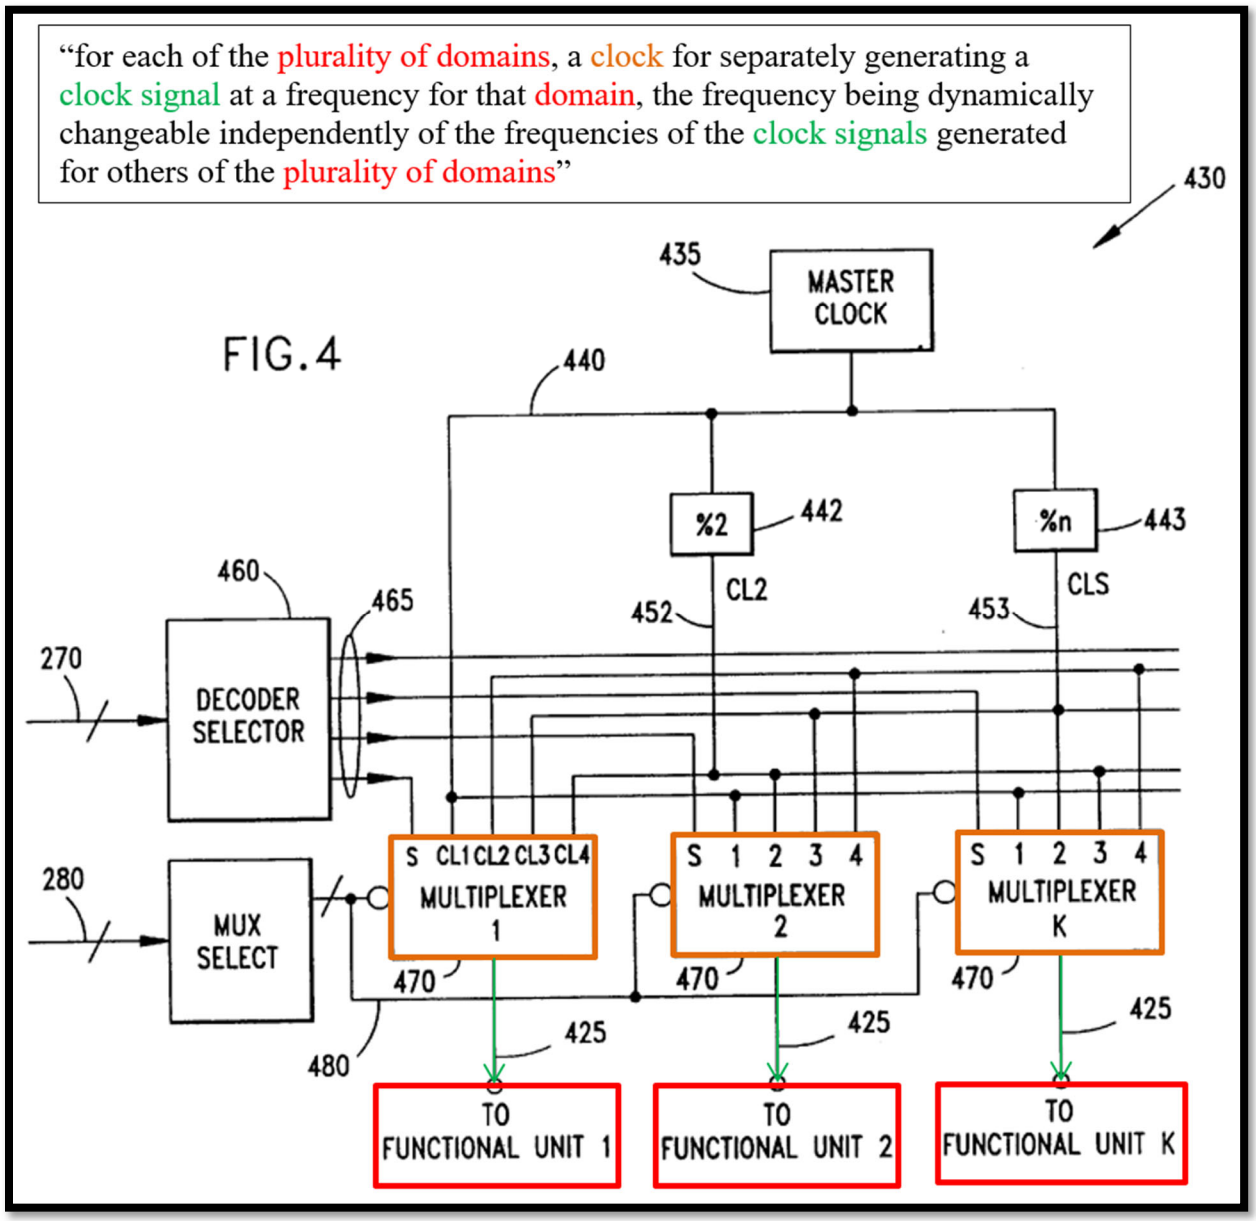

Figure 4 of *Georgiou*

Ex.-1008, p.3; Ex.-1002, ¶199. Each multiplexer 470 corresponds to the claimed "clock" for separately generating a clock signal at a frequency for a corresponding functional unit. Ex.-1002, ¶¶193, 197-201. Figure 4 shows clock selector 430 with

master clock 435 running at master clock frequency 440 used "as the common source and provides a common clock reference for a number of clock dividers[.]" Ex.-1008, 8:37-41; Ex.-1002, ¶¶199-201. This is similar to what is described in the '443 Patent: "The preferred embodiment uses four domains, one of which includes the L2 cache, so that domains may vary somewhat in size and **still be covered by a single clock**." Ex.-1001, 5:7-9; Ex.-1002, ¶200. In *Georgiou's* Figure 4, it can be seen that clock dividers can manipulate the master clock frequency by an integer value, such as divide-by-2 or divide-by-n. Ex.-1008, 8:42-48; Ex.-1002, ¶200. The system then "outputs a clock select signal 465 to a select input (S) to the multiplexer associated with the particular functional unit." *Id.* Clock select signal 465 is selected and latched into the appropriate multiplexor according to the functional unit. Ex.-1008, 8:58-61; Ex.-1002, ¶200.

"Each multiplexer 470 outputs the selected clock frequency which is coupled to the **particular functional unit** via the appropriate one of paths 425." Ex.-1008, 8:61-67; Ex.-1002, ¶201. *Georgiou* provides a "system and method for cooperatively **varying the** voltage and **operating frequency of a functional unit** in a single or multi-functional unit processor." Ex.-1008, 1:22-26; Ex.-1002, ¶197. Each multiplexer 470 provides a clock signal to its respective functional unit independent of the other clock signals, and the clock signals are dynamically changed depending on the individual needs of each functional unit at a given time. Ex.-1002, ¶¶197-201.

d. Element 1(d)

Georgiou teaches this element. Ex.-1002, ¶¶202-204. For example, *Georgiou* discloses "a thermal feedback system and method for cooperatively **varying the voltage** and operating frequency **of a functional unit** in a single or multi-functional unit processor to control heating." Ex.-1008, 1:22-26; Ex.-1002, ¶202. The output of the adjustable voltage regulator "is applied to the voltage supply of the functional unit[.]" Ex.-1008, 4:38-41; Ex.-1002, ¶202. This can be seen in annotated Figure 1 (excerpt) below.

