

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO., LTD.,
SAMSUNG ELECTRONICS AMERICA, INC., AND
QUALCOMM INCORPORATED,
Petitioners

v.

DAEDALUS PRIME LLC,
Patent Owner.

Case No. IPR2023-00617
U.S. Patent No. 8,898,494

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 8,898,494**

TABLE OF CONTENTS

	Page
I. INTRODUCTION.....	1
II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8	2
III. FEE AUTHORIZATION.....	5
IV. GROUNDS FOR STANDING	5
V. RELIEF REQUESTED	5
VI. THE CHALLENGED PATENT.....	6
VII. PROSECUTION HISTORY OF THE '494 PATENT	7
VIII. LEVEL OF ORDINARY SKILL IN THE ART	8
IX. PRIORITY DATE.....	8
X. CLAIM CONSTRUCTION	9
XI. THE PRIOR ART	9
A. Bose (Ex-1005).....	9
B. Fukatsu (Ex-1006).....	12
C. Naffziger (Ex-1007)	14
D. White (Ex-1008).....	16
E. Simeral (Ex-1009).....	18
XII. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS 20	
A. Ground 1: Claims 1 and 12 Are Rendered Obvious by Bose	20
1. Independent Claim 1	20
a. Element 1[pre]: A processor comprising:.....	20
b. Element 1[a]: an integrated circuit including:	21
c. Element 1[b]: a first core;.....	22
d. Element 1[c]: a cache memory;	22
e. Element 1[d]: a communication bus coupled to the first core, the communication bus to connect the first core and the cache memory;.....	23

TABLE OF CONTENTS

(continued)

	Page
f. Element 1[e]: a core workload monitor configured to determine a core workload for the first core; 25 g. Element 1[f]: a bus workload monitor configured to determine a bus workload for the communication bus; and 27 h. Element 1[g]: balancing control adapted to receive the bus workload from the bus workload monitor and to dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit and a comparison between the bus workload and a bus workload threshold, the power limit corresponding to a maximum thermal dissipation capacity for the integrated circuit, wherein a power consumption of one of the first core and the communication bus is to be reduced, the power consumption reduction to be limited to maintain operation of the one of the first core and the communication bus above a low limit. 29	
2. Dependent Claim 12: The processor of claim 1, wherein the balancing control is to receive a hint regarding a future workload and to dynamically tune the power allocation further based on the hint. 37	
B. Ground 2: Claims 3, 12, 14, and 15 Are Rendered Obvious by Bose Combined with Fukatsu 38	
1. A POSITA Would Have Been Motivated to Combine Bose and Fukatsu 38 2. Dependent Claim 3: The processor of claim 1, wherein the balancing control, in response to the integrated circuit operating at the power limit, being adapted to: increase frequency for the communication bus and cap frequency for the first core in response to the bus workload being greater than the bus workload threshold; and decrease frequency for the communication bus and increase the	

TABLE OF CONTENTS

(continued)

	Page
frequency for the first core in response to the bus workload being less than the bus workload threshold	42
3. Dependent Claim 12	45
4. Dependent Claim 14: The processor of claim 3, wherein the balancing control is to increase the communication bus frequency by a different amount than the first core frequency decrease.....	46
5. Dependent Claim 15: The processor of claim 3, wherein the balancing control is to further increase frequency for a second core, the second core and the communication bus having power planes tied together.	47
C. Ground 3: Claims 1, 3, 14, and 15 Are Rendered Obvious by Bose Combined with Naffziger.....	49
1. A POSITA Would Have Been Motivated to Combine Bose and Naffziger.....	49
2. Independent Claim 1	53
a. Elements 1[pre] through 1[f]: A processor comprising:	53
b. Element 1[g]: balancing control... ..	53
3. Dependent Claim 3	56
4. Dependent Claim 12	59
5. Dependent Claim 14	59
6. Dependent Claim 15	60
D. Ground 4: Claims 1, 3, 14, and 15 Are Rendered Obvious by White Combined with Simeral.....	62
1. A POSITA Would Have Been Motivated to Combine White and Simeral.....	62
2. Independent Claim 1	66
a. Element 1[pre]: A processor comprising:.....	66
b. Element 1[a]: an integrated circuit including:	67

TABLE OF CONTENTS

(continued)

	Page
c. Element 1[b]: a first core;.....	68
d. Element 1[c]: a cache memory;.....	69
e. Element 1[d]: a communication bus...;.....	70
f. Element 1[e]: a core workload monitor...;.....	72
g. Element 1[f]: a bus workload monitor...;	73
h. Element 1[g]: balancing control...;	76
3. Dependent Claim 3	83
4. Dependent Claim 14	86
5. Dependent Claim 15	87
XIII. THERE IS NO BASIS FOR ANY DISCRETIONARY DENIAL	
89	
XIV. CONCLUSION	90

LIST OF EXHIBITS¹

Ex-1001	U.S. Patent No. 8,898,494 (“’494 Patent”)
Ex-1002	Declaration of Dr. Trevor Mudge
Ex-1003	Curriculum Vitae of Dr. Trevor Mudge
Ex-1004	Prosecution History of U.S. Patent No. 8,898,494
Ex-1005	U.S. Patent No. 7,421,601, issued on September 2, 2008 (“Bose”)
Ex-1006	U.S. Patent Publication No. 2008/0005607 published on January 3, 2008 (“Fukatsu”)
Ex-1007	U.S. Patent No. 8,510,582, filed on July 21, 2010 (“Naffziger”)
Ex-1008	U.S. Patent No. 7,263,457, issued on August 28, 2007 (“White”)
Ex-1009	U.S. Patent Publication No. 2009/0150689, published on June 11, 2009 (“Simeral”)
Ex-1010	INTENTIONALLY LEFT BLANK
Ex-1011	INTENTIONALLY LEFT BLANK
Ex-1012	INTENTIONALLY LEFT BLANK
Ex-1013	INTENTIONALLY LEFT BLANK
Ex-1014	INTENTIONALLY LEFT BLANK
Ex-1015	INTENTIONALLY LEFT BLANK
Ex-1016	INTENTIONALLY LEFT BLANK
Ex-1017	INTENTIONALLY LEFT BLANK
Ex-1018	INTENTIONALLY LEFT BLANK
Ex-1019	INTENTIONALLY LEFT BLANK

¹ Four-digit pin citations that begin with 0 are to the branded numbers added by Samsung in the bottom right corner of the exhibits. All other pin citations are to original page, column, paragraph, or line numbers.

Ex-1020	INTENTIONALLY LEFT BLANK
Ex-1021	INTENTIONALLY LEFT BLANK
Ex-1022	U.S. Patent Publication No. 2011/0022356, published on January 27, 2011 (“Nussbaum”)
Ex-1023	U.S. Patent Publication No. 2013/0151869, published on June 13, 2013 (“Steinman”)
Ex-1024	U.S. Patent Publication No. 2012/0079290, published on March 29, 2012 (“Kumar”)
Ex-1025	U.S. Patent Publication No. 2009/0204831, published on August 13, 2009 (“Cousson”)
Ex-1026	U.S. Patent No. 7,949,817, filed on July 21, 2010 (“Sakarda”)
Ex-1027	U.S. Patent Publication No. 2010/0115304, published on May 6, 2010 (“Finkelstein”)
Ex-1028	Bharadwaj Amrutur, Nandish Mehta, Satyam Dwivedi, and Ajit Gupte, <i>Adaptative Techniques to Reduce Power in Digital Circuits</i> (July 2011)
Ex-1029	U.S. Patent Publication No. 2007/0250736 to Gaskins (“Gaskins”)
Ex-1030	Charles Lefurgy et al., <i>Energy Management for Commercial Servers</i> , Computer 39 (Dec. 2003)
Ex-1031	Euiseong Seo et al., <i>Energy Efficient Scheduling of Real-Time Tasks on Multicore Processors</i> , 19 IEEE Transactions on Parallel and Distributed Systems 1540 (Nov. 2008)
Ex-1032	Stefanos Kaxiras and Margaret Martonosi, <i>Computer Architecture Techniques for Power-Efficiency</i> , in <i>Synthesis Lectures on Computer Architecture</i> #4 (Morgan & Claypool 2008)
Ex-1033	Vasanth Venkatachalam and Michael Franz, <i>Power Reduction Techniques For Microprocessor Systems</i> , 37 ACM Computing Surveys 195 (2005)
Ex-1034	U.S. Patent No. 7,263,457 to White (“White”)

Ex-1035	U.S. Patent No. 8,438,416 to Kocev (“Kocev”)
Ex-1036	U.S. Patent No. 9,086,883 to Thomson (“Thomson”)
Ex-1037	U.S. Published Patent Application No. 2010/0073068 (“Cho”)
Ex-1038	U.S. Patent No. 8,397,090 to Gunther (“Gunther”)

I. INTRODUCTION

Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., and Qualcomm, Inc. (collectively, “Petitioner”) request *inter partes* review (“IPR”) of Claims 1, 3, 12, 14, and 15 of U.S. Patent No. 8,898,494 (“the ’494 Patent”) (Ex-1001), currently assigned to Daedalus Prime LLC (“PO”).

The ’494 Patent relates to “energy efficiency and energy conservation in integrated circuits,... [and] to balancing power and performance between processing elements and/or a communication bus.” Ex-1001, 1:12-16.

The prior art in this Petition demonstrates that Claims 1, 3, 12, 14, and 15 of the ’494 Patent involved well-known features and predictable design choices. This prior art also shows that a person of ordinary skill in the art (“POSITA”) would have used these components to address well-known issues relating to balancing performance and power between processing elements, such as processing cores and communication busses. Thus, Claims 1, 3, 12, 14, and 15 of the ’494 Patent would have been obvious over the prior art.

This Petition presents prior art and grounds that the examiner did not consider during prosecution. These grounds are more than reasonably likely to prevail, and this Petition should be granted.

II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8

Real Parties-in-Interest: Petitioner identifies the following real parties-in-interest: Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., and Co-Petitioner Qualcomm Incorporated, and Qualcomm Technologies, Inc.

Petitioner Qualcomm Incorporated further identifies, without conceding any such party is a Real Party-in-Interest, Mazda Motor Corporation, Mazda North American Operations, Mazda Motor of America, Inc., Mercedes-Benz Group AG, Mercedes-Benz AG, Mercedes-Benz USA, LLC, and Visteon Corporation as parties that Patent Owner has accused of infringing the '494 Patent based on incorporation of Qualcomm products in *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, 1:22-cv-01109 (D. Del.) and *Certain Semiconductors and Devices and Products Containing the Same, Including Printed Circuit Boards, Automotive Parts, and Automobiles*, ITC Inv. No. 337-TA-1332 (Violation).

Related Matters: Patent Owner has asserted the '494 Patent against Petitioner in (1) *Daedalus Prime LLC v. Arrow Electronics, Inc. et al.*, 1:22-cv-01107 (D. Del.); (2) *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, 1:22-cv-01109 (D. Del.); (3) *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, 1:22-cv-01108 (D. Del.); (4) *Daedalus Prime LLC v. Samsung Electronics Co., Ltd. et al.*, 2:22-cv-00352 (E.D. Tex.); (5) *Certain Integrated Circuits, Mobile Devices Containing the Same, and Components Thereof*, U.S. International Trade

Commission (“ITC”) Inv. No. 337-TA-1335 (Violation); and (6) *Certain Semiconductors and Devices and Products Containing the Same, Including Printed Circuit Boards, Automotive Parts, and Automobiles*, ITC Inv. No. 337-TA-1332 (Violation).

Lead and Back-Up Counsel:

- Lead Counsel:

William M. Fink (Reg. No. 72,332)
O’Melveny & Myers LLP
1625 Eye Street, NW
Washington, DC 20006
Telephone: (202) 383-5300
Fax: (202) 383-5414
E-Mail: tfink@omm.com

- Backup Counsel:

Ben Haber (Reg. No. 67,129)
Nicholas J. Whilt (Reg. No. 72,081)
Brian Cook (Reg. No. 59,356)
O’Melveny & Myers LLP
400 South Hope Street, 18th Floor
Los Angeles, CA 90071
Telephone: (213) 430-6000
Fax: (213) 430-6407
E-Mail: bhaber@omm.com
E-Mail: nwhilt@omm.com
E-Mail: bcook@omm.com
Counsel for Petitioners Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc.

