

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Mercedes-Benz USA, LLC,
Petitioner,

v.

Daedalus Prime LLC,
Patent Owner.

Case [To be assigned]
Patent 8,898,494

**DECLARATION OF DAVID WYATT IN SUPPORT OF
MERCEDES' PETITION FOR *INTER PARTES* REVIEW OF
CLAIMS 1-7, 13-15, 17-18 OF U.S. PATENT NO. 8,898,494**

Filed on behalf of Petitioner:
Celine Jimenez Crowson (Reg. No. 40,357)
Joseph J. Raffetto (Reg. No. 66,218)
Scott Hughes (Reg. No. 68,385)
Ryan Stephenson (Reg. No. 76,608)
Nicholas Rotz (Reg. No. 75,959)
Helen Y. Trac (Reg. No. 62,250)
HOGAN LOVELLS US LLP
555 13th Street N.W.
Washington, D.C. 20004
Telephone: 202.637.5600
Facsimile: 202.637.5710

TABLE OF CONTENTS

I.	INTRODUCTION	1
A.	Engagement	1
B.	Background and Qualification	1
C.	Information Considered.....	5
II.	background OF THE '494 Patent.....	6
III.	CLAIM CONSTRUCTION AND POSITA DEFINITION.....	9
A.	Claim Construction.....	9
B.	Definition of a Person of Ordinary Skill in the Art.....	10
IV.	UNDERSTANDING OF LEGAL STANDARDS.....	10
A.	Anticipation	10
B.	Obviousness.....	12
V.	THE ELEMENTS IN CLAIMS 1-7, 13-15, 17-18 OF THE '494 PATENT ARE DISCLOSED IN THE PRIOR ART	15
A.	Ground Based on Bose	15
1.	Ground 1 – Bose Discloses or Renders Obvious Claim 1	18
2.	Ground 2 – Bose in view of White and Simeral Renders Obvious Dependent Claims 3, 14-15	39
3.	Ground 3 – Bose in view of Anderson Renders Obvious Dependent Claims 2	57
4.	Ground 4 – Bose (as Modified by Anderson) in view of Naffziger Renders Obvious Dependent Claim 13	68
B.	Grounds Based on White.....	75
1.	Ground 5 – White Renders Obvious Claims 4, 7	75
2.	Ground 6 – White in View of Anderson Renders Obvious Claims 5-6	89
3.	Ground 7 – White (as modified by Anderson) in view of Naffziger Renders Obvious Claims 17-18.....	96
4.	Ground 8 – White in view of Simeral Renders Obvious Claims 1, 3, 14 and 15	99
5.	Ground 9 – White (as Modified by Simeral) in view of Anderson Renders Obvious Claim 2	116
6.	Ground 10 – White (as Modified by Simeral) in view of Anderson and Naffziger Renders Obvious Claim 13	117
VI.	CERTIFICATION.....	118
	APPENDIX A	119

I, David Wyatt, declare:

I. INTRODUCTION

A. Engagement

1. I have been retained on behalf of Mercedes-Benz USA, LLC (“**Mercedes**”) to offer technical opinions relating to U.S. Patent No. 8,898,494 (EX1001) (the “**494 Patent**”) and prior art references relating to its subject matter.¹

2. I have no financial interest in either party or in the outcome of this proceeding. I am being compensated for my work as an expert on an hourly basis at my standard consulting rate, and my compensation is being paid at \$595 per hour. My compensation is not dependent on the outcome of these proceedings or the content of my opinions.

B. Background and Qualification

3. I am the Chief Technology Officer of CardWare, Inc. (“CardWare”). I founded CardWare in 2011. CardWare is a FinTech startup that develops multi-function electronic devices capable of performing more secure payments, leveraging Smart Tokenization™, a technology to enhance convenience improve

¹ For brevity, when providing citations herein, I will generally refer to the name of the reference, rather than the exhibit number. The exhibit numbers associated with each reference can be found in Appendix A.

security and address fraud in payment cards and electronic devices, with more secure methods for conveying payment information on the existing payment rails, wherever used at online or in-store electronic payments facilities.

4. I am also the founder and Chief Technology Officer of PixelDisplay Inc. (“PixelDisplay”) since 2015, which develops new semiconductor devices and nanomaterials technologies, for application in enhancing solid-state lighting, computer-display, graphical rendering and computational performance.

5. I’ve been working as an engineer in the field of electrical engineering, computer engineering, computer science for over 35 years. Before founding my own companies, I worked at NVIDIA Corp. for 9 years in both Software & Hardware Groups, including 5 years as Principal Engineer, and 4 years as Distinguished Engineer, which is a role solely appointed by the CEO Jensen Huang. These roles included extensive work on technological solutions to the management of power and thermal challenges, across divisions from Mobile to Laptop to Workstation to Automotive to Data Center. As part of this work, I invented hybrid graphics technology for which several U.S. patents were obtained including U.S. Patent No. 8,570,331, and which is marketed as Optimus™. I also led the development of technologies for optimized power efficiency of combined CPU/GPU configurations, within a fixed thermal power budget, used both in systems from notebooks to data center rack-elements. Again, I was a named

inventor on several U.S. patents for related power management technologies, including U.S. Patent No. 8,826,048, and this technology was combined into marketed product features, such as BatteryBoost™ and MaxQ™.

6. Before NVIDIA, I worked at Intel Corp. for 8 years, including 4 years as Chief Engineer & Software Architect in the Chipset Group, and 4 years as Platform Architect in the Mobile Platform Group (“MPG”). My work at Intel involved Mobile Graphics Technology, the integration of Graphics Processors (including “GPU” cores), Chipsets (including the Memory Controller Hub “MCH”) and Processors (including “CPU” cores). For example, I worked on thermal and power management for semiconductor devices, including processors and chip packages. This included foundational work on what later became Intel’s Turbo Boost™ technology, that enabled optimal CPU/GPU performance in combined packages through coordinated management of device-level and system-level power within the available cooling capabilities.

7. I studied computer science at the University of Queensland, Australia, and also studied electrical engineering and computer electronics at the South Brisbane College of TAFE, Queensland, Australia.

8. I am a member of Institute of Electrical and Electronics Engineers (“IEEE”). I am currently serving in the Circuits and Systems (“CAS”) group, as a lead member of the Domain-Specific Architecture Steering Committee (“DSASC”)

for future AI Processor Architecture definition. I am also a member of the Society for Information Display (“SID”), an international organization for Display and Touch Technologies. Additionally, I am a member of US Payments and Security Technology Alliance, which is an international Security and Payments Industry Organization. I have also served as the General Chair of the ValleyML AI Conference, an international conference on advances in Artificial Intelligence and Machine Learning technologies, where I have led the conference tracks and dissemination of information on hardware and processors technologies advances for AI.

9. I am a named inventor of roughly 140 issued U.S. patents, including those in the technical fields pertinent to the '494 Patent. I have prepared and delivered presentations on the subjects of Integrated Graphics, and on Power and Thermal Management on behalf of Intel at public events such as Intel’s Developer Forum (“IDF”). I was a regular contributor to the Advanced Configuration and Power Interface (“ACPI”) specification, and the co-author of Appendix-B to the ACPI specification. The ACPI specification is an Open Standard jointly created by Intel and Microsoft and commonly adopted across the computer industry for more than two decades. And I also published a paper on thermal sensors for optimal thermal-power management for memory based on thermal management technologies developed while at Intel, in the Joint Electron Device Engineering

Council (JEDEC) Memory Standards group, in 2005, which was adopted in the SO-DIMM memory module standard, and became part of the memory device specification for DDR3-and-forward.

C. Information Considered

10. I have reviewed the '494 Patent and its prosecution history, as well as the other materials referenced in Appendix A. Counsel has informed me that I should consider these materials through the lens of one of ordinary skill in the art related to the '494 Patent, at the time of the earliest purported priority date of the '494 Patent, and I have done so during my review of these materials. I have been asked to assume, for purposes of this Declaration, that the '494 Patent has a priority date of December 15, 2011 (the “**Critical Date**”).

11. My analysis is based on my years of education, research, and work experience, as well as my investigation and study of relevant materials. In my analyses, I have considered the materials that I identify in this Declaration and those listed in Appendix A.

12. I may rely on these and additional materials to respond to arguments raised by the Patent Owner. I may also consider additional documents and information in further analyses—including documents that may not yet have been provided to me.

13. My review and assessment of the materials provided in this proceeding is ongoing, and I will continue to consider any new material as it is provided. I reserve the right to review, supplement, and amend my analyses based on new information and on my continuing review of the materials already provided.