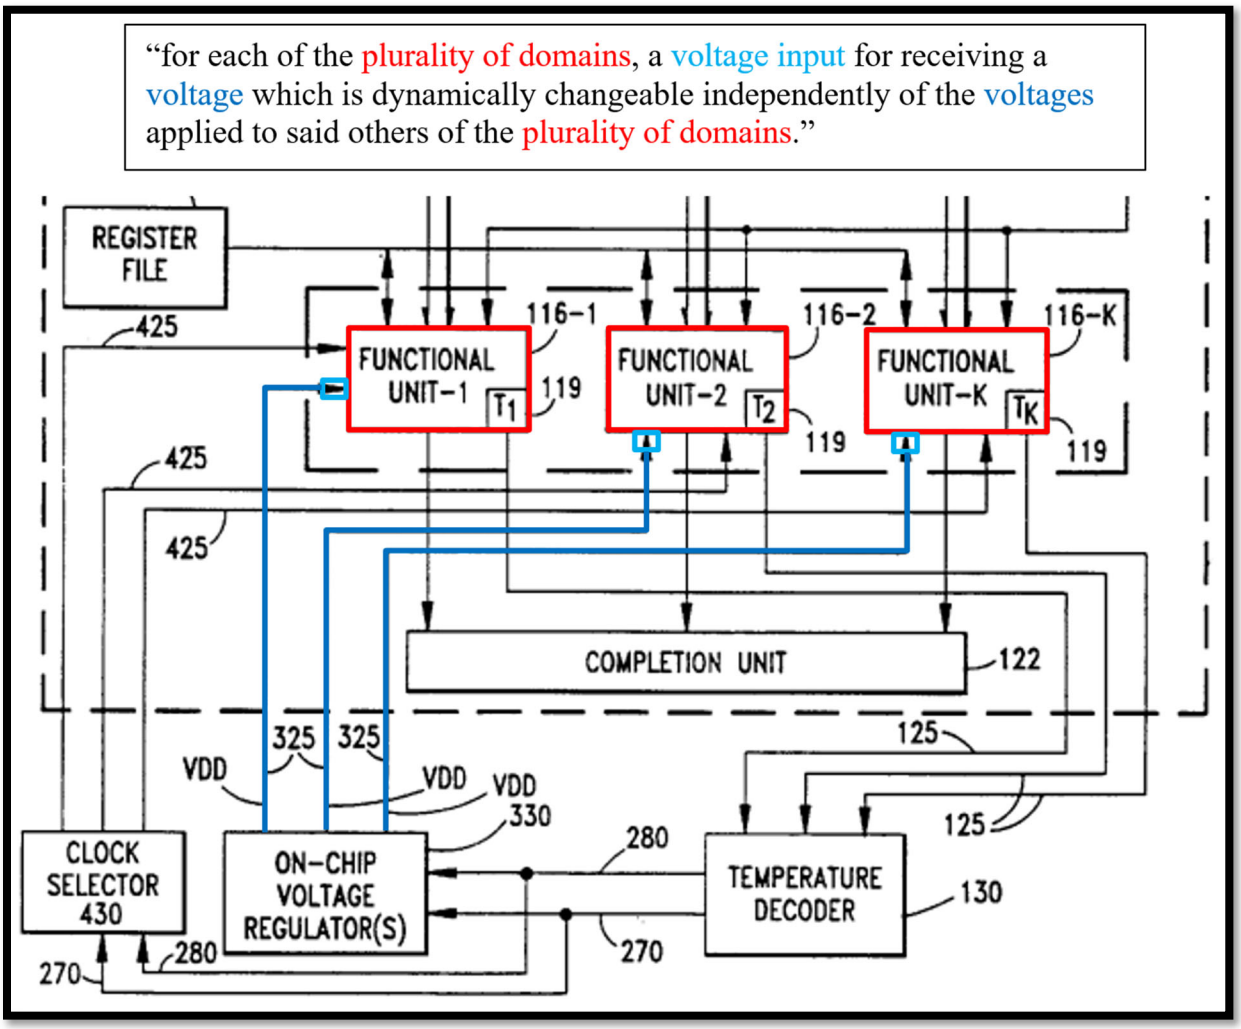

Figure 1 of *Georgiou* (Excerpted)

Ex.-1008, p.3 (excerpted); Ex.-1002, ¶202. Supply voltage 325 provides separate voltages to the inputs of functional unit 116-1, 116-2,…116-K. *Georgiou* teaches "the output 325 of the adjustable voltage regulator 330 is applied to the **voltage supply** of the functional unit which provided the thermal signal 125." Ex.-1008, 4:38-41 Ex.-1002, ¶¶202-204. To the extent that that *Georgiou's* voltage supply is

found not to correspond to a "voltage input," it was known in the art for microprocessors to use input pins to receive a voltage. Ex.-1002, ¶¶203-204; *see* Ex.-1013, p.138, 175; *see also* Ex.-1011, p.6. A POSITA would have applied known techniques of using voltage inputs (e.g., an input pin) to known devices (e.g., *Georgiou's* microprocessor) to yield the predictable result of receiving a voltage from a respective voltage output with a reasonable expectation of success. Ex.-1002, ¶¶203-204; *see KSR*, 550 U.S. at 415-421.