Daniel Leventhal (Reg. No. 59,576)
Richard Zembek (Reg. No. 43,306)
Darren Smith (Reg. No. 64,261)

Lu Ling (Reg. No. 78,121)
Norton Rose Fulbright US LLP
1301 McKinney, Suite 5100
Houston, TX 77010
Telephone: (713) 651-5151
Fax: (713) 651-5246
E-Mail: daniel.leventhal@nortonrosefulbright.com
E-Mail: richard.zembek@nortonrosefulbright.com
E-Mail: darren.smith@nortonrosefulbright.com
E-Mail: lu.ling@nortonrosefulbright.com
Counsel for Petitioner Qualcomm Incorporated

Eagle H. Robinson (Reg. No. 61,361)
Norton Rose Fulbright
98 San Jacinto Blvd., Ste. 1100
Austin, TX 78701
Telephone: (512) 474-5201
Fax: (512) 536-4598
E-Mail: eagle.robinson@nortonrosefulbright.com
Counsel for Petitioner Qualcomm Incorporated

Service Information: Petitioner consents to electronic service by email to

the following addresses:

- tfink@omm.com
- bhaber@omm.com
- nwhilt@omm.com
- bcook@omm.com
- daniel.leventhal@nortonrosefulbright.com
- richard.zembek@nortonrosefulbright.com
- darren.smith@nortonrosefulbright.com
- lu.ling@nortonrosefulbright.com
- eagle.robinson@nortonrosefulbright.com

III. FEE AUTHORIZATION

The PTO is authorized to charge any fees due during this proceeding to Deposit Account No. LA500639.

IV. GROUNDS FOR STANDING

Petitioner certifies that the '494 Patent is available for review, and Petitioner is not barred or estopped from requesting review.


V. RELIEF REQUESTED

Petitioner requests review and cancellation of Claims 1, 3, 12, 14, and 15 as unpatentable based on the following grounds, supported by a declaration from Dr. Trevor Mudge. Ex-1002 ¶¶1-30, 38-58.

Ground	Summary
1	Claims 1 and 12 are obvious over Bose (Ex-1005)
2	Claims 3, 12, 14, and 15 are obvious over Bose (Ex-1005) combined with Fukatsu (Ex-1006)
3	Claims 1, 3, 12, 14, and 15 are obvious over Bose (Ex-1005) combined with Naffziger (Ex-1007)
4	Claims 1, 3, 14, and 15 are obvious over White (Ex-1008) combined with Simeral (Ex-1009)

VI. THE CHALLENGED PATENT²

The '494 Patent purports to describe a processor for efficiently balancing performance and power between one or more cores and a communication bus based on measured workloads to alleviate bottlenecks while operating within thermal limits. Ex-1001, Abstract, 15:4-21. Figure 9 illustrates the '494 Patent's claimed processor:



Ex-1001, Fig. 9.³

Integrated circuit 900 (light blue) comprises multiple cores (e.g., CPU cores 905) (red and blue), cache or memory 960 (dark red), communication bus 950 (green), and power management service 920 (yellow). Ex-1001, Fig. 9, 2:21-24,

² Ex-1002 ¶¶31-35.

³ All annotations and emphasis have been added unless otherwise noted.

15:48-52, 16:10-41, 16:57-59, 17:11-15. Each core has an associated workload monitor. *Id.* at Fig. 9, 9:47-52, 16:14-56. Cache or memory 960 can be any known cache memory. *Id.* at Figs. 1, 9, 5:43-61, 15:59-67. Communication bus 950 (green) connects cores to shared resources such as cache or memory 960. Ex-1001, Fig. 9, 15:59-67. The communication bus includes any known bus or interconnect and is associated with its own workload monitor. *Id.* at Fig. 9, 15:38-58, 16:18-56.

Power management service 920 (yellow) balances power among the components based on their workloads. Ex-1001, Fig. 9, 16:57-59. For example, the '494 Patent discloses increasing the communication bus's frequency/power if its workload is too high, alleviating a bottleneck between the cores and cache, and counterbalancing that increase by decreasing the frequency/power of the cores. Ex-1001, 17:4-41; *see also id.*, Fig. 10, 18:35-19:46. In other examples, the core's power/frequency may be increased and be counterbalanced by a decrease in the power/frequency of another element to maintain operation of the integrated circuit below its thermal limit. *Id.* at 19:3-63.

VII. PROSECUTION HISTORY OF THE '494 PATENT⁴

The '494 Patent's file history is submitted as Ex-1004. The '494 Patent was filed on February 16, 2012, as a continuation-in-part of application No. 13/327,670,

⁴ Ex-1002 ¶¶36-37.

filed on December 15, 2011. After several office actions and responses, a notice of allowance issued on July 30, 2014. Ex-1004, 249, 257.

To the best of Petitioner's knowledge, none of Bose, Fukatsu, Naffziger, White, or Simeral was considered in the prosecution of the application that led to the '494 Patent, or in its parent application.

VIII. LEVEL OF ORDINARY SKILL IN THE ART

One of ordinary skill in the art would have had a bachelor's degree in electrical engineering, computer engineering, computer science, or a related field, and at least two years of experience in the research, design, development, and/or testing of *semiconductors, processors, power management, and related firmware and software*, or the equivalent, with additional education substituting for experience and vice versa. Ex-1002 ¶¶22-26, 38-58.

IX. PRIORITY DATE

The earliest possible priority for the '494 Patent is December 15, 2011. For purposes of this Petition, Petitioner takes no position on the proper priority date for each claim of the '494 Patent but reserves its rights to challenge the '494 Patent's priority date in other proceedings, including in the ITC. Assuming this earliest date, all prior art references asserted in this petition are prior art.

X. CLAIM CONSTRUCTION

Petitioner interprets the claims of the '494 Patent according to 37 C.F.R. § 42.100(b). To resolve the grounds presented in this Petition, Petitioner does not believe that any term requires explicit construction. Ex-1002 ¶¶87.⁵

XI. THE PRIOR ART

A. Bose (Ex-1005)⁶

Bose, titled “Method and system for controlling power in a chip through a power-performance monitor and control unit,” is U.S. Patent No. 7,421,601, issued on September 2, 2008. Ex-1005, Cover. Bose is prior art under at least pre-AIA 35 U.S.C. § 102(b).


Bose discloses “controlling power and performance in a microprocessor system.” Ex-1005, Abstract. Bose does so by integrating a power-performance monitor and control unit (“PMCU”) into a circuit, “where the power of the chip (or system) is maintained within a programmable limit; the peak temperature across monitored regions of the chip (or system) is maintained to be below a specified limit; and the net throughput performance of the chip (or system) is targeted for

⁵ Claim construction procedures have recently begun in the related district court or ITC matters. Petitioner reserves the right to revisit and revise these constructions in district court and the ITC as necessary. Additionally, Petitioner will request leave to submit any claim construction order as soon as it becomes available, so that it is timely made of record in the proceeding and can be considered by the Board. 37 C.F.R. § 42.100(b).

⁶ Ex-1002 ¶¶59-64.

maximization, within the specified constraints of power and temperature.” Ex-1005, 4:19-29.

Bose discloses processor 100 (light blue) with multi-core chip 104 comprising local components of cores 112-1 (red) and 112-2 (blue), cache hierarchy 106 (dark red), interconnect 108 (green), and a global power management control unit (PMCU) 110 (yellow):



Ex-1005, Fig. 1; 5:61-67, 6:18-24.

Each core, as well as interconnect and shared cache, has its own local PMCU (denoted by small boxes 114, 116, 118-1, and 118-2) to monitor activity “such as cache miss rates, branch misprediction rates, instructions per cycle (IPCs),...power


usage and temperature usage.” Ex-1005, Fig. 1, 6:56-7:2. Bose utilizes a hierarchical power management scheme where each component “(a) monitor[s] events, activity levels and direct power and/or temperature levels or suitable proxies thereof, within its domain of authority; (b) provide[s] direct response mechanisms to manage power within its domain of action...; [and] (c) report[s] on a summary of monitored results and actions taken, to the next higher authority in the power management hierarchy.” *Id.* at Fig. 1, 6:24-32. Thus, “each local PMCU...reports upwards to the on-chip global PMCU.” *Id.* at 6:32-34.

Bose teaches that “[b]ased on thermal packaging, power supply, and other considerations, the global PMCU 110 is given an overall chip 104 power budget, and is responsible for assigning the budget to each core element 112-1, 112-2, 112-R, 106, and 108.” Ex-1005, 7:18-21. While staying within this power budget, the global PMCU also strives to maximize performance: “the global PMCU 110 will choose the power modes and assign power for each core element 112-1, 112-2, 112-R, 106, and 108 such that all elements can run at the highest possible performance within the budget.” *Id.* at Fig. 1, 7:24-28.

B. Fukatsu (Ex-1006)⁷

Fukatsu, titled “Method of controlling information processing device, information processing device, program, and program converting method,” is U.S. Patent Publication No. 2008/0005607, published on January 3, 2008. Ex-1006, Cover. Fukatsu is prior art under at least pre-AIA 35 U.S.C. § 102(b).


Fukatsu teaches that by predicting the hit rate of a core’s onboard cache, the needed bus frequency connecting to share resources can be determined in advance, thereby eliminating delays “to optimally achieve both a reduction in the electric power consumption and an improvement in the processing efficiency.” Ex-1006, Fig. 8, ¶32, 95. Fukatsu’s information processing device (light blue) includes CPU 101 (red), which includes an onboard cache memory 160, controlling unit 161 (yellow), bus 105 (green), memory 104 (dark red), as well as power unit 151, clock supplying unit 150, and storage medium 111:



⁷ Ex-1002 ¶¶65-70.

Ex-1006, Fig. 3, ¶¶52, 98.

Fukatsu also discloses a multiprocessor configuration with a plurality CPUs (red and blue):



Ex-1006, Fig. 12, ¶¶113-14. Bus 105 (green) connects the various CPUs to shared resources, such as Memory 104.


Controlling unit 161 (yellow) “predicts a hit rate of the cache memory 160; determines a frequency of the clock signal in accordance with the predicted hit rate; and controls the clock supplying unit 150 so as to change to the determined frequency. The controlling unit 161 also determines a power voltage in accordance with the determined frequency. Ex-1006, Fig. 3 ¶¶55-57. Fukatsu teaches varying the bus frequency and voltage inversely to the CPU frequency and voltage since the CPU is performing its processing when the cache hit rate is high and is retrieving information from shared resources such as the memory when the cache hit rate is low. *Id.*, Fig. 8 ¶¶6-7, 57.

C. Naffziger (Ex-1007)⁸

Naffziger, titled “Managing current and power in a computing system” is U.S. Patent No. 8,510,582, filed on July 21, 2010. Ex-1007, Cover. Naffziger is prior art under at least pre-AIA 35 U.S.C. § 102(a). Ex-1007, Cover.

Naffziger discloses efficient power transfer on a die where a “power manager reallocates power credits across the die when it detects an activity level of a given CU [computation unit] is below a given threshold” or “exceed[s] a given current limit.” Ex-1007, Abstract. Naffziger discloses processor 700 (light blue) including “one or more processor cores 112” (red and blue), “packet processing logic [or crossbar switch] 116” (green), “shared cache memory subsystem 118” (dark red), and “power management unit 130” (yellow):

⁸ Ex-1002 ¶¶71-74



Ex-1007, Fig. 8, 16:45-51, 17:66-18:4.

Naffziger discloses that power management unit 130 (yellow) associates a total number of power credits for the die and allocates a separate number of power credits to each CU. Ex-1007, 4:1-5. It allocates power credits by determin[ing] which computation units are able to donate power credits [and an amount to donate,]...determine which computation units are able to receive power credits [and an amount to receive.]” *Id.* at Fig. 4, 11:3-13. Naffziger further discloses allocating power credits between units by measuring and determining an appropriate power-performance state (“P-state”) for each unit. *Id.* at 6:55-7:16. The highest P-state

represents the highest power consumption and operational voltage, while the lowest P-state represents the lowest. *Id.*

D. White (Ex-1008)⁹

White, titled “System and method for operating components of an integrated circuit at independent frequencies and/or voltages” is U.S. Patent No. 7,263,457, issued on August 28, 2007. Ex-1008, Cover. White is prior art under at least pre-AIA 35 U.S.C. § 102(b).