II. BACKGROUND OF THE '494 PATENT

14. I understand the '494 Patent was filed on February 16, 2012, and issued on November 25, 2014.

15. I have assumed, for purpose of this declaration, that the '494 Patent is entitled to the Critical Date of December 15, 2011.

16. The '494 Patent alleges that in the “prior art,” an integrated circuit with multiple devices, such as multiple processors usually implement power savings policies individually and separately as to each device, which limits the overall performance. ('494 Patent, 1:43-57.) The '494 Patent purports to provide a method for efficiently balancing power and performance between processing cores and/or a communication bus based on measured workloads so as to achieve maximum overall performance, particularly when the integrated circuit is optimized to operate within thermal power and power-consumption limits. (*Id.*, Abstract, 7:40-61, 15:4-21.) The '494 Patent begins with various examples for dynamically balancing power and performance between processing elements, such

as one CPU or GPU core or a collection of CPU or GPU cores. (*Id.*, 9:26-46.)

However, its claims are directed to balancing power/performance between one or more cores and a communication bus. (*Id.*, 15:4:21.) Figure 9 depicting an integrated circuit including multiple processing cores and a communication bus that implements the claimed method:




FIG. 9

(*Id.*, Fig. 9.)

17. The integrated circuit (“IC”) 900 includes a bus 950 (yellow) which connects GPU cores 910 (purple), CPU cores 905 (blue), and a cache memory 960 (grey). (*Id.*, 15:59-63.) IC 900 may be operating under a power limit, such as a thermal design power (“TDP”) limit. (*Id.*, 18:60-63.) As shown in Figure 9, the GPU cores 910 are associated with a GPU workload monitor 911 (orange), the

CPU cores 905 are associated with a CPU workload monitor 935 (blue), and the communication bus 950 is associated with a bus workload monitor 951 (green). (*Id.*, 16:14-19.) These workload monitors measure the workload of the respective components, determining its activity, level of usage, as well as projected power consumption. (*Id.*, 16:19-23.) The integrated circuit 900 further includes a power management service (“PMS”) 920 to balance power between GPU cores 910, CPU cores 905 and communication bus 950 based on those measured workloads. (*Id.*, 16:57-17:3; Fig. 10.)




FIG. 10

18. For example, the '494 Patent discloses that if the workload of communication bus 950 is too high, such as over the threshold, then the communication bus frequency/power is increased. (*Id.*, 17:4-15.) But if IC 900 has already been operating at its power limit, then “the increase in communication bus

frequency is counterbalanced in the CPU and/or GPU.” (*Id.*) The ’494 Patent also discloses that in this case, the power/frequency of both CPU cores 905 and GPU cores 910 can be capped, or only CPU or GPU cores are capped depending on their respective workload. (*Id.*, 17:16-41.)

19. As another example, the ’494 Patent discloses when GPU cores 910 are fully occupied, PMS 920 determines if the communication bus 950 is capable of accommodating the current workload, such as by comparing its workload to a threshold. (*Id.*, 19:20-26.) If the communication bus 950 is able to accommodate its workload, then the frequency of both the communication bus 950 and the CPU cores 905 is capped. (*Id.*, 19:30-41.) By capping bus frequency and CPU frequency, more power can be distributed to the GPU which is fully occupied and thus is a performance bottleneck. (*Id.*) On the other hand, if the communication bus 950 is not able to accommodate its workload, the cap on the communication bus frequency will be increased so that it can accommodate more work. (*Id.*, 19:26-30.) The ’494 Patent explains that the bus frequency cannot be reduced too low because the GPU still needs to work with the communication bus to obtain data requested from the cache. (*Id.*, 18:7-22.)

III. CLAIM CONSTRUCTION AND POSITA DEFINITION

A. Claim Construction

20. I understand that in an *inter partes* review, claims must be given their ordinary and customary meaning, as understood by one of ordinary skill in the art in light of the specification and prosecution history. I apply that standard in my analysis below to the words and phrases in the claims of the '494 Patent.

B. Definition of a Person of Ordinary Skill in the Art

21. A person of ordinary skill in the art (“**POSITA**”) in the field of the '494 would have had an undergraduate degree in electrical engineering, computer engineering, computer science, or similar disciplines, along with two years of professional experience working with the research, design and/or development of semiconductors, processors, power consumption and thermal power management, and related firmware and software or an equivalent level of skill, knowledge, and experience. The more education one has, the less experience needed to attain an ordinary level of skill. Similarly, more experience in the field may serve as a substitute for formal education.

IV. UNDERSTANDING OF LEGAL STANDARDS

A. Anticipation

22. I understand that a patent claim is invalid when the invention that it claims is not new. To establish that a claimed invention is not new (a.k.a., “not novel”), I understand that one may establish that a single publication, or other reference in the prior art discloses (explicitly or inherently) every element required

in a patent claim (*i.e.*, all features or “limitations” recited in the patent claim). I understand that a reference in the prior art “anticipates” a claimed invention if that reference discloses, either explicitly or inherently, every element of the claim.

23. I understand that “prior art” and “prior art reference” are legal terms of art referring to, for example, devices, methods, and publications that predate the earliest effective filing date of the claimed invention.

24. I understand that a prior art reference that anticipates a claim may disclose an element or limitation of a patent claim expressly or inherently. A prior art reference discloses an element or limitation of a patent claim inherently when the prior art’s disclosure necessarily requires or implies that the claimed element or limitation be present in the process, machine, manufacture, or composition disclosed. I understand there is still considered to be an inherent disclosure even if the author of the reference did not describe or understand the underlying inherent principle. I understand that one may establish that a prior art reference inherently discloses a claimed element or limitation through the use of other reference material or through testing.

25. I understand that the description in a written reference does not have to be in the same words as the claim, but all of the requirements of the claim must be there, either stated or necessarily implied, so that a POSITA looking at that one reference would be able to make and use the claimed invention.

26. I understand that any prior art reference that can be considered for purposes of determining whether it anticipates a patent claim may also be used to determine whether the reference renders that claim obvious, as discussed below.

B. Obviousness

27. I understand that, even if a claim is not fully disclosed in a single prior art reference, the patent claim is invalid if the invention would have been obvious to a POSITA at the time of the invention. In particular, I understand that a patent claim is normally invalid as obvious if it would have been an “ordinary innovation” within the relevant field to create the claimed product or method at the time of the invention.

28. I understand that the relevant portion of pre-AIA Section 103, subsection (a) of the Patent Act states:

“A patent may not be obtained through the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.”

29. I understand that, by way of example only, a claimed invention is obvious if:

it combines prior art elements according to known methods to yield predictable results;

it simply substitutes one known element for another to obtain predictable results;

it uses a known technique to improve similar devices (methods, or products) in the same way;

it applies a known technique to a known device (method, or product) ready for improvement to yield predictable results;

it was “obvious to try” in that the inventor chose from a finite number of identified, predictable solutions, with a reasonable expectation of success;

known work in one field of endeavor prompted variations of it for use in either the same field or a different one based on design incentives or other market forces and those variations were predictable to one of ordinary skill in the art; or

some teaching, suggestion, or motivation in the prior art led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

30. When considering obviousness, I understand that I am to: (i) decide the level of ordinary skill in the field that someone would have had at the time the

alleged invention was made; (ii) determine the scope and content of the prior art; (iii) determine what differences, if any, existed between the prior art and the Asserted Claims; and (iv) consider objective evidence of non-obviousness (also known as secondary considerations). Further, when considering obviousness, I understand that it is not necessary to seek out precise teachings, and it is permissible to consider the inferences, common sense, and creative steps that a POSITA (who is considered to have an ordinary level of creativity and is not an automaton) would employ.

31. I understand that objective evidence relevant to the issue of obviousness may also be considered. This type of evidence is sometimes referred to as “secondary considerations,” and may include evidence of commercial success, long-felt but unsolved needs, failure of others, and unexpected results. I understand that any secondary evidence must have a nexus to the relevant claims. For example, I understand that commercial success must have a nexus to features of the alleged invention not disclosed in the prior art, and in particular the prior art references which support an obviousness theory. In other words, I understand that commercial success is material only if it comes from the merits of the claimed invention beyond what the prior art disclosed. With respect to secondary considerations, I understand that Patent Owner bears the burden of proof, and I

have seen no evidence to date that any secondary considerations would establish non-obviousness.