Current/clock control signal 270, as seen in Figure 1, is coupled to "adjustable on-chip voltage regulators 330 to reduce the supply voltage 325." Ex.-1008, 4:35-36. The output (e.g., supply voltage 325) "is applied to the **voltage supply of the functional unit** which provided the thermal signal 125." Ex.-1008, 4:39-42; Ex.-1002, ¶¶202-204. "Thus, the **modulated voltage** and frequency **are applied to the functional unit to dynamically control heating** while maintaining synchronization and reliability." Ex.-1008, 10:48-51; Ex.-1002, ¶¶202-204. Each functional unit is provided a separate voltage that changes depending on the thermal feedback received from that functional unit. Ex.-1008, 5:5-7; Ex.-1002, ¶¶202-204. In this way, the voltages supplied are independent from one another, and are dynamically changeable (e.g., based on thermal feedback of each functional unit). Ex.-1002, ¶¶202-204.

2. *Claim 3*

Georgiou teaches this element. Ex.-1002, ¶¶205-206. For example, *Georgiou* includes functional units 116-1, 116-2, **through 116-K**, thus describing numerous functional units (e.g., an amount larger than four as functional units 116-1 and 116-2 are shown and a functional unit 116-K is provided with variable K indicating more than three functional units as, otherwise, the variable K would be "3"). Ex.-1002, ¶¶205-206. *Georgiou* provides at least four separate embodiments of the functional units, including "fixed-point units, floating point units, load/store units, and branch units found in typical general purpose and superscalar microprocessors." Ex.-1008, 3:58-61; *see supra* Section X(E)(1)(b). Ex.-1002, ¶¶205-206.

3. *Claim 8*

a. Element 8(a)

To the extent the preamble is limiting, *Georgiou* teaches *a method of operating a microprocessor*. *See supra* Section X(E)(1)(a); Ex.-1002, ¶207. *Georgiou* states the invention is directed to a "method for cooperatively varying the voltage and operating frequency of a circuit to control heating while maintaining synchronization." Ex.-1008, 3:19-21; Ex.-1002, ¶207.

b. Element 8(b)

As discussed for claim 1, *Georgiou* provides a plurality of functional units in its microprocessor. *See supra* Section X(E)(1)(b); Ex.-1002, ¶208.

c. Element 8(c)

As discussed for claim 1, *Georgiou's* functional units operate using an independent clock output by a respective multiplexer, where each functional unit is operated using its own respective clock. See *supra* Section X(E)(1)(c); Ex.-1002, ¶209.

d. Element 8(d)

As discussed for claim 1, *Georgiou's* functional units operate using an independent voltage, where each functional unit is operated using its own respective voltage generated based on each functional unit's respective thermal feedback. *See supra* Section X(E)(1)(d); Ex.-1002, ¶210.

e. Element 8(e)

This element would have been obvious to a POSITA in view of *Georgiou* since the microprocessor described by *Georgiou* operates its functional units using a plurality of respective independent frequencies at the same time. *See supra* Section X(E)(1)(c); Ex.-1002, ¶¶211-219. *Georgiou* discloses that microprocessor components such as the on-chip voltage regulators and "clock selector 430 (voltage-frequency pair) will **modulate the voltage 325 and corresponding clock frequency 425 respectively** to control the heat generated by the **associated functional unit while maintaining synchronization and reliability**." Ex.-1008, 9:4-11; Ex.-1002, ¶¶213. A POSITA would have recognized that the functional units of *Georgiou* operate asynchronously since they have different clocks while each

functional unit operates synchronously to allow the functional units to perform their operations. Ex.-1002, ¶¶211-219, 53-56.