White is directed to managing power and temperature of an integrated circuit and its components to protect the circuit from damage caused by overheating, by adjusting the operational frequency and/or voltage of components without simply shutting down the entire system. Ex-1008, Abstract, 2:15-34. Figure 2A depicts several components, including integrated circuit (light blue), multiple logic cores (red and blue), common bridge logic (green), and power management logic (yellow):

⁹ Ex-1002 ¶¶75-81.




FIG. 2A

Ex-1008, Figs. 2A. White teaches that the integrated circuit may include (or be coupled to) an L2 cache. Ex-1008, 24:30-41, 25:10-14, 25:45-49, Fig. 8.

Each logic core (red and blue) may be configured to operate at an independent voltage and frequency. Ex-1008, 6:44-49. Similarly, common bridge logic 110 (green), which is configured to communicate with the logic cores as well as with other elements internal and external to integrated circuit 100, may likewise operate at a frequency and/or voltage independent of other components. *Id.* at, 7:1-7, 14:32-38.

Power management logic 200 (yellow) can control, monitor, and adjust the operating frequency and/or voltage of various portions of integrated circuit 100, including logic cores 120 and common bridge logic 110. Ex-1008, 8:61-9:32, 14:32-38. Power management logic 200 may be part of a single set of logic, including

temperature control logic 210, capable of performing multiple functions. *Id.* at, 14:32-38, Fig. 2B.

E. Simeral (Ex-1009)¹⁰

Simeral, titled “Data Path Controller with Integrated Power Management to Manage Power Consumption of a Computing Device and Its Components” is U.S. Patent Publication No. 2009/0150689, published on June 11, 2009. Ex-1009, Cover. Simeral is prior art under at least pre-AIA 35 U.S.C. § 102(b).

Simeral discloses integrating power management functions into a data path controller to manage power consumption of a computing device. Ex-1009, Abstract. Specifically, Simeral discloses using ports for monitoring and/or quantifying power consumption of various components, modifying power consumption of components by adjusting their operational characteristics, and using profiles that specify permissible power consumption levels for components. Ex-1009, Abstract.

Figure 1A depicts several components, including processor (light blue), various devices (red and blue), bus (green), memory controller (dark red), and data path controller (yellow):

¹⁰ Ex-1002 ¶¶82-86.




FIG. 1A

Ex-1009, Fig. 1A.

Data path controller includes power management controller 106 for managing power consumption of components, including various devices and bus 130. Ex-1009 ¶19. Power management controller 106 monitors activity levels of components, compares activity levels to thresholds, and chooses performance profiles for components. *Id.* ¶¶24-27. Several general-purpose interface signal (“GPIS”) ports, such as 104a, 104b, and 104c, measure activity levels and GPIS port “k” does so for the bus. *Id.* ¶19.

XII. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS


A. Ground 1: Claims 1 and 12 Are Rendered Obvious by Bose

1. Independent Claim 1

a. Element 1[pre]: A processor comprising:

Bose discloses the preamble to the extent it is limiting. Ex-1002 ¶¶88-89.

Bose is directed to a “microprocessor system” (light blue), which is the claimed “processor.”




Ex-1005, Fig. 1. Bose’s systems and methods are “for controlling power and performance in a microprocessor system,” which “includes a monitoring and control system integrated into [the] microprocessor system.” *Id.* at 2:4-5, 4:21-24 (“Systems

and methods are described for controlling power consumption of a microprocessor chip or microprocessor-based system where the power of the chip (or system) is maintained within a programmable limit.”).

b. Element 1[a]: an integrated circuit including:

Bose discloses Element 1[a]. Ex-1002 ¶¶90-91.

Bose discloses that “the present invention can take the form of an integrated circuit chip.” Ex-1005, 4:67-5:3, 5:28-29. Bose further discloses that “the PMCU organization includes distributed or local components 114, 116, 118-1, 118-2...118-R and global components 110 engaged in the tasks of monitoring and controlling *in a multi-core chip 104*” (light blue). *Id.* at 5:64-67. Bose’s processor system comprising chip 104 (light blue) includes the claimed “integrated circuit”:




Ex-1005, Fig. 1.

c. Element 1[b]: a first core;

Bose discloses Element 1[b]. Ex-1002 ¶¶92-93.

Bose discloses that chip 104 comprises at least one core (red):




Ex-1005, Fig. 1. Bose describes “cores 112-1, 112-2, [and] 112-R” as part of the “management hierarchy” of the microprocessor and explains that local and global PMCUs monitor the cores to determine their activity. *Id.* at Fig. 1, 6:19-24, 5:64-67, 6:56-60.

d. Element 1[c]: a cache memory;

Bose discloses Element 1[c]. Ex-1002 ¶¶58, 94-96.

Bose discloses that chip 104 comprises L2/L3 cache 106 (dark red), which provides temporary storage, as part of the management hierarchy of the integrated circuit 104:




Ex-1005, Fig. 1, 5:21-22, 6:19-24, 5:64-67, 6:56-60.

A POSITA would have understood that L2/L3 cache hierarchy is a type of cache memory like the shared cache memory described by the '494 Patent. Ex-1002 ¶96; *see also* Ex-1001, 15:61-63.

- e. **Element 1[d]: a communication bus coupled to the first core, the communication bus to connect the first core and the cache memory;**

Bose discloses Element 1[d]. Ex-1002 ¶¶97-100.

Bose discloses a communication bus coupled to the first core and connecting that core to a cache memory. Bose discloses that chip 104 comprises several components, including an interconnect: “Each element or entity in this management hierarchy, e.g., cores 112-1, 112-2, 112-R or L2/L3 cache hierarchy 106, **interconnect 108 [(green)]**, local PMU facility 118-1, 118-2, 118-R, 114, 116, the chip-level global PMCU facility 110, or even the Operating System (OS) 102 itself, which has the highest authority in the hierarchy, is architected to provide the following general functions.” Ex-1005, Fig. 1, 6:19-24, 5:64-67, 6:56-60. Interconnect 108 (green) is a component of Bose’s chip 104:



Ex-1005, Fig. 1.

A POSITA would have understood that, while Figure 1 does not depict lines connecting interconnect 108 to core 112-1 and to L2/L3 cache hierarchy, Figure 1's arrowed lines represent the hierarchical organization of reports and commands between components, not their physical connections. Ex-1002 ¶¶99; Ex-1005, 3:59-61, 5:61-67, 6:18-38, 6:45-55. Based on Bose's depicted and described hierarchical organization and the well-known term of art "interconnect," it would have been obvious to a POSITA that Bose's interconnect 108 couples and connects to each core and to L2/L3 cache hierarchy and is the same type as described in the '494 Patent. Ex-1002 ¶¶57-58, 99-100.

f. Element 1[e]: a core workload monitor configured to determine a core workload for the first core;

Bose discloses Element 1[e]. Ex-1002 ¶¶101-104.

Bose discloses a core workload monitor. Specifically, Bose discloses that each core has an associated local PMCU monitor-and-control unit, which is a core workload monitor: "A set of *local PMCU monitor-and-control units are preferably associated with each physically monitored core*, storage and interconnect resource *on the chip* (or system)." Ex-1005, 3:24-28. Core 112-1 associates with local PMCU facility 118-1, whose functions include to "monitor events, activity levels and direct power and/or temperature levels or suitable proxies thereof, within its domain of authority." Ex-1005, 6:18-27. Local PMCU 118-1 (red) associates with core 112-1:




FIG. 1

Ex-1005, Fig. 1.

Bose's core workload monitor, PMCU 118-1, determines the workload for the first core. For example, Bose further discloses that local PMCU 118-1 monitors activities, performance statistics, power usage, and temperature usage in core 112-1:

The *local PMCUs* 114, 116, **118-1**, 118-2, 118-R *monitor activities in each core 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and interconnect 108 and carry out power-reducing optimizations as commanded by the global PMCU 110. The local PMCUs 114, 116, 118-1, 118-2, 118-R monitor performance statistics such as cache miss rates, branch misprediction rates, instructions per cycle (IPCs), etc...The *local PMCUs* 114, 116, **118-1**, 118-2, 118-R also *monitor information such as power usage and temperature usage and report the information to the global PMCU 110 periodically.**

Ex-1005, 6:56-7:3.

A POSITA would have understood that local PMCU 118-1 is a core workload monitor because monitoring activities, performance statistics, power usage, and temperature usage are examples of determining a workload of the core, as described by the '494 Patent. Ex-1002 ¶¶104; Ex-1001, 9:49-57 (“Determining a workload includes any known method for determining...use of a processing element....”).

g. Element 1[f]: a bus workload monitor configured to determine a bus workload for the communication bus; and

Bose discloses Element 1[f]. Ex-1002 ¶¶105-108.

Bose discloses a bus workload monitor. Specifically, Bose discloses that the interconnect has an associated local PMCU monitor-and-control unit: “A set of *local PMCU monitor-and-control units are preferably associated with each physically monitored* core, storage and *interconnect* resource *on the chip* (or system).” Ex-1005, 3:24-28. Interconnect 108 associates with local PMCU facility 116, whose general functions include to “monitor events, activity levels and direct power and/or temperature levels or suitable proxies thereof, within its domain of authority” and “report on a summary of monitored results and actions taken, to the next higher authority in the power management hierarchy.” Ex-1005, 6:18-37. Local PMCU 116 (green) associates with interconnect 108:




FIG. 1

Ex-1005, Fig. 1.

Bose's bus workload monitor, PMCU 116, determines the bus workload for the communications bus. For example, Bose further discloses that local PMCU 116 monitors activities, performance statistics, power usage, and temperature usage in interconnect 108:

The *local PMCUs* 114, **116**, 118-1, 118-2, 118-R *monitor activities in* each core 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and *interconnect 108* and carry out power-reducing optimizations as commanded by the global PMCU 110. *The local PMCUs* 114, **116**, 118-1, 118-2, 118-R *monitor performance statistics such as cache miss rates, branch misprediction rates, instructions per cycle (IPCs), etc....* The *local PMCUs* 114, **116**, 118-1, 118-2, 118-R also *monitor information such as power usage and temperature usage* and report the information to the global PMCU 110 periodically.

Ex-1005, 6:56-7:3; *see also* 6:18-38.

A POSITA would have understood that local PMCU 116 is a bus workload monitor because monitoring activities, events, performance statistics, power usage, and temperature usage are examples of determining a workload of the core and/or communication bus, as described by the '494 Patent. Ex-1002 ¶108; Ex-1001, 9:49-57; 17:22-29.

- h. Element 1[g]: balancing control adapted to receive the bus workload from the bus workload monitor and to dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit and a comparison between the bus workload and a bus workload threshold, the power limit corresponding to a maximum thermal dissipation capacity for the integrated circuit, wherein a power consumption of one of the first core and the communication bus is to be reduced, the power consumption reduction to be limited to maintain operation of the one of the first core and the communication bus above a low limit.**

- (1) 1[g.i]: balancing control adapted to receive the bus workload from the bus workload monitor**

Bose discloses 1[g.i]. Ex-1002 ¶¶109-112. Bose discloses that chip 104 comprises a global PMCU 110 (yellow), which performs the claimed “balancing control”:




FIG. 1

Ex-1005, Fig. 1.

Bose discloses that each local PMCU (e.g., core workload monitor and bus workload monitor in Elements 1[e] and 1[f]), reports a summary of monitored results to global PMCU 110: “Each element...is architected to...*report on a summary of monitored results* and actions taken, to the next higher authority in the power management hierarchy: e.g., each *local PMCU* 118-1, 118-2, 118-R, 114, *116 reports upwards to the on-chip global PMCU 110.*” Ex-1005, 6:18-27. Indeed, Bose discloses that “[t]he *local PMCU*s 114, *116*, 118-1, 118-2, 118-R also monitor information such as power usage and temperature usage and *report the information to the global PMCU 110* periodically.” Ex-1005, 6:56-7:3.

Thus, Bose's global PMCU 110 can receive a summary of monitored results from local PMCU 116 associated with interconnect 108 and is therefore a balancing control adapted to receive the bus workload. Ex-1002 ¶112.