V. THE ELEMENTS IN CLAIMS 1-7, 13-15, 17-18 OF THE '494 PATENT ARE DISCLOSED IN THE PRIOR ART

32. I have been asked to provide an analysis as to whether the elements of Claims 1-7, 13-15, 17-18 of the '494 Patent (the “**Challenged Claims**”) are disclosed in the primary prior art references: U.S. Patent No. 7,421,601 to Bose (“**Bose**”) (EX1005) and U.S. Patent No. 7,263,457 to White (“**White**”) (EX1006). I was also asked to consider these references in view of certain additional prior art for some of the limitations of the Challenged Claims, including U.S. Patent No. 8,601,300 to Anderson (“**Anderson**”) (EX1007), U.S. Patent Application No. 2009/0150689 (“**Simeral**”) (EX1008), and U.S. Patent No. 8,510,582 to Naffziger (“**Naffziger**”) (EX1009).

33. My analysis of these prior art references relative to the elements of the Challenged Claims—specifically, how and where the prior art references disclose the limitations of the challenged claims—is provided below. The citations that I have included are not intended to provide an exhaustive list, but rather provide examples of how the references disclose or teach the elements of these claims.

A. Ground Based on Bose

34. I have reviewed Bose. Bose is directed to “dynamic management and control of power and performance for an integrated circuit chip.” (Bose, 1:16-18.) Bose was filed on February 17, 2006 and issued on September 2, 2008.

35. My review of the '494 Patent file history shows that Bose was not considered during prosecution of the '494 Patent.

36. Bose discloses “a power/performance monitoring and control unit (PMCU)” that can improve instruction throughput and power management for an integrated circuit chip that operates under a power limit. (Bose, 1:17-22; 4:19-29.) The PMCU is designed to dynamically manage chip power “by exploiting the variability within a workload and between different workloads.” (*Id.*, 4:30-46.) Figure 1 depicts a high-level overview of PMCU organization in a system 100 with a multi-core chip 104. (*Id.*, 5:61-67; Fig. 1.)




FIG. 1

37. As shown in Fig. 1, system 100 comprises cores 112-1, 112-2 and 112-R (blue), a global PMCU 110 (red), interconnect 108 (yellow) and L2/L3 cache hierarchy 106 as well as local PMCUs 118-1, 118-2, 118-R, 114, 116 (green). (*Id.*) The local PMCUs 114, 116, 118-1, 118-2, 118-R monitor activity in each core 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and interconnect 108, such as “cache miss rates, branch misprediction rates, instructions per cycle (IPCs), etc.” (*Id.*, 6:56-62.) System 100 adopts a hierarchical power management scheme where “each local PMCUs 118-1, 118-2, 118-R, 114, 116 reports upwards to the on-chip global PMCU 110 through links or connections 126-1, 126-2, 126-R, 122, 124.” (*Id.*, 6:18-37.) The global PMCU 110 is given an overall chip 104 power budget based on thermal packaging, power supply and other considerations, and is


responsible for assigning the budget to each core 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106, and interconnect 108 with a goal to maximize performance. (*Id.*, 12:18-31.)

1. Ground 1 – Bose Discloses or Renders Obvious Claim 1

38. As viewed by a POSITA, Bose discloses or renders obvious Claim 1.

1[pre]: A processor comprising:

39. Bose discloses Claim 1[pre] to the extent the preamble limits the scope of the claim. Bose discloses “[a] power/performance monitor and control unit that can maximize instruction throughput for a *processor* chip for a given power envelope.” (Bose, 1:64-66.) Particularly, Bose discloses a microprocessor system 100, i.e., a processor that includes a multi-core chip 104 and incorporates its power-performance monitor and control unit (“PMCU”) in Fig. 1. (*Id.*, 4:21-24, 5:61-67, Fig. 1.)




A processor

FIG. 1

1[a]: an integrated circuit including:

40. Bose discloses Claim 1[a]. For example, Bose discloses the microprocessor system 100 includes a multi-core chip 104. (*Id.*, 5:61-67; Fig. 1) Bose further discloses “[t]he chip or chip set as described herein is preferably an *integrated circuit* chip.” (Bose, 5:28-29.) Thus, Bose discloses a processor including an integrated circuit, such as the multi-core chip 104:



A processor

FIG. 1

(*Id.*, Fig. 1 (annotated)).

1[b]: a first core;

41. Bose discloses Claim 1[b]. For example, Bose discloses a multi-core chip 104 which includes cores 112-1, 112-2, 112-R. (Bose, 5:61-67, 6:18-19; Fig.

1.) Thus, Bose discloses a first core, such as core 112-1:




FIG. 1

(*Id.*, Fig. 1 (annotated).)

1[c]: a cache memory;

42. Bose discloses Claim 1[c]. For example, Bose discloses the multi-core chip 104 including an “L2/L3 cache hierarchy” 106. (*Id.*, 5:61-67, 6:18-19; Fig. 1.)

In other words, the multi-core chip 104 as disclosed in Bose uses a multi-level cache memory system that includes Level 2 cache memory and Level 3 cache memory, with Level 1 cache being inside the core and the Level 2 and Level 3 cache being a separate unit from the cores, as is well-known to those familiar with cache architectures and was common practice in the art of the time. The “L2/L3 Cache Hierarchy” identified in Bose is well-known term in the art representing an

assumed hierarchical arrangement of cache memories, including Level 2 cache and Level 3 cache memory. Bose also discloses that embodiments of its invention “may include at least one processor coupled directly or indirectly to memory elements,” including “*cache memories* which provide temporary storage of at least some program code to reduce the number of times code is retrieved from bulk storage during execution.” (*Id.*, 5:16-24.) Bose not only discloses a cache memory, it also identifies the same multi-level hierarchy in Fig. 1 as that in the ’494 Patent. The system in the ’494 Patent comprises a first level D-Cache (a data cache) and data translation buffer 150, 151 (i.e., Level 1 cache memory) inside the processor cores, together with higher-level cache hierarchy 110 (i.e., Level 2 and Level 3 caches) being shared between the processor cores, as shown in its Fig. 1.




FIG. 1

(*Id.*, Fig. 1 (annotated).)

1[d]: a communication bus coupled to the first core, the communication bus to connect the first core and the cache memory;

43. As viewed by a POSITA, Bose discloses Claim 1[d]. For example, Bose discloses the multi-core chip 104 includes “interconnect” 108. (Bose, 5:61-67, 6:56-60, 7:4-11, 7:18-23, Fig. 1). From the arrangement in the Fig 1, it is clear the “interconnect” 108 is a major communication element as it has its own PMCU to monitor its power which indicates that it is at least as significant as other major elements such as the processing cores and the L2/L3 cache hierarchy. It is also clear from Fig 1 that the function of interconnect 108 is to interconnect the many cores (112-1 ... to 112-R) with the shared “L2/L3 Cache Hierarchy,” as discussed

above. In the '494 patent, this interconnecting functional unit is termed “a communication bus” and the functional synonym is re-enforced where the '494 Patent discloses that “[c]ommunication bus 950 includes *any known bus or interconnect.*” ('494 Patent, 15:38-39.)




FIG. 1

(*Id.*, Fig. 1 (annotated).)


44. A POSITA would have understood that the interconnect 108 is the communication interconnect coupled to the cores, including the first core 112-1, and connects some or all of the cores 112-1...112-R and the L2/L3 cache hierarchy 106. The term “interconnect” is commonly used in a multi-core integrated circuit to describe the infrastructure bus, fabric, or connection between processor cores

with the non-core (or “un-core” in Intel terminology) units, such as cache memories. Further, Bose discloses that the interconnect 108, the first core 112-1 and the L2/L3 cache hierarchy 106 are all “core element[s]” within the multi-core chip 104 that share an overall power budget and subject to the same power management and control system. (*Id.*, 5:64-67, 7:4-28.) Bose also discloses that the power budgeting of these components are controlled by the same global PMCU 110 based on their performance statistics for purpose of achieving maximum system performance. (*Id.*) Because these components are interconnected in the power domain and located in the same integrated circuit, a POSITA would have understood that these elements, including the cores 112, interconnect 108 and L2/L3 cache hierarchy represents a physically interconnected system where interconnect 108 connects at least one core and the cache memories.

45. Bose further discloses that “[i]t can be assumed that each core 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and interconnect 108 is capable of simultaneous multithreading (SMT).” (*Id.*, 6:62-65.) For these components to be capable of simultaneous multithreading, the L2/L3 cache hierarchy 106 has to be shared among the multiple cores, and based on the use of the well-known term “interconnect”, a POSTIA would have understood that the interconnect 108 connects cores, such as the first core 112-1 to the L2/L3 cache hierarchy.

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

46. Bose discloses Claim 1[e]. For example, Bose discloses the microprocessor system 100 includes a local power-performance monitor and control unit (“PMCU”) 118-1 (*i.e.*, a “core workload monitor”) that is associated with core 112-1. (Bose, 3:15-28, 5:61-67, 6:18-20, 6:56-60; Fig. 1.)