Georgiou's functional units are intra-unit synchronized. Ex.-1002, ¶¶212-214, 219. Each functional unit is operated based its respective frequency and voltage provided based on its respective thermal feedback. *See supra* Sections X(E)(3)(c)-X(E)(3)(d); Ex.-1002, ¶¶212-214. Components of a given functional unit operate synchronously using the same respective frequency and voltage provided to the given functional unit. Ex.-1002, ¶¶213-214, ¶219. A POSITA would have understood that such operation corresponds to intra-unit synchronized operation. Intra-domain synchronized operation was also known in the art. Ex.-1002, ¶¶212-214, 53-56; *see* Ex.-1022, p.1; *see also* Ex.-1007*,* 4:21-23. It would have been obvious to a POSITA to apply known intra-unit synchronization techniques to improve similar devices (e.g., *Georgou's* functional units) in the same way as described in the art to yield predictable results to maintain synchronization and reliability. Ex.-1002, ¶¶212-214, 219; *see KSR*, 550 U.S. at 415-421.

Georgiou's functional units are also inter-unit synchronized. Ex.-1002, ¶¶212-213, 215, 217-219. *Georgiou* teaches "[u]sing thermal sensors for feedback, the voltage swing and operating frequency of the circuit are cooperatively varied (also called voltage-frequency pairs) to reduce the power dissipation **without compromising** reliability or **system synchronization**." Ex.-1008, 2:21-26; Ex.-

1002, ¶215. Such inter-unit or system synchronization was also well known in the art. Ex.-1002, ¶¶213, 215-219; *see also* Ex.-1022, p.1 ("Obviously, modules cannot be synchronized if they all use different clocks, so we use an **asynchronous protocol** to communicate between them"); *see also* Ex.-1023, 6:18-21 ("**synchronization schema** must be implemented when there is a need to communicate data between PEs (or with other system elements) that operate at different frequencies.") To the extent *Georgiou* is found not to describe how its inter-unit synchronization is implemented, a POSITA would have known to apply known techniques for inter-unit synchronization (e.g., an asynchronous protocol, a synchronization schema, etc.) to improve the system synchronization of *Georgiou's* functional units to yield predictable results. Ex.-1002, ¶¶213, 215-219, 53-56; *see KSR*, 550 U.S. at 415-421.

f. Element 8(f)

As discussed for elements 1(b) and 1(c), *Georgiou's* functional units operate using an independent clock and voltage, where each functional unit's clock and voltage is determined based on its operations independent of other functional units. See *supra* Sections X(E)(1)(b)-X(E)(1)(c); Ex.-1002, ¶220.

4. *Claim 12*

As discussed for claim 3, *Georgiou* discloses at least four loads *See supra* Section X(E)(2); Ex.-1002, ¶221.

F. Ground 6: Claims 4 and 9 are obvious over *Georgiou* in view of *Halepete*

1. *Claim 4*

A POSITA would have found it obvious to combine *Georgiou* with *Halepete* for similar reasons for combining *Shenai* with *Halepete*. *See supra* Section X(A)(3)); Ex.-1002, ¶¶222-224. *Shenai* and *Georgiou* are both directed to microprocessors with a plurality of functional units. *See supra* Sections X(A)(1)(b), X(E)(1)(b). Ex.-1002, ¶223. *Georgiou* also teaches critical timing paths. Ex.-1002, ¶223; Ex.-1008, 1:49-54 ("For example, a digital chip is typically characterized by a critical timing path. Reducing the supply voltage will cause the propagation of a signal through the critical timing path to slow until a malfunction occurs."). *Georgiou* varies the voltage to maintain these critical timing paths. *Id.*, 2:28-29. Ex.-1002, ¶223

A POSITA would have been motivated to combine *Georgiou's* multi-functional-unit microprocessor with *Halepete's* control software capable of determining necessary clock frequency and voltage levels (and consequently adjusting the clock frequency and voltage level) as the combination would have been nothing more than combining prior art elements according to known methods to yield predictable results, and thus there is a reasonable expectation of success, etc. Ex.-1002, ¶¶222-224; *see KSR*, 550 U.S. at 415-421; *see also supra* Section X(B)(1) (discussing why a POSITA would have combined *Halepete* with a similar multi-load microprocessor)

2. *Claim 9*

a. Element 9(a)

Georgiou renders obvious the method of claim 8. *See supra* Section X(E)(3); Ex.-1002, ¶225.

b. Element 9(b)

As discussed for claim 4, *Georgiou* in view of *Halepete* renders obvious monitoring critical timing paths to identify which operating characteristics can be reduced based critical timing paths. *See supra* Section X(F)(1); Ex.-1002, ¶226.

c. Element 9(c)

As discussed for claim 4, *Georgiou* in view of *Halepete* renders obvious reducing a clock frequency and voltage for a functional unit based on its operational needs. *See supra* Section X(F)(1); Ex.-1002, ¶227.