- (2) **1[g.ii]: to dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit...the power limit corresponding to a maximum thermal dissipation capacity for the integrated circuit**

Bose discloses 1[g.ii]. Ex-1002 ¶¶113-114. Bose discloses that “global PMCU 110 is given an overall chip 104 power budget, and is responsible for assigning the budget to each core element 112-1, 112-2, 112-R, 106, and 108,” according to “thermal packaging, power supply and other considerations.” Ex-1005, 7:18-23. The global PMCU monitors for events exceeding acceptable system limits, “such as over the budget of power, over voltage or overheating emergencies.” *Id.* at 9:39-43. Thus, the global PMCU 110 optimizes power between the cores and interconnect within this power budget. *Id.* at 7:25-28 (“global PMCU 110 will choose the power modes and assign power for each core element **112-1**, 112-2, 112-R, 106, and **[interconnect] 108** such that all elements can run at the highest possible performance **within the budget.**”), 11:60-66 (the global PMCU performs and maintains “monitoring of one or more of a maximum power of a chip (or system) within a programmable limit, a peak temperature across monitored regions of the chip (or

system) below a specified limit, and/or a net throughput performance of the chip (or system) within the specified constraints of power and temperature.”); 7:8-11, 8:11-17, Figs. 2, 7.

A POSITA would have understood such power optimizations as the global PMCU dynamically tuning power allocation between the core 112-1 and interconnect 108. Ex-1002 ¶¶42-47, 114. Bose’s overall chip power budget is a power limit for the integrated circuit. *Id.* For example, Bose explains that “the peak temperature limit of the chip (or system) may be a preset design specification, determined by the package/cooling solution engineered for the chip (or system).” Ex-1005, Fig. 7, 12:18-21. Thermal packaging or package/cooling solution sets a maximum thermal dissipation capacity for the integrated circuit, which in turn discloses the power limit for the integrated circuit. Ex-1002 ¶114.

(3) 1[g.iii]: and a comparison between the bus workload and a bus workload threshold

Bose discloses 1[g.iii]. Ex-1002 ¶¶115-123. Global PMCU 110 calculates for each core a “power budget” and “power mode.” Ex-1005, 8:35-40. As their names suggest, the “power budget” is the power available to a core, and the “power mode” is the power assigned to the core in a given workload operational scenario. Ex-1002 ¶¶115-117. Each “power mode” has an associated “power budget.” The global PMCU “is responsible for assigning the budget to each core element” within an

overall chip power budget and sending “the associated power budget” to each core. Ex-1005, 7:18-21, 8:19-21. A POSITA would have thus understood Bose’s power budget as defining the workload threshold for a given workload operating in an associated “power mode.” Ex-1002 ¶¶115-117. For example, if a core operating in a certain “power mode” uses substantially more or less power than the associated “power budget,” then the global PMCU may switch that core to a different, more appropriate power mode.

Bose further teaches a predicted power and performance for each defined power mode included in its “Power/Performance Trade-off Table 206.” In another aspect, these predicted values disclose the claimed workload thresholds. For example, the performance optimizations performed by the global PMCU can place the various components, such as the interconnect and cores, “into different power modes based on their power-performance trade-off characteristics,” where “[a] power mode may be defined to be higher if it is expected to consume more power and to achieve higher performance, and to be lower if it is expected to consume less power and to achieve lower performance, relative to other power modes.” Ex-1005, 7:11-17; 9:25-30, 9:59-10:24. Thus, the “power budget,” as discussed above, may be considered a workload threshold as well as the predicted values included in the Trade-off Table 206 may be a workload threshold. A POSITA would have thus understood each component’s power mode to be representative of the workload at

which the core is currently executing. Ex-1002 ¶118. And, if the workload threshold (i.e., the associated “power budget” for that power mode) is too high or too low in comparison with the current workload, then the global PMCU may switch that core into a different power mode more appropriate for its current workload. Additionally (or alternatively) this switch over may be based on the defined power modes in Trade-off table 206, which would also have been considered another workload threshold. *Id.*

Bose teaches reading real time statistics for the workload. Bose teaches that “global PMCU 210 stores history information” (such as “performance statistics” and “performance events”) “periodically supplied from units of all the local PMCUs 240 in real time in the Phase History Tables 213.” Ex-1005, Fig. 2, 7:41-61. The global PMCU reads those real time environmental statistics from the Phase History Table 213. *Id.* at 8:11-12. A POSITA would have understood such real time environmental statistics, performance statistics, and performance events as measuring the current workload because they measure contemporaneous operating statistics of each component.

A POSITA would have understood that Bose suggests comparing those real time statistics to the power budget assigned to each component, as determined in Power/Performance Trade-off Table 206, because doing so would yield the most accurate and informative way for the system to adjusting power allocation between

components without exceeding their thresholds. Ex-1005, 8:11-18; Ex-1002 ¶¶120-122. Specifically, “[b]y reading real time environmental statistics from the Phase History Table,” “Global PMCU controller 216 can dynamically perform optimizations for performance and power from a globally optimal standpoint.” Ex-1005, 8:11-18, Fig. 2. Indeed, “[t]he Power/Performance Trade-off Table 206 includes the predicted power and performance for each power mode,” which “global PMCU controller 216 reads...to assist in calculating the power modes and power budgets for each core.” *Id.*, 8:35-40. The local PMCUs operate with the newly assigned power modes (corresponding to the claimed workload), with their associated “power budgets” (corresponding to the claimed workload threshold), and when a local PMCU reports new real time information “such as over the budget of power, the global PMCU takes corrective action by switching that core to a power mode more appropriate for the current workload. Ex-1005, 9:36-43.

Although Bose discloses that the global PMCU calculates the power budget and power mode for each core but does not explicitly state the same for the interconnect, a POSITA would have understood that the same calculation applies to the interconnect. Ex-1002 ¶123. Phase History Table 213 receives history information, performance statistics, and performance events “[f]rom all Local PMCUs,” which includes local PMCU 116 on interconnect 108. Ex-1005, Fig. 2, 6:56-60. Such history information, performance statistics, and performance events

from the interconnect includes workload of the interconnect. Ex-1002 ¶123. After calculating, the global PMCU sends the associated power budget and power mode “[t]o each Local PMCU,” which includes local PMCU 116 on interconnect 108. Ex-1005, Fig. 2, 6:56-60. Indeed, the global PMCU “choose[s] the power modes and assign[s] power for each core element,” including for interconnect 108. *Id.*, 7:25-27, 7:9-11, 7:19-20.

- (4) **1[g.iv]: wherein a power consumption of one of the first core and the communication bus is to be reduced, the power consumption reduction to be limited to maintain operation of the one of the first core and the communication bus above a low limit.**

Bose discloses 1[g.iv]. Ex-1002 ¶¶124-125. Bose discloses reducing power consumption of either the core or communication bus to above a low limit. Bose discloses a low limit of when a core is powered off, as its system maximizes performance by “use and deployment of one or more additional cores (*that would otherwise be powered off* or excluded from the design in the absence of a dynamic power management method such as described herein).” Ex-1005, 2:43-48. Bose discloses operating above such a low limit through its low power mode for the global PMCU. For example, Bose teaches that a “power mode may be defined to be higher if it is expected to consume more power and to achieve higher performance, and *to*

be lower if it is expected to consume less power and to achieve lower performance, relative to other power modes.” *Id.* at 7:11-17.

Bose teaches reducing power allocated to components to maintain their operation without disabling them entirely. Ex-1005, 2:43-48. Also, “if the core power target is not met,” the local PMCU “will inform the global PMCU to assign a lower power mode to this core,” and “[t]he local PMCU may decide to **dramatically reduce activities** on the core if the power budget is severely insufficient.” *Id.* at 11:47-53; *see also* 7:65-8:10. A POSITA would have understood reducing components’ allocated power budgets to lower power modes while maintaining operation of a core or interconnect above a low limit, namely when they are powered off. Ex-1002 ¶125.

2. Dependent Claim 12: The processor of claim 1, wherein the balancing control is to receive a hint regarding a future workload and to dynamically tune the power allocation further based on the hint.

Bose discloses Claim 12. Ex-1002 ¶¶126-127.

Bose discloses the balancing control receiving a hint regarding future workload and dynamically tuning power allocation based on a hint. Bose discloses global PMCU “employ[ing] a phase prediction technique” in order “[t]o better assign power to each core element 112-1, 112-2, 112-R, 106, and 108.” Ex-1005, 7:4-7. Bose explains: “To perform effective optimizations, the global PMCU 210

relies on...[t]he Phase Prediction Table 204...*to perform accurate phase prediction, in which the word phase represents the durations and values of the power and performance related statistics* such as those provided by the Phase History Table 213.” *Id.* at 8:19-26. A table predicting changes and duration of distinct workload phases based on power and performance statistics from all local PMCU’s provides a hint to the global PMCU about future workloads and enables the global PMCU to dynamically tune power allocation accordingly.


B. Ground 2: Claims 3, 12, 14, and 15 Are Rendered Obvious by Bose Combined with Fukatsu

1. A POSITA Would Have Been Motivated to Combine Bose and Fukatsu

A POSITA would have been motivated to combine the teachings of Bose and Fukatsu and would have had a reasonable expectation of success in doing so because each relates to the same well-known issues. Ex-1002 ¶¶42-58, 128-136. Both teach methods for controlling power and performance in a processor to improve processing efficiency and reduce power consumption and making trade-offs between power performance and processing efficiency. Ex-1005, Abstract, 7:8-17; Ex-1006, Abstract, ¶¶5-8, 95, Figs. 1, 8.


Moreover, both Bose and Fukatsu teach similar components and structures, with shared resources connected to processor cores through a bus or interconnect. Bose teaches depicts processor 100 (light blue), cores 112-1 (red and blue),

interconnect 108 (green), shared cache 106 (dark red), and global PMCU 110 (yellow):



Ex-1005, Fig. 1.

Likewise, Fukatsu teaches processing device (light blue), CPUs 114a (red and blue), bus 105 (green), shared memory 104 (dark red), and controlling unit 161 (yellow):



Ex-1006, Fig. 12. Although Fukatsu discloses memory 104 while Bose discloses L2/L3 cache, a POSITA would have recognized their similar teachings to connect the bus/interconnect to the shared memory/cache resource, as Fukatsu’s bus 105 connects to shared memory 104 and Bose’s interconnect connects to L2/L3 cache. Ex-1002 ¶¶57-58, 129-133.

A POSITA would have been motivated to combine the teachings of Fukatsu with Bose. Ex-1002 ¶134. For example, while Bose discloses a local power management control unit for managing power allocated to an interconnect, Fukatsu provides teachings regarding how to adjust an interconnect based on measured and predicted core activity (for which Fukatsu uses the core’s hit rate for its on-board cache). Ex-1006, Fig. 8, ¶¶6-7, 95. By adjusting the interconnect clock frequencies based on measured and predicted core activity, a POSITA would have recognized that Fukatsu’s teaching would improve Bose’s power distribution efficiency. Ex-1002 ¶134. Specifically, by not allocating excess power to an interconnect when the

core is heavily occupied, a POSITA would have understood that additional power is available to the cores themselves, while allocating additional power to the interconnect when significant traffic is expected (e.g., calls to the shared L2/L3 cache) would prevent the interconnect from becoming a bottleneck to overall processor performance. Ex-1006, Fig. 8, ¶¶6-7, 95; Ex-1002 ¶134.¹¹

Additionally, a POSITA would have a reasonable expectation of success in combining the teachings of Fukatsu and Bose. For example, a POSITA would have understood how to implement Fukatsu’s teachings to adjust interconnect frequency based on core activity with Bose’s power management control units. Ex-1002 ¶¶135-136. As Bose explains, the PMCUs “include a consolidated software layer (possibly including a more generalized combination of firmware, operating system (OS) and hypervisor layers).” Ex-1005, 3:18-20. Additionally, “a set of local PMCU monitor and-control units are preferably associated with each physically monitored core, storage and interconnect resource on the chip.” *Id.*, 3:24-27. And Fukatsu teaches techniques to monitor core and bus activity using simple-to-implement hardware or software, which does not require “special hardware”: “[w]ith this structure, mounting special hardware used for monitoring an actual cache hit rate is

¹¹Although Fukatsu explicitly discloses a memory and storage medium as examples of shared resources connected to the bus, a POSITA would have understood that other shared resources, such as the L2/L3 cache disclosed by Bose would also be connected using Fukatsu’s bus. Ex-1002, ¶¶___; Ex-1005, Fig. 1; Ex-1006, Fig. 12.

unnecessary, and *it is possible to easily predict a hit rate by software or hardware.*”

Ex-1006 ¶12. Accordingly, since Bose’s PMCUs are designed to be implemented using generalized software layers, and Fukatsu teaches using standard and non-specialized software and hardware, a POSITA would have had a reasonable expectation of success in implementing these teachings together; doing so would be nothing more than applying predictable solutions involving routine hardware and software functionalities. Ex-1002 ¶¶135-136.