Z

(*Id.*, Fig. 1 (annotated).) Bose discloses “one objective of PMCU is to address the challenge of reducing power by exploiting the variability within a workload and between different workloads.” (*Id.*, 4:30-33.) Specifically, Bose discloses that the local PMCU 118-1 functions to “monitor events, activity levels and direct power and/or temperature levels or suitable proxies thereof, within its domain of authority,” and to monitor “performance statistics such as cache miss rates, branch misprediction rates, instructions per cycle (IPCs) etc.” of the core 112-1. (*Id.*, 6:18-27, 6:56-7:3.) The ’494 Patent discloses that monitoring these same things are examples of determining a workload of a core, stating “determining a workload includes any known method for determining activity of hardware, a workload of software, a number of instructions executed or to execute, or any other metric for determining use of a processing element.” (’494 Patent, 16:20-23.) Thus, the local PMCU 118-1 a “core workload monitor”, since its function is to monitor and evaluate performance statistics that reflect the workload of the cores, including the first-core 112-1.

47. Local PMCU 118-1 also determines the workload for the first core.

For example, Bose discloses:

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.

(Bose, 6:56-7:3.) The local PMCU's not only monitor performance statistics, but also report the information to the global PMCU 110, which means the local PMCU's must determine what performance statistics to report, i.e., determines the workload of the components being monitored. Thus, PMCU 118-1 determines a core workload for the first core.

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

48. Bose discloses Claim 1[f]. For example, Bose discloses the microprocessor system 100 includes a local PMCU 116 that is associated with interconnect 108. (Bose, 3:24-28, 5:61-67, 6:18-24, 6:56-60, 7:4-7, Fig. 1.)




FIG. 1

(*Id.*, Fig. 1 (annotated).) Like the local PMCU 118-1 (i.e., the core workload monitor), Bose discloses that PMCU 116 functions 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.”

(*Id.*, 6:18-37.) Bose further discloses that the local PMCU 116 monitors “performance statistics such as cache miss rates, branch misprediction rates,

instructions per cycle (IPCs) etc.” of the interconnect 108. (*Id.*, 6:56-62.) Thus, the local PMCU 116 is a “bus workload monitor” that determines a workload of the bus, as it is monitoring the relevant performance statistics (such as cache misses) that reflect the workload of a communication bus that interconnects the cores to the cache memories.

49. Local PMCU 116, i.e., the bus workload monitor also determines the workload for the communication bus. For example, Bose discloses:

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.

(Bose, 6:56-7:3.) The local PMCUs not only monitor performance statics, but also report the information to the global PMCU 110, which means the local PMCUs must determine what performance statistics to report, i.e., determines the workload of the components being monitored. Thus, PMCU 116 determines a bus workload for the communication bus.

1[g]: balancing control adapted [1] to receive the bus workload from the bus workload monitor and [3] 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, [2]

the power limit corresponding to a maximum thermal dissipation capacity for the integrated circuit,

50. Bose discloses Claim 1[g]. For example, Bose discloses the microprocessor system 100 includes a chip-level global PMCU 110, which functions as “balancing control” of all of the functional sub-units. (Bose, 4:19-29, 5:61-67, 6:18-37, 6:56-7:31, Fig. 1; *see also* 8:11-18, Fig. 2.). In particular, Bose discloses that the local PMCUs monitor information including “power usage and temperature usage” and report this information “to the global PMCU 110.” (*Id.*, 6:56-7:31.) A POSITA would have recognized that the global PMCU receiving the data from the Local PMCUs attached to each sub-unit to optimally balance sub-unit needs (including performance), within thresholds (including package constraints) and budget (including a power limit), for increased system performance (i.e. optimally balanced).




FIG. 1

(*Id.*, Fig. 1 (annotated).)

51. First, Bose discloses balancing control receives the bus workload from the bus workload monitor. Specifically, Bose disclose that local PMCU 118-1 (i.e., the core workload monitor) and local PMCU 116 (i.e., the bus workload monitor) reports upwards to the global PMCU 110 “a summary of monitored results and actions taken” through links or connections 126-1 and 124. (*Id.*, 6:18-34.) As illustrated above, local PMCU 118-1 monitors the workload of core 112-1, and local PMCU 116 monitors the workload of interconnect 108. Thus, Bose discloses the global PMCU 110 (“balancing control”) receives the bus workload from local

PMCU 116 (the “bus workload monitor”) and the core workload from local PMCU 118-1 (the “core workload monitor”).

52. Second, Bose discloses the balancing control dynamically tunes power allocation between the first core and the communication bus based on a power limit for the integrated circuit, and the power limit corresponds to a maximum thermal dissipation capacity for the integrated circuit. Bose discloses that the global PMCU 110 “can dynamically manage chip power between the components ... to meet the **power limit** with the expectation of increased system performance.” (*Id.*, 4:42-45.) Bose also discloses that the power limit is “due to packaging/cooling constraints.” (*Id.*, 4:34-35.) Additionally, Bose discloses “[t]he **peak temperature limit** of the chip ... may be a preset design specification, determined by the package/colling solution engineered for the chip.” (*Id.*, 2:41-43.) A POSITA would have recognized that where Bose describes “packaging/cooling constraints” to fit within “[t]he peak temperature limit of the chip,” Bose is expressing the “power limit” as similarly expressed in the ‘494 Patent. Such power limit is a standard element of design specification of integrated circuits recognizable to any POSITA, commonly expressed in the “thermal dissipation power” or “thermal design power” (“TDP”). The TDP defines both the processor device’s intended maximum thermal dissipation, and package-external cooling capacity required, in order to stay within device operational temperature range.

With respect to the multi-core chip 104, Bose also discloses that the global PMCU 110 is responsible for assigning power to each core element 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and interconnect 108 based on information collected from local PMCUs and “the overall chip 104 power budget.” (*Id.*, 6:56-7:23; *see also* 7:9-11, 7:18-21, 7:25-27, 3:40-44, 7:65-8:10.) Thus, Bose discloses the global PMCU 110 (“balancing control”) will dynamically tune power allocation between the first core (i.e., core 112-1) and the communication bus (i.e., interconnect 108), based on the power limit for the integrated circuit, and the power limit corresponding to a maximum thermal dissipation capacity (i.e., TDP) for the integrated circuit.

53. Third, Bose discloses that the balancing control compares the bus workload and a bus workload threshold in dynamically tuning power allocation between the first core and communication bus. Bose discloses that the global PMCU 100 will “choose the power modes and assign power [budget] for each core element,” including the first core 112-1 and the interconnect 108 based on “the information [provided by local PMCUs 114, 116, 118-1, 118-2, 118-R] and the objective of the system 100.” (*Id.*, 7:4-31.) The power modes reflect optimizations based on power-performance trade-off characteristics. (*Id.*) “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.*) Thus, each power mode represents the current or expected workload of a component operating in this particular work mode. Bose further discloses that the global PMCU 210 includes a Power/Performance Trade-off Table 206, which “includes the *predicted power and performance for each power mode*,” and the global PMCU 210 uses the predicted power and performance to “assist in calculating the power modes and power budgets for each core.” (*Id.*, 7:41-45, 8:35-43.) A POSITA would have understood that the predicted power and performance represents a workload threshold for a given component working in the associated “power mode.”

54. Bose also discloses that when assigning power mode, the global PMCU 100 considers “the *real time* statistics (stored e.g., in the Phase History Table 213)” and “the power/performance trade-off of each mode for each thread in each core (stored e.g., in the Power/Performance Trade-off Table 206).” (*Id.*, 8:11-12, 9:3-24.) The real time statistics include bus workload monitored by local PMCU 116 (the “bus workload monitor”). (*Id.*, 6:56-60, 7:41-61.) (“The global PMCU 210 stores history information periodically supplied from units of *all the local PMCUs 240* in real time in the Phase History Tables 213.”)

55. A POSITA would have understood that Bose compares a workload threshold as exemplified by the predicted power and performance of each power mode and the monitored workload to adjust the power mode of a component. For

example, Bose discloses that when the actual total power exceeds the total power budget, a lower power mode is assigned for a component “if the power budget assigned to the [component] is insufficient and increases to higher power modes if the [components] are *exceeding* the power budget, such that the sum of power usage for all components fits into the overall power budget.” (*Id.*, 9:14-30.) Bose explains that the power mode is adjusted based on “monitored power and performance levels” of the local components, and adjustment of power mode means decreasing or increasing power level “with corresponding decrease or increase in net performance” of a component. (*Id.*, 12:30-41, Fig. 7.) Bose further discloses that “the amount of such change in power or performance is available from an (updatable) power-performance tradeoff table,” such as the Power/Performance Trade-off Table 206 that includes the predicted power and performance for each power mode. (*Id.*, 8:35-47, 12:42-45.) Thus, a POSITA would have understood that in the Bose system, selecting or adjusting power mode for a component entails comparing a workload threshold, such as the predicted power and performance associated with each power mode with real time statistics, such as the monitored workload of a component.