G. Ground 7: Claims 5 and 10 are obvious over *Georgiou* in view of *Kranich*

1. *Claim 5*

Kranich discloses "a plurality of processing units[.]" Ex.-1009, 9:4-5; Ex.-1002, ¶¶228-234. A POSITA would have recognized the teachings of *Kranich* to apply to functional units, such as functional units within a single processor. Ex.-1002, ¶¶229-235. Figure 3A, as shown annotated below, illustrates "interconnection between two processors 12A-12B." Ex.-1009, 10:6-7; Ex.-1002, ¶¶229-230.

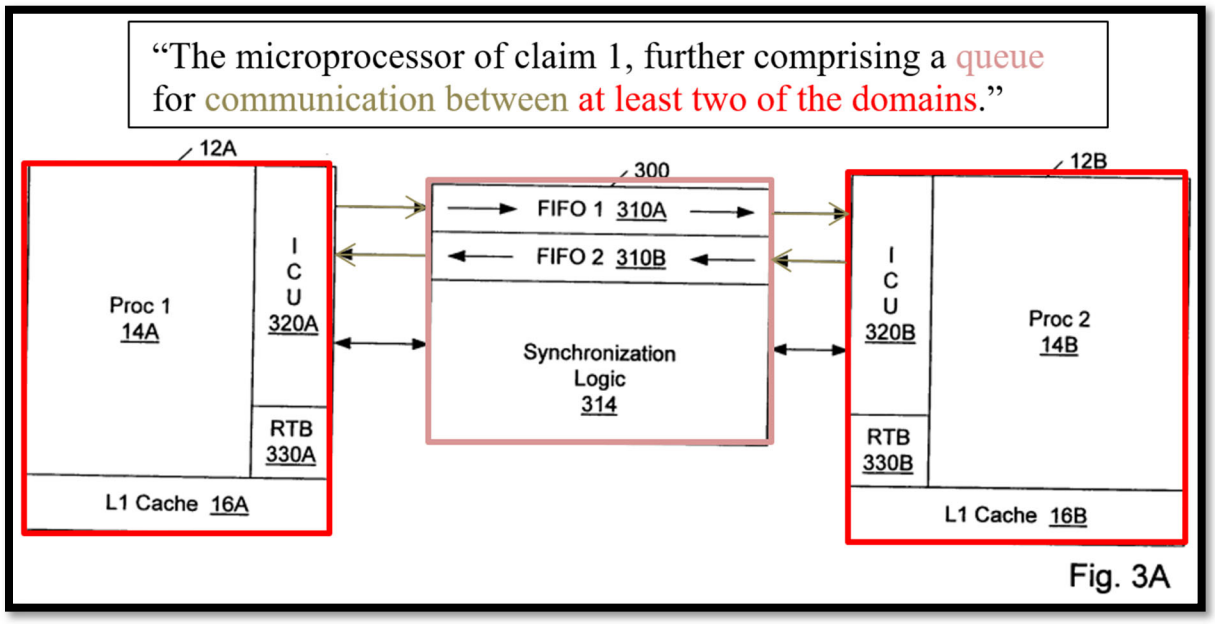

Figure 3A of *Kranich*

Ex.-1009, p.26; Ex.-1002, ¶229. "[T]hread control device 300" includes "first-in-first-out (hereinafter FIFO) buffers 310A-310B, and synchronization logic 314" to facilitate intercommunication between processors 12A and 12B. Ex.-1009, 10:6-11; Ex.-1002, ¶230. The interprocessor communication units (ICU) 320A and/or 320B communicates with FIFO buffers 310A and 310B. The ICUs of "each processor 12 are coupled to thread control device 300 which facilitates communication between processors 12A and 12B." Ex.-1009, 10:12-13; Ex.-1002, ¶230.