2. **Dependent Claim 3: The processor of claim 1, wherein the balancing control, in response to the integrated circuit operating at the power limit, being adapted to: increase frequency for the communication bus and cap frequency for the first core in response to the bus workload being greater than the bus workload threshold; and decrease frequency for the communication bus and increase the frequency for the first core in response to the bus workload being less than the bus workload threshold**

Bose-Fukatsu disclose Claim 3. Ex-1002 ¶¶137-141.

Bose teaches balancing power allocation between components when operating at a power limit. For example, Bose teaches that “if the actual total power is not below or equal to the total power budget,” the global PMCU “assigns a lower power mode for a core if the power budget assigned to the core is insufficient and increases to higher power modes if the cores are exceeding the power budget, such that the sum of power usage for all components fits into the overall power budget.” Ex-1005, 9:14-32; Ex-1002 ¶¶42-47, 137.


Fukatsu also recognizes varying the frequency and voltage of both the CPU (i.e., core) and bus to “achieve both an improvement in processing efficiency and a reduction in electric power consumption.” Ex-1006 ¶7. Fukatsu explains that it is optimal to vary the CPU’s voltage and frequency inversely to the bus’s voltage and frequency. Fukatsu discloses about the prior art:

The clock controlling unit 0103 controls a frequency of a clock 0105 which is to be supplied to the CPU 0101 in accordance with the cache hit rate information 0104. In doing so, ***the clock controlling unit 0103 increases the clock frequency of the CPU 0101*** when the cache hit rate is high, and ***decreases the clock frequency of the CPU 0101*** when the cache hit rate is low. At the same time, the clock controlling unit 0103 ***decreases clock frequencies of a bus*** and a memory connected to the CPU 0101 when the cache hit rate is high, and ***increases the clock frequencies of the bus*** and the memory when the cache hit rate is low.

Ex-1006 ¶6. The rationale for this inverse relationship is that “the CPU effectively performs its processing when the hit rate is high, and the speed of the bus and the memory is increased when the hit rate is low.” *Id.* ¶¶6-7. Indeed, “by controlling the frequencies 204 of the CPU 101 based on the cache hit rate 202 predicted by the cache hit rate predicting unit 201, the controlling unit 161 ***decreases the frequencies 204 of the CPU 101 when the cache hit rate 202 is low, and increases the frequencies 204 of the CPU 101 when the cache hit rate 202 is high***” providing “optimal[] control [of] electric power consumption.” *Id.* ¶57.

Fukatsu’s Figure 8 depicts the inverse relationship between CPU clock frequency (middle graph) and bus frequency (bottom graph):

FIG. 8



Ex-1006, Fig. 8. In blue, when the CPU clock frequency increases, bus frequency decreases or sets an upper limit. Ex-1006, Fig. 8; Ex-1002 ¶140. Conversely, in red, when the CPU clock frequency decreases or sets an upper limit, bus frequency increases. Ex-1006, Fig. 8; Ex-1002 ¶140.

A POSITA would have been motivated to combine the teachings of Fukatsu with those of Bose to achieve a more efficient processor for balancing power and performance and specifically to maximize the power allocated to the busiest

functional units, whether core or interconnect. Ex-1002 ¶141; Ex-1006 ¶¶5-7, 86, 95. It would have been obvious to a POSITA to apply the technique of frequency balancing of the bus 105 and memory 104 and the CPU in Fukatsu to the cache interconnect structure of Bose. Ex-1002, ¶141. A POSITA would have had a reasonable expectation of success in implementing this inverse core/interconnect power allocation because Bose already teaches allocating power between an interconnect and the core using a combination of global and local power management control units. Fukatsu's teachings are thus a complementary implementation detail that would have been implemented in Bose using the same hardware and routine software programming. Ex-1002 ¶141; Section XII.B.1.

3. Dependent Claim 12

Bose-Fukatsu disclose Claim 12. Ex-1002 ¶¶142-144.

In addition to the teachings of Bose described in Ground 1, Fukatsu discloses the balancing control receiving a hint regarding future workload and dynamically tuning power allocation based on the hint. Fukatsu teaches that controlling unit 161 determines frequency and power voltage based on a predicted hit rate: "controlling unit 161: *predicts a hit rate* of the cache memory 160; *determines a frequency of the clock signal in accordance with the predicted hit rate*; and controls the clock supplying unit 150 so as to change to the determined frequency[,]*...determines a power voltage in accordance with the determined frequency*, and controls the

power unit 151 so as to change to the determined power voltage.” Ex-1006 ¶55. It would have been obvious to a POSITA that a predicted hit rate of the cache memory would provide a hint to the controlling unit about future workloads. Ex-1002 ¶143.

A POSITA would have been motivated to combine the teachings of Bose and Fukatsu and would have had a reasonable expectation of success in making the combination for the same reasons in Section XII.B.1. Moreover, a POSITA would have expected such a combination to be successful because Fukatsu’s prediction technique would work seamlessly with Bose’s already-disclosed use of a Phase Prediction Table. Ex-1005, 3:7-14; Ex-1002 ¶144.

4. Dependent Claim 14: The processor of claim 3, wherein the balancing control is to increase the communication bus frequency by a different amount than the first core frequency decrease.

Bose-Fukatsu disclose Claim 14. Ex-1002 ¶¶42-52, 145-148.

A POSITA would have expected that in the Bose-Fukatsu system, when balancing power between multiple elements, the resulting frequency adjustments would be in different amounts for at least two reasons.

First, a POSITA would have understood that the relationship between frequency and voltage is not uniform between computer components and would have expected that frequency-voltage relationship to be different between a core and bus. Ex-1002 ¶¶42-47, 147 (citing Ex-1022). Even assuming that voltage is


conserved (e.g., the circuit is at a power limit), the change in frequency resulting from adjusting between a single core and bus would not be in the same amount. Indeed, even the '494 Patent recognizes that “increasing a bus frequency [may result] in less extra power consumption than the same reduction in frequency for cores.” Ex-1001, 17:44-54. This is consistent with what a POSITA would have known at the time of the '494 Patent. Ex-1002 ¶147.

Second, even if frequency and voltage are assumed to have a one-to-one relationship in different components, a POSITA would have understood that the multi-component adjustments in Bose and Fukatsu would include scenarios in which an increase in the power/frequency allocated to one component was offset by corresponding decreases to the power/frequency in *multiple* components. Ex-1002 ¶148. For example, Bose discloses allocating power/frequency between at least three cores, an interconnect, and a cache. Similarly, Fukatsu, discloses allocating power/frequency between at least a CPU, a bus, and a memory. This is a common scenario in the art as well. *See* Ex-1002 ¶148(citing Ex-1022).

5. **Dependent Claim 15: The processor of claim 3, wherein the balancing control is to further increase frequency for a second core, the second core and the communication bus having power planes tied together.**


Bose-Fukatsu disclose Claim 15. Ex-1002 ¶¶48-52, 149-153.

Bose teaches a processor with a second core. Figure 1 shows cores 112-1 (red) and 112-2 (blue) on Bose's processor 100 (light blue):



Ex-1005, Fig. 1, 5:61-67, 6:18-24.

Fukatsu teaches a similar “multiprocessor configuration,” which has “plural sets of a CPU” (red and blue), in which the second core and the communication bus have their power planes tied together:



Ex-1006, Fig. 12.

Fukatsu explains that “a clock 102 in each of the CPU 114a-114c is independently controlled” and further explains that “it is necessary to arbitrate details of the setting of the frequencies between the CPU clock, the bus clock and the memory clock, or to fix these frequencies.” Ex-1006 ¶114. A POSITA would have understood that Fukatsu’s multiple power planes are tied together because changes to the frequency/voltage of shared bus 105 and shared resources such as memory 104 would necessitate accounting for these changes across various power planes. *Id.*; Ex-1002 ¶¶42-52, 152.


C. Ground 3: Claims 1, 3, 14, and 15 Are Rendered Obvious by Bose Combined with Naffziger.

1. A POSITA Would Have Been Motivated to Combine Bose and Naffziger

A POSITA would have been motivated to combine the teachings of Bose and Naffziger and would have had a reasonable expectation of success in doing so


because each relates to the same well-known issues. Ex-1002 ¶¶42-58, 154-158. Both relate to similar methods for controlling or managing power consumption and performance in a processor, and both specifically describe transferring or adjusting power allocation between cores or dies for efficiently managing or controlling the processor's power consumption. Ex-1005, Abstract, 7:8-17, 24-31; Ex-1007, Abstract, 1:6-38.

Moreover, both Bose and Naffziger have similar components and structures. Bose's teaches processor 100 (light blue), cores 112-1 (red and blue), interconnect 108 (green), shared cache 106 (dark red), and global PMCU 110 (yellow):



Ex-1005, Fig. 1.

Likewise, Naffziger teaches processor 700 (light blue), cores 112 (red and blue), crossbar switch 116 (green), shared cache memory subsystem 118 (dark red), and power management unit 130 (yellow):



Ex-1007, Fig. 8. Given Bose's and Naffziger's similar components and structures, a POSITA would have understood that they are compatible. Ex-1002 ¶¶48-58, 154-156.

A POSITA would have been motivated to combine the teachings of Bose with Naffziger to create a more efficient and better functioning processor for controlling or managing power consumption and performance in a processor. Ex-1002 ¶157. First, a POSITA would have been motivated to utilize Naffziger's use of thresholds

in determining whether the activity level of a computational unit requires allocating additional power to that unit. Ex-1007, Fig. 4, 10:46-57, 11:3-13. A POSITA would have recognized that this feature would provide for more efficient power allocation between components. Ex-1002 ¶157. Second, while Bose describes assigning power modes to local elements (*see* Ex-1005, 6:65-7:17, 7:24-31), a POSITA would have been motivated to use Naffziger’s pre-defined P-states as a more effective way to predictably adjust power allocation between components (*see* Ex-1007, Fig. 4, 10:46-57, 11:3-13). Ex-1002 ¶157. A POSITA would have also recognized that these discrete P-states also ensure that performance stays within certain minimum levels required for operation. Ex-1007, Fig. 2, 6:65-7:4; Ex-1002 ¶157.

Additionally, a POSITA would have a reasonable expectation of success in combining the teachings of Naffziger and Bose. For example, a POSITA would have understood how to implement Naffziger’s pre-defined power state teachings with Bose’s power management control units and the cores they control. Ex-1002 ¶158. As Naffziger explains, the controlled “computation[al] unit 120 may be one of *general purpose processors*, graphics processors, special-purpose coprocessors, and so forth.” Ex-1007, 3:41-44. Naffziger further explains that “[t]aking *cores within modern microprocessors* as an example of a computation unit 120, these cores *allow flexibility* regarding its operational mode.” Ex-1007, 5:49-51. As Bose explains, the PMCU’s “include a consolidated software layer (possibly including a more

generalized combination of firmware, operating system (OS) and hypervisor layers).” Ex-1005, 3:18-20. Accordingly, since Naffziger envisions that its teachings will be applied to “flexible,” “general purpose,” “modern microprocessors,” and Bose’s PMCU’s are designed to be implemented using generalized hardware and software layers, a POSITA would have had a reasonable expectation of success in implementing these teachings together; doing so would be nothing more than applying predictable solutions involving routine hardware and software functionalities. Ex-1002 ¶158.

2. Independent Claim 1

a. Elements 1[pre] through 1[f]: A processor comprising:

Bose discloses Elements 1[pre] through 1[f] for the reasons in Ground 1. Ex 1002 ¶159.

b. Element 1[g]: balancing control...

Bose discloses Element 1[g] for the reasons in Ground 1. A POSITA seeking to implement the power management scheme of Bose would have found this feature obvious in view of Naffziger, which also teaches aspects of this limitation.