56. Moreover, Bose discloses that “the local PMCU 315 operate with the newly assigned power budget and power modes 310, and report the new real time information to” the global PMCU 210, which will determine if the real time

information reported by the local PMCU indicates “events such as *over* the budget of power, *over* voltage or overheating emergencies” (*id.*, 9:36-43.) and herein again with the words “over” or “exceeds,” Bose is describing comparing a workload threshold with a monitored workload like that in the ’494 patent.

1[h]: [1] wherein a power consumption of one of the first core and the communication bus is to be reduced, [2] 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.

57. Bose discloses Claim 1[h]. First, Bose discloses reducing power consumption of a component, such as core 112-1 (“first core”) or interconnect 108 (“communication bus”). For example, Bose discloses that the global PMCU 110 selects and adjusts power mode for each core element 112-1 (“first core”), 112-2, 112-R, 106, and 108 (“communication bus”) based on the information provided by the local PMCUs 114, 116, 118-1, 118-2, 118-R and the objective of the system 100. (*Id.*, 7:4-13.) And, “[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.*, 7:13-17.) Bose further discloses that when the actual total power exceeds the total power budget, a lower power mode is assigned for a component “if the power budget assigned to the [component] is insufficient.” (*Id.*, 9:14-30.) Bose also discloses that each time the global PMCU sets a core power budget for a component, the local PMCU for the component “will inform

the global PMCU to *assign a lower power mode* to this [component]” “if the core power target is not met.” (*Id.*, 11:48-51.) And, “[t]he local PMCU may decide to dramatically *reduce* activities on the [component] if the power budget is severely insufficient.” (*Id.*, 11:51-53.)

58. Second, Bose discloses that the operations of the components in Bose system, including the cores and the interconnect 108 (i.e., the communication bus) will be maintained over a low limit of being powered off. Bose also discloses that in performing local optimization of a component, i.e., to fit actual power usage to that specified by the assigned power budget and make the component operate in the voltage and frequency specified by the assigned power mode, the local PMCU “assumes that the core is already optimized for power with ideal *clock gating* techniques.” (*Id.*, 11:3-11.) The term “clock-gating” is commonly used in the industry and understood to a POSITA to mean declining to feed a clock signal to a portion of the integrated circuit while that portion is *powered-on*, in order to reduce active power consumed within that portion. When applying the clock-gating technique, the clock signal to the gated portion is alternatively supplied and declined, creating a power-reduced operational state. A POSITA would also recognize that this is typically the lowest power-reduced operational state where the component is active and operational, above the lower limit of inactive sleep states, such as power-off. Thus, a POSITA would have understood that in the Bose

system, Bose assumes power reduction methods where the components will always remain powered and operational, i.e., their power state and operational capability will be maintained, but they may be intermittently clock gated to reduce power consumption.

59. Additionally, Bose discloses

The performance may be targeted for maximization by one or more of the following: (a) 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); and (b) at a given response point, and given a choice of *power mode allocations per core*, make an allocation such as to maximize the chip (or system)—level throughput performance.

(*Id.*, 2:44-51.) Thus, Bose discloses that its invention is meant to achieve performance goals, which could not be achieved with components simply being turned off.

2. Ground 2 – Bose in view of White and Simeral Renders Obvious Dependent Claims 3, 14-15

a. Dependent Claim 3

60. As viewed by a POSITA, Bose in view of White (EX1006) and Simeral (EX1008) renders obvious Claim 3. I have reviewed White. White was issued on August 28, 2007. I have also reviewed Simeral. Simeral was published on June 11, 2009.

3[a]: The processor of claim 1, wherein the balancing control, in response to the integrated circuit operating at the power limit, being adapted to:

61. Bose discloses Claim 3[a]. *See supra* ¶¶50-56 (Claim 1[g]).

3[b]: 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

62. Bose in view of White and Simeral renders Claim 3[b] obvious. White is titled “System and Method for Operating Components of an Integrated Circuit at Independent Frequencies and/or Voltages,” and like Bose is directed to power management of a multi-core integrated circuit. (White, 2:15-44.) White discloses a multi-core integrated circuit which includes a common bridge logic, which is a communication bus that communicates with the multiple logic cores of the integrated circuit, where the operating voltage and/or frequency of each logic core and the common bridge logic can be independently adjusted for power management and temperature control. (White, Abstract, 2:22-44, 6:34-40.) An example of the multi-core integrated circuit is depicted in Fig. 2A:




FIG. 2A

(*Id.*, Fig. 2A (annotated).)

63. As shown in Fig. 2A, IC 100 includes multiple logic cores, such as logic cores 120A and 120B (blue), each of which may be configured to operate at an independent voltage and frequency. (White, 6:43-49, 10:11-24.) IC 100 also includes a common bridge logic 110 (yellow), which communicates with the logic cores and other components, such as bus, peripheral and/or memory subsystem.

(*Id.*, 7:1-7.) White teaches common bridge logic 110 is configured to operate at a voltage or frequency that is different from and independent of the voltage or frequency at which the logic cores 120A and 120B are operating for purpose of power management. (*Id.*, 7:7-11, 7:49-56.)

64. White further teaches that when “common bridge logic 110 may be more of a *bottleneck to overall performance* than any individual logic core,” power management logic 200 may supply a higher voltage and frequency to

common bridge logic 110 “to reduce any potential bottlenecking of common bridge logic 110.” (*Id.*, 9:36-49.) At the same time, the power management logic 200 supplies ***lower voltages and frequencies*** to the individual logic cores 120 to reduce overall power consumption, i.e., capping frequency for the first core. (*Id.*, 9:50-53.) Thus, the system can maintain the optimal operating performance, while not exceeding overall power consumption. (*Id.*) Thus, by disclosing a bus is the bottleneck and preferentially deserving increased operational budget, White discloses both increasing the voltage and frequency for the communication bus, and capping (i.e., limiting) the voltage and frequency for the first core, in response to the bus workload being greater than the bus workload threshold.

65. Simeral further teaches adjusting operational characteristics of a bus so as to modify its power consumption based on a comparison of the activity level of the bus with a threshold. (Simeral, Abstract, [0007], [0021], [0024], [0026].) Simeral is titled “Data Path Controller with Integrated Power Management to Manage Power Consumption of a Computing Device and Its Component” and is directed to managing power consumption for computing devices. (Simeral, Abstract.) Simeral discloses “a data path controller with embedded power management for computing devices” in a computing device that can provide power management functions. (*Id.*, [0007], [0019], Fig. 1.) Simeral’s data path controller 100 includes multiple general-purpose interface signal (“GPIS”) ports that

“monitor and quantify activity levels of a processor or a peripheral device for power management purposes,” including GPIS port (“k”) 191 that monitors activity associated with bus 130. (*Id.*, [0019], Fig. 1A.) Simeral explains that the term “activity level” can represent the amount of workload associated with a component. (*Id.*, [0020].) Simeral further explains that its data path controller can adjust operational characteristics of a component (such as voltages or frequencies) to ensure that its activity level complies with a threshold activity level set forth in a performance profile that specifies permissible power consumption levels for the component. (*Id.*, [0020], [0023]-[0024].)




FIG. 1A

(*Id.*, Fig. 1A (annotated).)

66. As an example, Simeral discloses that the power management controller 106 within the data path controller 100 monitors the activity level of a component, and “compares that monitored activity level to a threshold activity level set forth in the performance profile.” (*Id.*, [0024].) Simeral explains that “[a] performance profile can include a data structure that describes the operational characteristics that the select components should have to comply with a performance profile,” and “[e]xamples of operational characteristics include

voltages, *frequencies*, data throughput on buses, etc.” (*Id.*, [0026].) Simeral further discloses that “[d]epending on the monitored activity level of [a] bus, power management controller 106 ... can adjust the width of the bus to the appropriate activity level, as set forth in one of performance profiles 110.” (*Id.*)

67. Simeral also discloses if it is determined, that the monitored activity level “falls out of compliance,” the power management controller 106 will take actions to modify the behavior of the component, such as adjusting its operational characteristics. (*Id.*, [0020], [0024], [0026].) Specifically, “the power management controller can recalibrate activity levels of other less essential components to compensate for the increased activity level of the noncompliant component, especially when the noncompliant component is essential to meet the performance objective defined by the performance profile.” (*Id.*, [0024].) Simeral provides an example of such recalibration, that is to reduce activity level of a GPU when the computing device is operating under a performance profile that optimizes music play back by a CD-ROM drive. (*Id.*, [0026].)