Similar to *Sherburne*, it would have been obvious to a POSITA to include a well-known method (e.g., implementing FIFO based buffers) as described in *Kranich* (and in *Sherburne*) to facilitate the communication between the functional

units of *Georgiou*. Ex.-1002, ¶¶228-234; *see supra* Section X(C)(1). Since *Georgiou*'s microprocessor is maximizing the usage of each functional unit with the power provided, the functional units would have benefited from communicating necessary voltages and/or frequencies to one another to efficiently complete tasks and reduce any wait times (e.g., slack). Such FIFO buffers were used in the art. Ex.-1002, ¶¶228-234, 41-44.

Georgiou is directed determining operating parameters "of a functional unit in a multi-functional unit processor such as a superscalar microprocessor." Ex.-1008, 3:48-53; Ex.-1002, ¶232. *Kranich* is directed to a "processor architecture containing multiple closely coupled processors in a form of symmetric multiprocessing system." Ex.-1009, p.1, Abstract; Ex.-1002, ¶232. Therefore, *Georgiou* and *Kranich* are in the same field of endeavor of operating a plurality of units of a processing system. Ex.-1002, ¶232. A POSITA would have known to combine *Georgiou's* functional units with known buffers (e.g., as described in *Kranich*) according to known methods (e.g., FIFO implementations) to yield predictable results (e.g., inter-unit communication) with a reasonable expectation of success. *Id.*, ¶¶228-234; *see KSR*, 550 U.S. at 415-421.

2. *Claim 10*

As discussed for claim 5, *Georgiou* in view of *Kranich* provides a buffer for communication between two functional units. *See supra* Section X(G)(1); Ex.-1002, ¶¶235.

XI. ARGUMENTS FOR DISCRETIONARY DENIAL SHOULD BE REJECTED

A. Section 325(d) Is Inapplicable Because the Asserted Art Was Never Evaluated During Examination.

The Board should not deny institution under § 325(d) because the art asserted here was not identified during prosecution or before the Examiner, and is not cumulative of art that was. As set forth below, the Examiner either (1) was not presented with the same or substantially the same art or arguments as Petitioner's, or (2) materially erred in allowing the Challenged Claims. *Advanced Bionics, LLC v. Med-El Elektromedizinische Gerate GmbH*, IPR2019-01469, Paper 6 at 8 (P.T.A.B. Feb. 13, 2020) (citing *Becton, Dickinson, & Co. v. B. Braun Melsungen AG*, IPR2017-01586, Paper 8 (P.T.A.B. Dec. 15, 2017).

Advance Bionics Analysis Step 1 - Neither "the same [nor] substantially the same" art or arguments were previously presented to the Office during prosecution of the Challenged Claims. *Advanced Bionics*, IPR2019-01469, Paper 6 at 8. *Shenai, Halepete, Sherburne, Traylor, Georgiou,* and *Karnich* were never cited during prosecution of the '443 Patent, let alone considered by the Examiner or used in any rejection. *See* Ex.-1003 generally. These references are also not substantially the

same or cumulative of references considered during examination. During Examination, the issued claims were allowed in a first action Notice of Allowance.

Advance Bionics Analysis Step 2 - To the extent the Board disagrees and determines that the Advance Bionics Analysis Step 1 does not favor institution, discretionary denial still is not warranted because the Examiner must have necessarily overlooked disclosures of the art that was examined, constituting material error. *Advanced Bionics*, IPR2019-01469, Paper 6, 10 (listing silence as evidence of error). As stated above in detail, both *Shenai* and *Georgiou* alone or in combination with *Halepete*, *Sherburne*, *Traylor*, *or Kranich,* render obvious each of the Challenged Claims. To the extent any reference that was examined could be considered cumulative of these prior art references, the Examiner should have rejected the Challenged Claims at least under Section 103, and maintained the rejection(s).

Accordingly, the Board should reach the merits of this petition, and institute review of all Challenged Claims, especially in light of the accompanying expert testimony, which was not before the Office during prosecution.