- (1) **1[g.ii]: to dynamically tune power allocation...based on a power limit for the integrated circuit...the power limit corresponding to a maximum thermal dissipation capacity...**

Naffziger discloses 1[g.ii]. Ex-1002 ¶¶160-161. Naffziger teaches “[a] power target may be assigned to die 102” and “is a thermal design power value for the die 102.” Ex-1007, 3:46-50. “The thermal design power (TDP), which may be also referred to as a thermal design point, represents a maximum amount of power a cooling system in a computer is able to dissipate.” *Id.*; *Id.* at 7:25-29. And Naffziger further teaches “[a] sum of the associated power credits equals the total number of power credits for die 102.” *Id.* at 4:6-11. Naffziger further teaches that “[t]he power management unit 130 may allocate a separate given number of power credits to each computation unit 120.” *Id.* at 4:1-5, 5:12-40. Thus, Naffziger teaches allocating power credits between computational units and that the number of power credits available to be allocated depends on the TDP for the integrated circuit, which a POSITA would have understood corresponds to a maximum amount of power a cooling system in a computer is able to dissipate. Ex-1002 ¶161.


- (1) **1[g.iii]: and a comparison between the bus workload and a bus workload threshold**

Naffziger discloses 1[g.iii]. Ex-1002 ¶162. Naffziger teaches that its “power management unit 130 may have one large table...configured to track reported activity levels of each computation unit, compare the reported activity levels to

corresponding thresholds, determine which computation units are able to donate power credits, determine an amount of power credits that each qualified donating computation unit is able to donate, determine which computation units are able to receive power credits, and determine an amount of power credits that each qualified receiving computation unit is able to receive.” Ex-1007, 11:3-13, Fig. 4, 10:46-57. A POSITA would have readily applied this threshold comparison to Bose’s similar power management scheme because Bose’s tables for calculating power budgets and power modes includes performance statistics from all local PMCUs, which includes the workload of the interconnect. Ex-1005, Fig. 2, 6:56-60, 7:41-61, 8:45-51; Ex-1002 ¶162.

(2) 1[g.iv]: ...maintain operation...above a low limit

Naffziger discloses 1[g.iv]. Ex-1002 ¶163. Naffziger teaches discrete P-states, which correspond to different performance levels based on their frequency and voltage. These discrete operational states include P-state P₄, which Naffziger teaches “may correspond to a discrete state with a lowest performance of all the supported discrete states and comprises the lowest operational frequency,” and “may correspond to a discrete state with a lowest power consumption of all the supported discrete states and comprises the lowest operational voltage.” Ex-1007, 6:65-7:4. Figure 2 depicts the lowest P-state P₄ (teal):



Ex-1007, Fig. 2. A POSITA would have been understood Naffziger's P-state P₄ as disclosing the lowest functioning operational state of computation units, which is above a non-operational low limit. Ex-1002 ¶163. Also, a POSITA would have been motivated to utilize Naffziger's P-States to provide predictable performance to a system like that described by Bose. *Id.* For example, Naffziger teaches that its P-states reflect predictable performance of computation units: "[t]ypically, the endpoint discrete states represented by P-states P₀ and P₄ define a region of predictable performance," so "configuring a processor to support multiple P-states, or operating points, along the non-linear frequency versus voltage curve may provide stable, optimal utilization of power and delivery of performance for the semiconductor chip, such as a processor." *Id.* at 7:10-16.

3. Dependent Claim 3

Bose-Naffziger disclose Claim 3. Ex-1002 ¶¶42-47, 164-170.

Bose teaches portions of Claim 3 for the reasons in Ground 2. Ex 1002, ¶165.


A POSITA seeking to implement the power management scheme of Bose would have found this feature obvious in view of Naffziger, which also teaches aspects of this limitation.

Naffziger teaches increasing or decreasing power credits to computation units depending on their activity levels. For example, “[w]hen a given computation unit does not have a high or moderate workload” and “its activity level may decrease below a given threshold...the power management unit 130 may redistribute the power credits,” for example by “*lend[ing] power credits of a determined inactive computation unit to a computation unit that is determined to be highly active.*” Ex-1007, 5:28-37. POSITA would have understood that decreasing the power credits allocated also necessarily caps the computational unit’s power levels. Ex-1002 ¶166.

Conversely, Naffziger teaches returning power credits to a computation unit when that unit’s activity level rises above a threshold:

If the die has reassigned power credits and the power management unit detects a donating computation unit substantially increases its activity level (conditional block 612) or otherwise requires additional power, then in block 614, the power management unit may reallocate the power credits across the on-die computation units. ... If the respective *activity level has risen above the threshold*, then the *power management unit may convey a respective number of power credits from each of the previously receiving computation units to the previous donating computation unit.*

Ex-1007, 16:29-44. Figure 7 depicts returning power credits (red) when detected activity level is above a threshold for a given computation unit:



Ex-1007, Fig. 7.

Although Naffziger teaches increasing and decreasing (or setting an upper limit of) power credits for computation cores only, a POSITA implementing Bose would have naturally extended this technique to the interconnect as well based on Bose's recognition that the global PMCU manages the interconnect's power by assigning a portion of the power budget to the interconnect. Ex-1005, 7:8-31; Ex 1002, ¶¶56-57, 166-169.

4. Dependent Claim 12


Bose-Naffziger disclose Claim 12 for the reasons in Ground 1. Ex-1002 ¶171.

5. Dependent Claim 14

Bose-Naffziger disclose Claim 14. Ex-1002 ¶¶42-52, 172-177.

In addition to the teachings of Bose in Ground, 2, Naffziger teaches allocating different amounts of power credits to computation units depending on each's production of thermal energy during high workloads. Naffziger teaches that "[t]he power management unit 130 may allocate a separate given number of power credits to each computation unit 120." Ex-1007, 4:1-11. As an example, "three computation units may produce different thermal energy levels during high workloads," where "[a] first computation unit may produce three times the thermal energy produced by the third computation unit under high workloads" and "[a] second computation unit may produce two times the thermal energy produced by the third computation unit under high workloads." *Id.* at 5:12-18. "The power management unit 130 may be configured to adjust the number of power credits for each of the computation units 120 over time." *Id.* at 4:1-11. A POSITA would have understood that adjusting and allocating power credits among computation units in turn affects their operational frequencies. Ex-1007, Fig. 2, 6:55-7:16; Ex 1002 ¶¶42-47, 173.

Further, Naffziger discloses a non-linear power-frequency relation for each computation unit:




Ex-1007, Fig. 2, 6:55-59. Boosting to a higher P-state increases frequency by a smaller amount than frequency decreases when throttling to a lower P-state. *Id.* A POSITA would have understood that such a non-linear power-frequency relation means frequencies of different components necessarily change by different amounts when throttling one and boosting the other to different P-states. Ex-1007, 7:65-8:15; Ex 1002, ¶¶175-176.

6. Dependent Claim 15

Bose-Naffziger disclose Claim 15. Ex-1002 ¶¶48-52, 178-182.

Like Bose in Ground 1, Naffziger teaches a second core. Ex-1007, 16:45-51.

Naffziger's multiple cores 112 (red and blue) are on processor 700 (light blue):



Ex-1007, Fig. 8.

Naffziger further teaches coupling computation units to the same voltage plane. In Figure 8, “the CPUs 115 share a same voltage regulator 110a.” Ex-1007, Fig. 8, 17:60-61. A POSITA therefore would have understood that computation units could be tied to the same power plane. Ex-1002 ¶180.

Although Naffziger teaches tying together voltage planes of computation units (but not with crossbar switch 116), based on Bose’s teaching of similarly allocated power between its cores and interconnect, it would have been obvious to a POSITA that the same method could be applied to its packet processing logic or crossbar switch 116, and by extension to Bose’s interconnect. Ex 1002 ¶181.

D. Ground 4: Claims 1, 3, 14, and 15 Are Rendered Obvious by White Combined with Simeral

1. A POSITA Would Have Been Motivated to Combine White and Simeral

A POSITA would have been motivated to combine the teachings of White and Simeral and would have had a reasonable expectation of success in doing so because both relate to similar methods of managing power in circuits by adjusting power consumption of circuit components. Ex-1002 ¶¶42-58, 183-187; Ex-1008, Abstract; Ex-1009, Abstract. Specifically, White teaches managing power by adjusting the power consumption and operating frequencies of logic cores and common bridge logic to avoid overheating the integrated circuit, Ex-1008, Abstract, 2:15-34, while Simeral similarly discloses adjusting the operational characteristics of the bus in comparison to a threshold activity level. Ex-1009, Abstract, ¶¶21, 24. A POSITA would have been motivated to combine White's teachings with Simeral's teachings for managing power by monitoring activity associated with the bus in comparison to a threshold and accordingly adjusting circuit components activity levels because Simeral's teachings would improve White's ability to respond quickly to power-related events, such as spikes in bus activity, thus increasing the number and duration of power-saving opportunities. Ex-1009, Abstract, ¶¶4-9, 19-24; Ex 1002, ¶183.

Moreover, both White and Simeral teach similar components and structures. White teaches integrated circuit 100 (light blue), cores 120A and 120B (red and blue), common bridge logic 110 (green), and power management logic 200 (yellow):




FIG. 2A

Ex-1008, Fig. 2A.

Likewise, Simeral teaches processor (light blue), various devices (red and blue), bus (green), and data path controller (yellow):




FIG. 1A

Ex-1009, Fig. 1A. Given White's and Simeral's similar components and structures, a POSITA would have understood that they are compatible. Ex-1002 ¶¶184-185.

A POSITA would have been motivated to combine the teachings of Simeral and White. Ex-1002 ¶186. For example, while White discloses adjusting the power allocated to the common bridge logic, Simeral teaches adjusting the bus's operational characteristics (such as voltage, frequency, bandwidth, and data throughput) based on the bus's monitored activity level, also to improve power usage. Ex-1009 ¶¶20, 26. By adjusting the bus based on measured activity level, a POSITA would have recognized that Simeral's teaching would have improved White's power distribution efficiency. Ex-1002 ¶186. Specifically, by implementing

teachings of Simeral about bus workload thresholds, a POSITA would have understood that the combination would have more effectively realized White's goal of avoiding a bottleneck at the bridge logic. Ex-1008, 9:39-53; Ex-1009, ¶¶21, 24; Ex-1002 ¶186.


Additionally, a POSITA would have a reasonable expectation of success in combining the teachings of White and Simeral. For example, a POSITA would have understood how to implement Simeral's teachings related to bus workload thresholds with White's common bridge logic. Ex-1002 ¶187. As White explains, its system uses a "NorthBridge," which "is a common component of many processors." Ex-1008, 2:48-49. Moreover, "*adjusting voltages and frequencies* used by common bridge logic may be included *as part of changing the operating characteristics* (i.e. voltage and frequency) of the logic cores in a multi-core integrated circuit." And Simeral specifically explains that its teachings may be applied, for example, using "controller 610 [that] provides Northbridge functionalities." Ex-1009 ¶42. Simeral also specifically teaches "a controller for *adjusting operational characteristics of the component* for modifying the power consumed by the component to comply with a performance profile." Ex-1009 Abstract. Accordingly, a POSITA would have expected success in implementing Simeral's teachings combined with White's teachings, because White specifically envisions controllers like Simeral's to be used to control its common bridge logic. Ex-1002 ¶¶56, 187.

2. Independent Claim 1

a. Element 1[pre]: A processor comprising:

White discloses the preamble to the extent it is limiting. Ex-1002 ¶¶48-55, 188-190.

White discloses a computer system that includes microprocessor 1100. Ex-1008, 24:30-41. Figure 8 depicts processor 1100 (light blue):



Ex-1008, Fig. 8.

White’s “integrated circuit may include multiple logic cores,” and “integrated circuit 100 may represent a *multi-core processor* and logic cores 120A and 120B may represent *processor cores* that execute instructions supplied to integrated circuit 100.” Ex-1008, 6:41-56. Figure 1 depicts a multi-core processor represented by integrated circuit 100 (light blue):




FIG. 1

Ex-1008, Fig. 1.

b. Element 1[a]: an integrated circuit including:

White discloses Element 1[a]. Ex-1002 ¶¶48-55, 191-192.