68. Thus, Simeral teaches in response to the monitored activity level being above the threshold activity level, the activity level of the component can be increased while the activity level of other components can be reduced. Simeral further teaches that “[s]uch recalibration maintains an overall power consumption of the computing device compliant with a global activity threshold or watermark.”

(*Id.*) Simeral also teaches that “[a]ctivity levels can be expressed in terms of operating characteristics associated with a component, such as voltage, *frequency*, temperature, bandwidth, or any other characteristic that is measurable and can be altered or adjusted to influence power consumption.” (*Id.*) Thus, Simeral teaches increase frequency for a first component and reduce frequency for a second component in response to the workload of the first component being less than its workload threshold, and Simeral also teaches the first component could be a bus, and the second component could be a processor.

69. It would have been obvious to a POSITA to apply the teachings of White and Simeral to Bose. As discussed above, Bose discloses that the global PMCU 210 considers the real time statistics in comparison with the predicted power and performance for each power mode for assigning the power mode for the interconnect 108. *See supra* ¶¶50-56 (Claim 1[g]). Simeral discloses a similar idea, i.e., the power management controller compares the monitored activity level of components, i.e., real time statistics with the threshold activity level set forth in a performance profile, and adjust the behavior of the components accordingly so as to maintain the overall power consumption under a limit. Thus, a POSITA would have been motivated to combine the teaching of adjusting the frequency of components so as to modify its power consumption based on a comparison of a monitored workload with a workload threshold, as already contemplated by Bose.

Bose discloses that its local PMCU uses dynamic voltage frequency scaling (DVFS) technique, which indicates that frequency can be increased or decreased when power mode is being adjusted. (Bose, 11:11-15) Thus, a POSITA would have understood that the Bose system is already operable to adjust frequency as taught by White and Simeral and thus, it would have just been implementation of known techniques to improve a similar system.

70. Similarly, White also contemplates adjusting the operating voltages and/or frequencies of individual components so as to maintain the total power consumption of all the components under a power limit. Particularly, both White and Bose share the same structure where multiple cores are connected to the same interconnect so as to share the higher-level cache memory. Thus, A POSITA would have been motivated to combine the teachings of increasing operating frequency of the interconnect while reducing the operating frequency of multiple cores when the interconnect is determined to be the performance bottleneck as this is the same issue the Bose system may encounter and Bose's PMCU organization is similarly adjusting the operating characteristics of the components in order to adjust their power mode.

71. Further, a POSITA would have had a reasonable expectation of success in applying the teachings in White and Simeral to Bose's system. As discussed above, Bose's system shares similar multi-core integrated circuit

structure as that in White. Further, Bose already discloses a bus workload monitor as that in Simeral, and also contemplated comparing the real time data, such as the monitored workload with a workload threshold in order to adjust power mode.

Bose's PMCU organization is also designed to dynamically adjust power mode of each component within a power limit, which includes adjusting the frequency of each component to fit the assigned power mode. Thus, there is no technical difficulty in combining the teachings in White and Simeral, and a POSITA would have had a reasonable expectation of success as these systems are contemplating managing power consumption in similar systems in a similar way.

3[c]: 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.

72. Bose in view of White and Simeral renders Claim 3[c] obvious. Bose as modified by White and Simeral discloses a system to change frequency in response to a comparison between the bus workload and bus workload threshold. *See supra* ¶¶62-71 (Claim 3[b]). Simeral further discloses that in this system, the data path controller 100 monitors the activity level of a bus, such as 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 decrease (i.e., “throttle down”) the bus width.” (Simeral, [0021].) Simeral also discloses that when it is determined the monitored activity level of a component is

not compliant with a threshold, operating characteristics such as voltages, frequencies, data throughput on buses can be adjusted so that the component can comply with its threshold activity level. (*Id.*, [0020], [0024], [0026].) Thus, a POSITA would have understood that when the monitored activity level of a bus, like that in Bose, is below a threshold, the frequency of the bus can be decreased so as to comply with the threshold. Also as discussed above, Simeral discloses that the activity level of some components could be reduced so as to compensate the increased activity level of other components. *See supra* ¶¶62-71 (Claim 3[b]). As an example, Simeral discloses that when the computing device is operating under a performance profile that emphasizes video play-back quality, “the power management controller 106 will adjust activity levels of other components that are subordinate to the GPU.” (Simeral, [0026].) On the other hand, when “graphics pipeline is not essential to” music playback by a CD-ROM drive, the activity level of a GPU core could be reduced. (*Id.*) Thus, a POSITA would have understood that when the activity level of a bus is reduced, the additional power budget could be used to compensate the activity level of another component, such as a processing core that is more essential for the application profile, and by increasing the frequency of the processing core.

73. Bose also renders Claim 3[c] obvious. Bose discloses that the when the actual total power exceeds the total power budget, the global PMCU “assigns a

higher power mode if there is excess power as compared to the power budget, such that the sum of power usage for all components fits into the overall power budget.” (Bose, 10:5-21.) A POSITA would have understood that when the bus workload is lower than the bus workload threshold, it means there is excess power as compared to the power budget assigned to the communication bus, and thus the power mode or power budget of another component, such as the first core 112-1 could be increased. A POSITA would also have understood that when the power mode or power budget of the first core 112-1 is increased, the power mode or power budget of the communication bus (i.e., the interconnect 108) needs to be reduced so that the sum of power usage for all components fits into the overall power budget. Also, a POSITA would have understood that reducing the power mode or power budget of the bus in this situation will not affect the overall system performance as when the bus workload is already lower than the workload threshold. Further, a POSITA would have understood that one way to increase the power mode of a component is to increase the operating frequency of the component, and one way to reduce the power mode is to decrease the operating frequency. Thus, Bose renders this element obvious.

b. Dependent Claim 14

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.

74. As viewed by a POSITA, Bose in view of White and Simeral renders obvious Claim 14. Bose as modified by White (and Simeral) discloses an integrated multi-core processor device with balancing control to increase the communication bus frequency and decrease first core frequency. *See supra* ¶¶50-56 (Claim 1[g]), ¶¶62-71 (Claim 3[b]). White further discloses that in doing so, the operating voltage and/or frequency of each core and the common bridge logic (i.e., a communication bus) within a multi-core integrated circuit can be ***independently*** adjusted for power management and temperature control. (White, Abstract, 2:22-44, 6:34-40.) White further discloses that the operating frequency of the common bridge logic can be different from all the logic cores, such as being higher than the operating frequency of all the logic cores, or being “an integer multiple of the operating frequency of one or more of the logic cores.” (*Id.*, 4:17-30.) Particularly, White discloses that when logic cores 120A, 120B “consume significantly more power than common bridge logic 110,” the power management logic 200 supplies “a higher voltage and frequency to the common bridge logic 110 to reduce any potential bottlenecking of common bridge logic 110,” “while supplying lower voltages and frequencies to the individual logic cores 120 ***to reduce overall power***

consumption.” (*Id.*, 9:36-53.) Because White discloses that the increase in operating voltage and frequency of the common bridge logic is offset by corresponding decreases to the operating voltages and frequencies of *multiple cores*, a POSITA would have understood that the frequency of the common bridge logic could be increased by a different amount than the decrease of the voltage and frequency of *a single core*.

75. Further, as White describes setting a lower voltage and frequency for one unit (the core) while separately setting a higher voltage and frequency for another unit (the bus), White is teaching *independent* hardware control as in the ’494 Patent, which means a change to the operating characteristics of one component does not necessarily require a change of the same amount to the operating characteristics of another component. Thus, White discloses that in response to the system operating at the power limit, the operating frequency of the communication bus is increased by a different amount than the frequency decrease of the first core as the frequencies of multiple cores are decreased to offset the necessary increase of bus frequency to alleviate a bottleneck that is unique to the bus and central to reasonable system performance, and that the operating frequencies of the other components are independently adjusted without a direct correlation.

c. **Dependent Claim 15**

15: The processor of claim 3, wherein [1] the balancing control is to further increase frequency for a second core, [2] the second core and the communication bus having power planes tied together.