B. Any Secondary Considerations Cannot Overcome the Strong Evidence of Obviousness

The Board, at the institution phase, has repeatedly held that evidence of secondary considerations presented by the Patent Owner should be addressed in a

trial where the parties may develop, and the Board may consider, a full record. *See*, e.g., *Tristar Products, Inc. v. Choon's Design, LLC*, IPR2015-00883, Paper 6 at 26 (P.T.A.B. Mar. 9, 2016). That is the appropriate course here, especially given that "the inventions represent[] no more than 'the predictable use of prior art elements according to their established functions,'" and, thus any "secondary considerations are inadequate to establish nonobviousness as a matter of law." *Wyers v. Master Lock Co.,* 616 F.3d 1231, 1246 (Fed. Cir. 2010) (quoting *KSR*, 550 U.S. at 417. Institution is Proper Under Section 314(a)

The merits of this Petition are strong, which alone warrants institution. Petitioner also hereby stipulates that, if institution is granted, Petitioner will not raise in the co-pending litigation any defense based on the same grounds raised, or that could have been raised, in this Petition. This stipulation precludes discretionary denial under Section 314(a). Director Vidal, Memorandum, "Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation," at 4-5, 7-8 (June 21, 2022) ("the PTAB will not discretionarily deny institution if Petitioner presents a *Sotera* stipulation.").

XII. PRESERVATION OF RIGHTS

Petitioner reserves all rights to any benefit of changes resulting from challenges based on the recent Supreme Court decision in *Loper Bright*. *Loper Bright Enters. v. Raimondo*, No. 22-451, slip op. at 35 (U.S. June 28, 2024).

XIII. CONCLUSION

For the reasons above, Petitioner requests institution of IPR of the Challenged Claims based on all grounds.

Respectfully submitted,

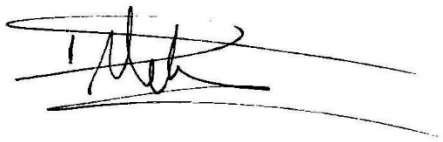

Dated: August 2, 2024

Dinesh N. Melwani (Reg. No. 60,670)
Ankit Aggarwal (Reg. No. 67,882)
William Uhr (Reg. No. 71,282)
Bookoff McAndrews, PLLC
2000 Pennsylvania Avenue NW, Suite 4001
Washington, D.C. 20006
202-808-3497

Attorneys for Petitioner
Lenovo (United States) Inc.

CERTIFICATE OF COMPLIANCE

Pursuant to 37 C.F.R. § 42.24(d), the undersigned certifies that the foregoing Petition for *Inter Partes* Review of U.S. Patent No. 7,089,443 B2 contains, as measured by the word-processing system used to prepare this paper, 13,362 words. This word count does not include the items excluded by 37 C.F.R. § 42.24.

Dated: August 2, 2024

Respectfully submitted,

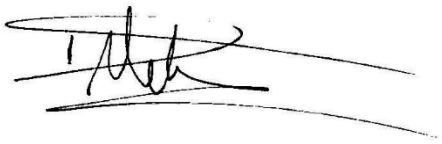

By: _____
Dinesh N. Melwani (Reg. No. 60,670)
Counsel for Petitioner

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. § 42.6(e) and 37 C.F.R. § 42.105(a), I hereby certify that on August 2, 2024, I caused a true and correct copy of the foregoing "PETITION FOR INTER PARTES REVIEW OF U.S. PATENT NO. 7,089,443 B2" and supporting exhibits to be served via Federal Express on the Patent Owner at the following correspondence address of record and/or likely to effect service:

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

Prince Lobel Tye LLP
One International Place
Suite 3700
Boston, MA 02110

Dated: August 2, 2024

Respectfully submitted,

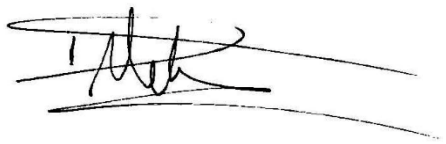

By: _____
Dinesh N. Melwani (Reg. No. 60,670)
Counsel for Petitioner