White discloses “a system that includes a multi-core *integrated circuit 100*.” Ex-1008, 6:41-56. White’s “*integrated circuit* may include multiple logic cores,” and “*integrated circuit 100* may represent a multi-core processor and logic cores 120A and 120B may represent processor cores that execute instructions supplied to integrated circuit 100.” *Id.* at 6:41-56. Figure 1 depicts integrated circuit 100 (light blue):




FIG. 1

Ex-1008, Fig. 1.

c. Element 1[b]: a first core;

White discloses Element 1[b]. Ex-1002 ¶¶193-194.

White discloses that its “integrated circuit may include multiple logic cores, such as *logic cores 120A* and 120B” and “integrated circuit 100 may represent a multi-core processor and logic cores 120A and 120B may represent processor cores that execute instructions supplied to integrated circuit 100.” Ex-1008, 6:41-56.

Figure 1 depicts logic core 120A (red):





FIG. 1

Ex-1008, Fig. 1.

d. Element 1[c]: a cache memory;

White discloses Element 1[c]. Ex-1002 ¶¶58, 195-196.

White discloses microprocessor 1100 “may [] be coupled to an optional **L2 cache 1128**.” Ex-1008, 25:10-14. L2 cache 1128 (dark red) coupled to processor 1100 (light blue):



Ex-1008, Fig. 8. “L2 cache 1128 may be separate from microprocessor 1100, integrated into a cartridge (e.g., slot 1 or slot A) with microprocessor 1100, or even integrated onto a semiconductor substrate with microprocessor 1100.” Ex-1008, 25:45-49. Because processor 1100 may include integrated circuit 100, a POSITA would have understood that the L2 cache may be part of integrated circuit 100. Ex-1002, ¶¶196; *see also* Ex-1008, 6:41-56, 24:30-41, 25:45-49.

e. Element 1[d]: a communication bus...;

White discloses Element 1[d]. Ex-1002 ¶¶56-58, 197-199.

White discloses a communication bus coupled to the first core. White discloses that its “integrated circuit multi-core integrated circuit or processor may include a *common bridge or interface unit, such as common bridge logic 110, configured to communicate with the logic cores as well as with other elements or peripherals either internal or external to integrated circuit 100.*” Ex-1008, 7:1-7.

White's common bridge logic 110 (green) connects to logic core 120A (red) and other elements:




FIG. 1

Ex-1008, Fig. 1. Because the common bridge logic (green) connect the cores to other elements, a POSITA would have understood that it is a type of communication bus, as described by the '494 Patent. Ex-1002, ¶¶___; Ex-1001, 15:44-45 (“bus 950 may include any known type of interconnect”); Ex-1008, 7:1-7, 2:45-67.

White further teaches that the communication bus connects the first core and cache memory. Given the common bridge logic's function of connecting the logic cores to various elements of integrated circuit 100, it would have been obvious to a POSITA to connect the cores to logic cache 1128. Ex-1002 ¶199. Indeed, whether the L2 cache is on-chip or off-chip, the common bridge logic would have been the

most natural way to connect the L2 cache and logic cores. Ex-1002 ¶199; Ex-1008, 7:1-7, 2:45-67, 6:41-56, 24:30-41, 25:45-49, Figs. 1, 8.

f. Element 1[e]: a core workload monitor...;

White discloses Element 1[e]. Ex-1002 ¶¶200-201.

White discloses a core workload monitor that determines workload for the first core. White discloses power management logic 200 “*configured to monitor and adjust the operating frequency and/or voltage of various portions of integrated circuit 100, such as the logic cores 120 and common bridge logic 110.*” Ex-1008, 14:32-38. The core and power management logic are in the integrated circuit. *Id.*, Claim 1 (“An integrated circuit, comprising: a plurality of logic cores,...and power management logic”). Power management logic 200 (yellow) connects to core logic 120A (red):




FIG. 2A

Ex-1008, Fig. 2A. A POSITA would have understood power management logic 200 as disclosing the claimed “core workload monitor” because the power management logic 200 monitors the operating frequency and voltage of logic core 120A.¹² Ex-1002 ¶201.

g. Element 1[f]: a bus workload monitor...;

White and Simeral disclose Element 1[f]. Ex-1002 ¶¶202-206.

White discloses a bus workload monitor that determines workload for the communication bus. White discloses using power management logic 200 that is “*configured to monitor and adjust the operating frequency and/or voltage of various portions of integrated circuit 100, such as the logic cores 120 and common bridge logic 110.*” Ex-1008, 14:32-38. The common bridge and power management logic are in the integrated circuit. *See id.* Claim 1 (“An integrated circuit, comprising:...a common interface unit...and power management logic”). Power management logic 200 (yellow) connects to common bridge logic 110 (green):

¹² Although White’s power management logic 200 (yellow) also corresponds to the claimed “bus workload monitor” and “balancing control” recited in Elements 1[f] and 1[g], respectively, White explains that the power management logic may be part of a single set of logic that can perform multiple functions. Ex-1008, 14:32-38.




FIG. 2A

Ex-1008, Fig. 2A. A POSITA would have understood power management logic 200 as disclosing the claimed “bus workload monitor” because the power management logic 200 monitors the operating frequency and voltage of common bridge logic 110.¹³ Ex-1002 ¶203. Moreover, White teaches recognizing whether the common bridge logic is “a bottleneck to overall performance,” which a POSITA would have understood as White’s power management logic monitoring the workload of the common bridge logic to make such a determination. *Id.*; Ex-1008, 9:39-53.

Simeral also teaches a bus workload monitor that determines workload for the communication bus. Simeral discloses that its GPIS ports “are power-monitoring ports that can be *configured to monitor and quantify activity levels*,” or “load[s],” for purpose of managing power in components. Ex-1009 ¶¶19, 30. In particular, a POSITA would have understood “GPIS port (‘k’)” to correspond to the claimed bus

¹³ See *supra* n.4.

workload monitor since it “can *monitor activity associated with bus 130.*” *Id.* ¶19; Ex-1002, ¶204. GPIS port (“k”) 191 (yellow) connects to bus 130 (green):




FIG. 1A

Ex-1009, Fig. 1A. “[A]ctivity level” refers to “a quantifiable load associated with a component, where the amount of load correspondingly affects consumption of power by either that component or another component.” *Id.* ¶20.

A POSITA would have been motivated to implement the teachings of Simeral’s bus-monitoring GPIS port with White’s power management logic and common bridge logic, particularly given White’s observation that the bridge logic may consume significant power and contribute to exceeding the thermal power limit for the circuit. Ex 1002 ¶205. The combined White-Simeral teachings would have

been implemented by applying Simeral's technique of measuring the bus's activity, corresponding to the bus's power consumption, in White's power management logic 200, allowing it to effectively monitor the activity of the bus in real-time and make power-management decisions about adjusting operating frequency and/or voltage of common bridge logic 110 to balance consumption and performance. Ex-1009, ¶20; Ex-1008, 14:32-38. This functionality would have furthered the goals of White by further enabling White to monitor and control the bridge logic so that it does not bottleneck overall operation. Ex-1008, 9:30-53.

A POSITA would have had a reasonable expectation of success because both White and Simeral disclose a bus, and both disclose a component for monitoring the voltage, frequency, or activity level of the bus, each of which are particularly suited for working together. Ex-1002 ¶206; Section XII.D.1.

h. Element 1[g]: balancing control...;

White and Simeral disclose Element 1[g]. Ex-1002 ¶207. White's power management logic is the claimed "balancing control" and is in the integrated circuit. Ex-1008, Claim 1. A POSITA would have been motivated to combine White with Simeral for the reasons below.

- (1) 1[g.ii]: ...to dynamically tune power allocation...based on a power limit for the integrated circuit..., the power limit corresponding to a maximum thermal dissipation capacity for the integrated circuit

White discloses 1[g.ii]. Ex-1002 ¶¶208-210. White's power management logic controls and adjusts the operating characteristics of the cores and common bridge logic. “[P]ower management logic 200 may be configured to control the operating characteristics, such as the operating voltage and/or *operating frequency*, of each of multiple logic cores 120 of integrated circuit 110,” and it “may be configured to *adjust the* voltages and/or *frequencies supplied to common bridge logic 110.*” Ex-1008, 8:61-9:32, 14:32-38. Power management logic 210 (yellow) connects to logic core 120A (red) and common bridge logic 110 (green):




FIG. 2A

Ex-1008, Fig. 2A.

White teaches that these adjustments are based, at least in part, on the temperature of the integrated circuit—in particular if the integrated circuit reaches a

maximum operating temperature: “[I]f the operating temperature for either an individual logic core, or for the entire integrated circuit, is above *a preferred, maximum, or safe, operating temperature*, temperature control logic 210 may be configured to reduce the operating frequency for one or more of the logic cores of the multi-core integrated circuit in order to reduce the overall operating temperature of the system.” Ex-1008, 16:7-13, 2:22-26, 12:59-63. A POSITA would have understood White’s reference to the circuit’s maximum operating temperature as corresponding to the claimed power limit. Ex-1002 ¶209. This is because Claim 1 defines the power limit as the maximum thermal dissipation capacity. A POSITA would have understood that these quantities are the same because operating a circuit above a thermal dissipation capacity may result in over-heating and thus corresponds to White’s teaching that its temperature control and power management may both be part of a single logic for monitoring and adjusting frequency and/or voltage of components to stay within the system’s overall power and temperature limits. Ex-1008, Abstract, 14:32-38; Ex-1002 ¶209.

Moreover, White discloses that such adjustments may apply to the common bridge logic because “adjusting voltages and frequencies used by common bridge logic may be included as part of changing the operating characteristics (i.e., voltage and frequency) of the logic cores in a multi-core integrated circuit.” Ex-1008, 2:22-61, 7:46-56.

- (2) **1[g.i] and 1[g.iii]: ...adapted to receive the bus workload from the bus workload monitor...to dynamically tune power allocation...based on... a comparison between the bus workload and a bus workload threshold**

White combined with Simeral disclose 1[g.i] and 1[g.iii]. Ex-1002 ¶¶211-216.

White teaches dynamically tuning power allocation between the cores and common bridge to avoid bottlenecking at the bridge. For example, “allowing a common bridge that interfaces to multiple logic cores to operat[e] at a voltage and/or frequency independent from each of the logic cores may allow the cores to operate at a lower voltages and frequencies while operating the common bridge at a higher voltage and/or frequency than the logic cores, *thereby minimizing any potential bottlenecking at the common bridge.*” Ex-1008, 7:46-56. A POSITA would have understood “bottlenecking” as referring to a comparison of the bridge’s monitored frequency and/or voltage with its operating capacity, because when the bridge’s monitored frequency exceeds its operating capacity, it causes a bottleneck to overall performance. Ex-1002 ¶¶42-47, 211; Ex-1008, 9:39-53. A POSITA implementing White’s workload-related teachings would have turned to additional references for further implementation details of dynamic tuning of power allocations, like Simeral, and a POSITA would have been motivated to combine White with Simeral as explained in Section XII.D.1.

For example, like White's power management logic, Simeral's data path controller 100 and its power management controller 106 disclose the claimed balancing control. Ex-1002 ¶212. Power management controller 106 of data path controller 100 receives feedback signals from GPIS k 191 (yellow):




FIG. 1A

Ex-1009 Fig. 1A. The GPIS ports “send [] feedback signals back to power management controller 106 for purposes of power management.” Ex-1009 ¶30.

Simeral further discloses data path controller 100 as dynamically tuning power allocation because it “manage[s] power consumption for any component” and “react[s] to power-related events in real-time or (nearly so).” Ex-1009 ¶¶20, 23. It does so through power management controller 106, which “can selectably monitor

and adjust operational characteristics of various components,” such as power consumption. *Id.* ¶¶30, 20.¹⁴ Such real-time power adjustment between components is dynamically tuning power allocation. Ex-1002 ¶213. Indeed, this is the same type of real-time power tuning discussed in the ’494 Patent where “a dynamic balance is continuously stuck [sic] between devices.” Ex-1001, 7:56-61.

Simeral further discloses such dynamic power allocation based on a comparison between the bus workload and a bus workload threshold, by increasing/decreasing operational characteristics such as the “bandwidth of a bus.” Ex-1009 ¶21. For example, recognizing that “the amount of data passing over the bus at any given instance represents an activity level,” Simeral teaches that data path controller 100 manages power by “*monitoring the activity level of the bus (e.g., the data transfer rate of the bus),*” and “*if the data rate is below some threshold (i.e., low amounts of data are passing over the bus),* then data path controller 100 can either decrease (i.e., ‘throttle down’) the bus width or shut down one or more bus controllers, thereby conserving power that would otherwise be expended by maintaining a higher bandwidth.” *Id.*

¹⁴ Simeral teaches that “operational characteristics” are functional attributes of a component that can be controlled to alter the activity level of the component” such as “voltage, frequency, temperature, bandwidth, or any other characteristic that is measurable and can be altered or adjusted to influence power consumption.” Ex-1009, ¶¶20, 31.