76. As viewed by a POSITA, Bose in view of White and Simeral renders obvious Claim 15. As discussed above, Bose discloses a second core 112-2 in addition to the first core 112-1, and the global PMCU (“balancing control”) adjusts power allocation between the first core 112-1, second core 112-2 and interconnect 108 based on the power limit of the system and the workload of each component. *See infra* Claim 2. Bose also discloses that when the actual total power is above the total power budget, “a lower power mode [is assigned] for a core if the power budget assigned to the core is insufficient” and “higher power mode [is assigned to cores] if the cores are exceeding the power budget, such that the sum of power usage for all components fits into the overall power budget.” (Bose, 9:14-30.) Bose also explains that “[b]ased on the power budget and power mode received through connections 523 and 524, the Local PMCU Controller 512 performs local optimizations to fit actual core power usage to that specified by the power budget, and ***operate in the voltage and frequency specified by the power mode.***” (*Id.*, 11:3-7.) In other words, Bose discloses that when the power mode of each component is adjusted, the voltage and frequency of that component will be adjusted accordingly. A POSITA would have understood that when a core is set to

a higher power mode, it means that more of the power budget is allocated to the core, and thus, the operating voltage and/or frequency of the core are increased. Importantly, as Bose discloses this control is being applied locally on a per-core basis, and this power distribution scheme is equally and independently applicable to every core in Bose's system. (*Id.*, 7:4-31, 9:3:30) Thus, Bose discloses the global PMCU 110 ("balancing control") can, and will, further increase (or decrease) the frequency for a second core 112-2 in response to the integrated circuit 104 operating at the power limit.

77. Bose does not explicitly disclose the second core and communication bus having power planes tied together. However, White discloses a multicore processor system in which the second core and the communication bus have power planes tied together. The '494 Patent does not disclose how power planes are tied together. It is well-known in the art that there are a number of ways power planes can effectively be tied together, and thus set the same voltage at the same time to multiple cores within a semiconductor device. A POSITA would have recognized that this can be accomplished in a number of ways in semiconductor devices, for example: 1) by sharing one power-plane between sub-units within the chip, 2) by providing an electrical connection between power-planes across the chip, 3) by externally connecting together the rails electrically such that each internally separated power-planes is ultimately connected to one common external source, or

4) by operationally connecting together, an effectively simultaneous setting of disparate external power sources, where such sources are operable to synchronously set the power for separate power-planes, such a configuration would have no practical difference with a configuration which had the power planes electrically or physically connected together.

78. As discussed above, White discloses a multi-core integrated circuit 100, which includes multiple logic cores, such as logic core 120A (“first core”) and logic core 120B (“second core”), that are connected to a common bridge logic 100 (“communication bus”). (White, 6:1-6, 6:41-56, 7:1-7.) White further discloses that a power management logic 200 is configured to adjust the operating characteristics, such as operating voltage and operating frequency of the logic cores and the common bridge loci 100 “in order to adjust the overall power consumption for integrated circuit 100.” (*Id.*, 8:61-9:53.) Specifically, White discloses that “[i]n one embodiment, 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.” (*Id.*, 9:32-36.) In other words, White discloses that when the operating voltage and/or frequency of a logic core is changed, the operating voltage and/or frequency of communication bus must change in the same way at the same time, which indicates that the logic core and the communication bus are intended to be logically co-controlled in sync

together, identical in result as if the power planes were electrically or physically tied together. Additionally, White discloses that *the same voltage regulator 130* is configured to supply voltages to the logic cores and common bridge logic 110. (*Id.*, 7:23-31; *see also* 8:32-34.) White also discloses that the operating voltages of the logic cores and the common bridge logic 100 are independently adjustable. (*Id.*, 6:41-7:22.) Thus, a POSITA would understand that in White the power planes of the logic cores and the power plane of the common bridge logic are tied together because the same voltage regulator 130 externally connects the internally separated power-planes of these components.

79. Moreover, White discloses that the same frequency regulator 140 is configured to “supply clock frequency signals to logic cores, such as logic cores 120A and 120B,” and also configured to “provide clock frequency signals to common bridge logic 100.” (*Id.*, 7:31-38; *see also* 8:37-40 (“In one embodiment, a clock signal at a single frequency may be coupled to logic cores 120 and each core may generate a specific internal operating frequency derived from this signal.”)) This also indicates power planes of logic cores and the common bridge logic are tied together in the clocking domain, which is the same as the example provided by the ’494 where clock generators are shared between communication bus and CPU cores. (’494 Patent, 17:60-66.)




FIG. 1

(White, Fig. 1 (annotated).) A POSITA would have recognized that, like the power planes in White, the power planes the second core and interconnect (“communication bus”) in Bose could also be tied together. As discussed in White, a POSITA would have recognized that the operating frequencies/voltages of the cores and interconnect in Bose could be adjusted in step so they can be logically controlled in sync, as if the power planes were physically tied together.


3. Ground 3 – Bose in view of Anderson Renders Obvious Dependent Claims 2

2: *The processor of claim 1, further comprising [1] a second core and a second core workload monitor configured to determine a second core workload for the second core, [3] wherein the first core includes a host processing core, the second core includes a graphics processing core, and the communication bus includes a ring interconnect, and [2] wherein the balancing control is further adapted to dynamically tune power allocation between the first core, the second*

core, and the communication bus based on the power limit and at least the bus workload, the core workload, and the second core workload.

80. As viewed by a POSITA, Bose in view of Anderson renders obvious Claim 2.

81. First, Bose discloses a second core and a second core workload monitor. Bose discloses a management hierarchy in the multi-core chip 104 which includes cores 112-1, 112-2, 112-R. (Bose, 5:61-67, 6:18-19; Fig. 1.) As discussed above, Bose discloses an integrated circuit including a first core, such as core 112-1. *See supra* ¶41(Claim 1[b]). Similarly, Bose discloses the integrated circuit includes a second core, such as core 112-2.



(*Id.*, Fig. 1 (annotated).) Bose also discloses the microprocessor system 100 includes a local PMCU 118-2 that is associated with core 112-2. (Bose, 5:61-67, 6:18-20, 6:56-60; Fig. 1.) As discussed above with respect to the first core, the local PMCU 118-2 monitors performance statistics that reflect the workload of the second core 112-2, and thus is a second core workload monitor. *See supra* ¶¶46-47 (Claim 1[e]).

82. Second, Bose discloses that the chip-level global PMCU 110 (“balancing control”) dynamically adjusts power allocation between the first core 112-1, the second core 112-2, and the interconnect 108 (“communication bus”) based on the power limit, workload of the interconnect 108, workload of core 112-1, workload of core 112-2. *See supra* ¶¶50-56 (Claim 1[g]). The local PMCU 118-2 (“second core workload monitor”) functions in the same way as the local PMCU 118-2 (“core workload monitor”) and the local PMCU 116 (“bus workload monitor”), all of which report workload of the respective component upwards to the global PMCU 110. (Bose, 6:18-37, 6:56-62.) As discussed above with respect to Claim 1[g], Bose discloses that the global PMCU 110 “can dynamically manage chip power between the components ... to meet the *power limit* with the expectation of increased system performance.” (*Id.*, 4:42-45.) Bose also discloses that the global PMCU 110 is responsible for assigning power to each core element 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and interconnect 108 based on

information collected from local PMCUs for each component and “the overall chip 104 power budget.” (*Id.*, 6:56-7:23; *see also* 7:18-21, 7:65-8:10.) Thus, Bose discloses global PMCU 110 (“balancing control”) dynamically tunes power allocation between a first core 112-1, second core 112-2 and interconnect 108 (“communication bus”) based on the overall power limit for the integrated circuit chip.

83. Also as discussed above, Bose discloses that in assigning power, the global PMCU 100 will “choose the power modes for each core element,” including the core 112-1, core 112-2 and the interconnect 108 based on “the information [provided by local PMCUs 114, 116, 118-1, 118-2, 118-R] and the objective of the system 100.” (*Id.*, 7:4-17, 24-28.) Bose further discloses that when assigning power mode, the global PMCU 100 considers “the ***real time*** statistics (stored e.g., in the Phase History Table 213)” and “the power/performance trade-off of each mode for each thread in each core (stored e.g., in the Power/Performance Trade-off Table 206).” (*Id.*, 9:3-24.) The real time statistics include bus workload monitored by local PMCU 116, and core workload of the first and second cores monitored by local PMCUs 118-1 and 118-2, respectively. (*Id.*, 7:50-53.) (“The global PMCU 210 stores history information periodically supplied from units of all the local PMCUs 240 in real time in the Phase History Tables 213.”)

84. Third, while Bose does not specifically disclose that the first core includes a host processing core, the second core includes a graphics processing core, and the communication bus includes a ring interconnect, this would have been obvious over Bose in view of Anderson. Anderson, titled “System and Method for Managing Thermal Energy Generation in a Heterogeneous Multi-core Processor” was filed on September 29, 2011, and issued on December 3, 2013. I have reviewed Anderson.