Moreover, after determining the activity level at which a component is operating, “power management controller 106 *compares that monitored activity level to a threshold activity level* set forth in the performance profile.” *Id.* “If the monitored activity level is in concert with the performance profile, then no action need be taken to alter the activity level as monitored[,] [b]ut if the monitored activity level falls out of compliance, then power management controller 106 initiates corrective action.” *Id.* Thus, a POSITA would have understood comparing the bus activity level to the bus threshold activity level as disclosing the claimed comparison between the bus workload and bus workload threshold and managing power based on that comparison. Ex-1002 ¶215.

The combined White-Simeral teachings would have been implemented by applying Simeral’s technique of comparing the bus activity level to a threshold activity level in White’s power management controller. A POSITA would have been motivated to do so because such a comparison would be an effective way to accomplish a goal of White to minimize potential bottlenecking at the common bridge. *See* Ex-1008, 9:25-53.

(3) 1[g.iv]: ...maintain operation...above a low limit.

White discloses 1[g.iv]. Ex-1002 ¶¶217-218. White discloses reducing the power consumption of either the core or the communication bus to above a low limit

to maintain operation of that core or communication bus. As discussed above, White teaches reducing the operating characteristics of the cores for *inter alia* power management and temperature management. This reduction maintains White above a low limit, which corresponds to shutting down its components. Ex-1008, 2:22-34. But White teaches away from shutting down its components, instead preferring reducing frequency/voltage to save power and decrease the temperature “***rather than simply shutting down a processor that is operating too warmly*** (as is generally done in the prior art).” *Id.* at 2:22-34.

White further teaches lowering the frequency/voltage of the various components to some minimum point but not shutting them down. For example, White discloses that “the operating frequency and/or voltage of a single logic core may be reduced through various stages until ***reaching some minimum level*** and then the operating frequency and/or voltage of one or more other logic cores may be adjusted through those same stages until the operating temperature of integrated circuit 100 is reduced to a level at or below a specified maximum operating temperature.” Ex-1008, 22:15-22.

3. Dependent Claim 3

White-Simeral disclose Claim 3. Ex-1002 ¶¶42-47, 219-225.

White teaches operating at its power limit and increasing or decreasing (or setting an upper limit of) frequency of its core and bus depending on a comparison

of the bus workload compared to its threshold. White teaches that its power management logic may balance voltage and frequency among components when its integrated circuit is operating at a maximum operating temperature. For example, “[i]f a *current temperature for the processor is determined to be at, above, or near such a maximum operating temperature, one or more logic cores of the processor, or other integrated circuit, may be adjusted* to operate at a lower frequency and/or voltage in order to lower the operating temperature of the processor.” Ex-1008, 2:34-44.

White observes that “logic cores 120A and 120B may consume significantly more power than common bridge logic 110, while at the same time, common bridge logic 110 may be more of a bottleneck to overall performance than any individual logic core.” Ex-1008, 9:25-53. Accordingly, White teaches that “power management logic 220 may be configured to *supply a higher voltage and frequency to common bridge logic 110* to reduce any potential bottlenecking of common bridge logic 110, and therefore to help maintain a high overall performance, *while supplying lower voltages and frequencies to the individual logic cores 120* to reduce overall power consumption.” *Id.*

White further discloses decreasing frequency of the bridge and increasing frequency of a logic, under appropriate circumstances. For example, “temperature control logic 210 may lower the operating voltage and/or frequency for one logic

core that may be operating too warmly, *while raising the operating voltage and/or frequency for another logic core that has a lower operating temperature.*” Ex-1008, 23:53-63. And “temperature control logic 210 may also adjust common bridge logic 110 for compatibility with the new operating characteristics of the logic cores,” which a POSITA would have understood as decreasing the frequency of the bridge. *Id.*; Ex-1002 ¶222. A POSITA would have further understood that by decreasing these units, White also discloses capping their upper limit. *Id.*

A POSITA would have been motivated to look for another reference with obvious implementation details, such as Simeral, and further would have been motivated to combine the teachings of White and Simeral. Simeral teaches that data path controller 100 decreases the bandwidth of the bus, depending on the bus’s activity level compared to its activity level threshold. Ex-1002 ¶ 223. “[B]andwidth of a bus is an operational characteristic” that “can be controlled to either increase or decrease the size of the bandwidth,” depending on “the amount of data passing over the bus at any given instance.” Ex-1009 ¶21. Other operational characteristics include voltages, frequencies, and data throughput on buses. *Id.* ¶¶26, 31. “[I]f the data rate is below some threshold (i.e., low amounts of data are passing over the bus), then data path controller 100 can...decrease (i.e., ‘throttle down’) the bus width.” *Id.* ¶21. A POSITA would have understood decreasing or throttling down bus width when the bus’s data rate is below some threshold as disclosing the claimed

decreasing (or setting an upper limit of) frequency for the communication bus in response to the bus workload being less than the bus workload threshold, particularly since Simeral discloses “frequency” as one of the “operational characteristics” that can be adjusted. Ex-1009, ¶¶26, 31; Ex-1002 ¶223.

Simeral also discloses responding to the bus workload being greater than the bus workload threshold. For example, if the monitored activity level falls out of compliance, then power management controller 106 initiates corrective action, including modifying the behavior of the noncompliant component to within its threshold activity level. Ex-1009 ¶24. A POSITA would have understood out-of-compliance activity level as disclosing the claimed in response to the bus workload being greater than the bus workload threshold. Ex-1002 ¶¶224.

4. Dependent Claim 14

White-Simeral disclose Claim 14. Ex-1002 ¶¶226-227.

White discloses increasing the communication bus frequency by a different amount than the first core frequency decrease. White discloses that its logic cores and common bridge logic may consume different amounts of power. Ex-1008, Abstract. For example, “*logic cores 120A and 120B may consume significantly more power than common bridge logic 110.*” Ex-1008, 9:36-53. Under those circumstances, “power management logic 220 may be configured to *supply a higher voltage and frequency to common bridge logic 110* to reduce any potential

bottlenecking of common bridge logic 110, and therefore to help maintain a high overall performance, while *supplying lower voltages and frequencies to the individual logic cores 120* to reduce overall power consumption.” *Id.* at 9:36-53. Because an increase in frequency to the bridge would be offset by corresponding decreases to multiple cores, a POSITA would have recognized that the communication bus frequency could be increased by a different amount. Ex-1002, ¶¶42-47. Moreover, White’s disclosure that “the common bridge logic may be configured to operate at voltages and frequencies independent of current operating voltages and frequencies of the logic cores” would lead a POSITA to have understood that a change to the frequency of one does not require a change by the same amount to the other. Ex-1008, 2:53-67; Ex-1002 ¶¶42-47, 227. Accordingly, a POSITA also would have understood that, when supplying a higher frequency to the common bridge while supplying a lower frequency to a core, the amounts of the common bridge’s frequency increase and core’s frequency decrease may differ and need not be the same amount. Ex-1002 ¶227.

5. Dependent Claim 15

White-Simeral disclose Claim 15. Ex-1002 ¶¶48-55, 228-231.

White discloses increasing frequency for a second core because its “integrated circuit may include multiple logic cores, such as *logic cores 120A and 120B*” and that “each logic core may operate at a voltage and/or frequency independently from

the voltage and/or frequency at which other logic cores...may operate.” Ex-1008, 6:41-56. White further teaches that its “power management logic 200 may be configured to control the operating characteristics, such as the operating voltage and/or operating frequency *of each of multiple logic cores 120* of integrated circuit 110.” Ex-1008, 8:59-65, 9:54-10:49, 20:2-18, Cl. 16. Figure 2A depicts logic cores 120A (red) and 120B (blue):




FIG. 2A

Ex-1008, Fig. 2A.

White further discloses the second core and the communication bus having power planes tied together because the same frequency regulatory supplies clock frequency to all cores and the common bridge logic. For example, a multi-core integrated circuit may include “frequency regulator 140, configured to supply clock

frequency signals to logic cores, such as logic cores 120A and 120B,” and “frequency regulator 140 may also be configured to provide clock frequency signals to common bridge logic 110.” Ex-1008, 7:31-38. Frequency regulator 140 (tan) connects to logic core 120B (blue) and to common bridge logic 110 (green):

A POSITA would have understood that circuit components regulated by the same frequency regulator as having power planes tied together, because the same regulator supplies the same clock frequency signals to each of the circuit components, as in the '494 Patent. *See* Ex-1001, 17:60-62; Ex-1002 ¶231. Additionally, a POSITA would have understood White’s disclosure that “power management logic 200 may adjust the voltages and/or *frequencies supplied to common bridge logic 110 in step with the operating* voltages and/or *frequencies of the logic cores*” as teaching power planes tied together. Ex-1008, 9:32-36; Ex-1002 ¶231.

XIII. THERE IS NO BASIS FOR ANY DISCRETIONARY DENIAL

There is no basis for discretionary denial under §325(d) or §314(a). None of the prior art references asserted in this Petition appear in the file history of the '494 Patent or parent application. Nor has any IPR or other post-issuance challenge ever been made against the '494 Patent. Additionally, the only other proceedings relating to the '494 Patent with scheduled hearing dates are in the ITC (the concurrent District court litigations have been statutorily stayed pending the ITC proceedings).

Under current USPTO guidance, concurrent ITC proceedings are not a basis to discretionarily deny IPR institution. *See* USPTO Memorandum by Katherine Vidal, June 21, 2022, *Interim Procedure For Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation*, at 5-7.

XIV. CONCLUSION

The unpatentability grounds presented above are reasonably likely to prevail, *inter partes* review should be instituted, and Claims 1, 3, 12, 14, and 15 should be cancelled as unpatentable over the prior art.

Respectfully Submitted,

/s/ Ben Haber
Ben Haber (Reg. No. 67,129)

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner certifies that this petition includes 13,995 words, as measured by Microsoft Word, exclusive of the table of contents, mandatory notices under § 42.8, certificates of service, word count, and exhibits.

CERTIFICATE OF SERVICE (37 C.F.R. § 42.6(e)(1))

The undersigned hereby certifies that the above document was served on February 27, 2023, by filing this document through the Patent Trial and Appeal Board P-TACTS System, as well as delivering a copy via express mail upon the following attorneys of record for the Patent Owner and Petitioner:

KINNEY & LANGE, P.A.
312 South Third Street
Minneapolis, MN 55415

A courtesy copy was sent to the below counsel via electronic mail:

Andrea Leigh Fair
Charles Everingham, IV
Claire Abernathy Henry
1507 Bill Owens Parkway
Longview, TX 75604
Phone: (903) 757-6400
Fax: (903) 757-2323
Email: andrea@wsfirm.com
Email: ce@wsfirm.com
Email: claire@wsfirm.com

Michael T. Renaud
Mint Levin Cohn Ferris Glovasky & Popeo PC
One Financial Center, Suite 100
Boston, MA 02111
Phone: (617) 542-6000
Fax: (617) 542-2241
Email: mtrenaud@mintz.com

Brian E. Farnan
Michael J. Farnan
Farnan LLP
919 North Market Street, 12th Floor

Wilmington, DE 19801
Phone: (302) 777-0300
Fax: (302) 777-0301
Email: bfarnan@farnanlaw.com
Email: mfarnan@farnanlaw.com

Deanna Tanner Okun
AMS Trade LLP
1133 Connecticut Avenue NW, Suite 1200
Washington, DC 20036
Phone: (202) 467-6300
Fax: (202) 466-2006
Email: okun@adduci.com

Dated: February 27, 2023

Respectfully submitted,

/s/ Ben Haber

Ben Haber (Reg. No. 67,129)
O'Melveny & Myers LLP
400 South Hope Street, 18th Floor
Los Angeles, CA 90071
Telephone: (213) 430-6000
Fax: (213) 430-6407
E-Mail: bhaber@omm.com
Attorney for Petitioner