85. Anderson, like Bose, is directed to power management of an integrated circuit with a multi-core processor. Anderson discloses “methods and systems for controlling and/or managing thermal energy generation on a portable computing device that contains a heterogeneous multi-core processor.” (Anderson, Abstract.) Anderson discloses that a chip “may be comprised of one or more distinct processing components generally referred to herein as ‘core(s)’.” (*Id.*, 4:3-5.) These “cores” in Anderson are the same components as the “cores” described in Bose, as in each reference the cores are the processing cores in a multi-core chip. In Fig. 5A, Anderson discloses a chip 102 that includes an applications central

processing unit (“CPU”) 110 with multiple CPU cores, each of which is a host processing core,² such as core 0, core 1 and core N.




FIG. 5A

(*Id.*, 9:66-10:7.), Fig. 5A (annotated).) A POSITA would have understood that the Application CPU cores (i.e., core 0, core 1 and core N) in Anderson’s system are host processing cores as they are the collective central processor of a typical multi-core processor chip running the host operating system and applications for a portable computing device. Anderson discloses the multiple CPU cores in the context of portable computing devices, and refers to the CPU 110 as the

² In the ’494 Patent, the term central processor and host processor are used interchangeably, and refer to a processor other than the graphics processor. (’494 Patent, 1:50-52, 3:15-21.)

“applications CPU.” (*Id.*, 1:19-67, 10:1-33, Fig. 5A.) Anderson also discloses and illustrates in Fig. 5A that the applications CPU 110 includes operating system (“O/S”) module 207. (*Id.*, 10:10-12, Fig. 5A.) A POSITA would understand that the applications CPU 110 is running the host operating system as indicated by element 207, and thus CPU 110 in its entirety could be considered as the host processor, and all cores inside CPU 110 are host processing core.

86. Anderson further discloses that the same chip 102 comprises multiple graphics processors, including a first graphics processor 135A, a second graphics processor 135B, a third graphics processor 135C, and a fourth graphics processor 135D. (*Id.*, 10:30-57.) A POSITA would have understood that when a graphics processor or a graphics processing unit (“GPU”) is integrated into a chip with a CPU, it no longer requires the usual peripheral logic that would otherwise exist as a standalone GPU device. Because the usual peripheral logic such as add-in bus interface, GPU local memory interface or other I/O support, are no longer needed, a graphics processor, or a graphics processing unit (“GPU”) correlates to graphics core when integrated into a combined processor. Thus, Anderson discloses an integrated circuit including a first core which is a host CPU core, and a second core which is a GPU core.

87. Anderson further discloses with respect to Fig. 5B that the CPU cores, i.e., the first core 222, the second core 224 through to the Nth core 230 of the CPU

110 “may implement message or instruction passing via network topologies such as bus, *ring*, mesh and corssbar topologies.” (*Id.*, 11:29-36.) That is, Anderson discloses that CPU cores in a multicore processor can be connected with a communication bus network that is a “ring interconnect.”




FIG. 5B

(*Id.*, Fig. 5B.)

88. It would have been obvious to apply these teachings and include both CPU and GPU cores as well as ring interconnect in the multicore processor system of Bose. Bose and Anderson are in the same field of art, as they are both directed to thermal power management in a multi-core processor. While Bose does not explicitly disclose the different types of cores such as CPU and GPU cores in its system, a POSITA would have understood that the multi-core processor like that in Bose would often have a GPU core if graphics processing is required. A POSITA would have understood that Bose's PMCU organization can easily accommodate such different types of cores by, for example, adding another local PMCU to monitor the workload of a GPU core when a GPU core is present. The GPU core can be treated as other core elements, like CPU cores, cache memories and interconnect. The global PMCU can similarly choose a power mode for the GPU core like it does for the other core elements based on the monitored workload. Thus, a POSITA would have had a reasonable expectation of success in applying the Bose power management mechanism to an integrated circuit with both GPU core and CPU core as no substantial change is required. That is, the system would monitor a GPU core and CPU core in a multicore processor in the same way it describes doing so for cores 112-1 through 112-R. In fact, cores 112-1 through 112-R would simply include cores that were CPU cores and GPU cores.

89. A POSITA would have also been motivated to apply Bose's power management technique in a processor with heterogeneous cores, including CPU and GPU cores, as disclosed in Anderson. As disclosed in Anderson, efficiently balancing thermal energy generation and system performance could be tricky with a system that has heterogeneous processing components "because the various processing components ... are not created equal." (Anderson, 1:64-67.) Because in Bose's hierarchical power performance management, each local component is associated with a local PMCU to monitor its workload and adjust its operating characteristics based on the power mode selected by the global PMCU, the individual features and capability of each component would be taken into consideration and accommodated.

90. A POSITA would have also been motivated to implement a ring interconnect as disclosed in Anderson in Bose's system. A ring interconnect connects all the processing units in a circular arrangement and allow them to exchange data under certain communication protocol. As compared to traditional point-to-point bi-directional or multi-point bus topologies, a ring interconnect can substantially reduce the number of physical interconnects (i.e., bus routings) to exchange data among all the processing units, and thus significantly reduce the semiconductor die area needed for bus routing. As discussed above, the interconnect 108 in Bose's system interconnects the major elements in the multi-

core chip 104, which would typically be located at separate locations inside the chip. As more CPU and GPU cores, caches, interfaces and other sub-units are integrated into ever more complex combined packages, the traditional bus structures (such as bi-directional linear bus, deeply multiplexed bus, or crossbar-switch bus topologies) become too complex and expensive, in terms of the physical area of the semiconductor die that they consume, which could otherwise have been used for increasing performance (e.g. with more cores, larger cores, or larger caches). A POSITA would have recognized, as had occurred during the evolution of large multi-core processors in networking chips (e.g. internet router chips of the late 1990's), that a ring topology bus is not only more efficient in terms of the physical space and ease of routing, it becomes necessary as the cores are ideally physically arranged in a ring around the package. Indeed, a POSITA in the art of the time would have already recognized that a ring-like circular arrangement of cores is not only the optimal physical placement for units on a semiconductor device, but also necessary for achieve highest efficiency in terms of die area and performance. A POSITA would also have had a reasonable expectation of success in implementing a ring interconnect in the Bose system since Anderson discloses such a ring topology implementation, applied in a multi-core processor system as disclosed in Bose, and a ring topology bus is architecturally simpler, with lower engineering complexity, and clearer path to successful implementation.

4. Ground 4 – Bose (as Modified by Anderson) in view of Naffziger Renders Obvious Dependent Claim 13

13: The processor of claim 2, wherein the balancing control is to increase a frequency of the second core and to reduce a frequency of the first core when the second core is to execute a producer thread and the first core is to execute a consumer thread to consume information produced by the producer thread.

91. As viewed by a POSITA, Bose (as modified by Anderson) in view of Naffziger renders obvious Claim 13. I have reviewed Naffziger. Naffziger, titled “Managing Current and Power in a Computing System” was filed on July 21, 2010.

92. Naffziger is directed to a system and method for efficient power transfer on a die that includes two or more computing units (CUs). (Naffziger, Abstract.) Like Anderson, Naffziger discloses a processor with both GPU cores and CPU cores. Specifically, Naffziger discloses with respect to Fig. 1 that a die 102 comprises a processing block 160 which includes computation units 120a-12d, and “[e]ach computation unit 120 may be one of general-purpose processors, graphics processors, special-purpose coprocessors, and so forth.” (*Id.*, 3:21-44.) More specifically, Naffziger discloses in Fig. 8 a processor 700 that includes both general-purpose central processing units (CPUs) 115a, 115b and graphics processing unit (GPU) 170, where GPU 170 comprises a graphics processor core 172 and CPUs 115a, 115b comprise processor cores 112. (*Id.*, 16:45-58, 17:9-13.)




FIG. 8

(*Id.*, Fig. 8 (annotated).)


93. A POSITA would have understood that the CPU cores and the GPU cores as disclosed in Naffziger are the same CPU cores and GPU cores described in the '494 Patent, which can form a consumer-producer relationship with either the GPU cores being the producer and the CPU cores being the consumer or the CPU cores being the producer and the GPU cores being the producer. ('494 Patent, 7:62-8:17, 9:32-39, 15:8-11.) Naffziger discloses that the CPU cores and GPU cores are all connected to packet processing logic 116, which can be “crossbar

switch” as shown in Figure 8. (Naffziger, 17:49-59, Fig. 8.) It is well-known in the art that the crossbar switch is a structure to allow the CPU cores to communicate with the GPU cores. Naffziger further discloses in Fig. 8 that the GPU 170 and CPUs 115 share the same cache memory subsystem 118, and GPU 170 also has direct access to the cache memory subsystem 114 that are integrated into the CPUs. (*Id.*, 17:33-48.) And, because “the GPU 170 may now to able to directly access both local memories 114 and 118 and main memory via memory controller 122,” latency for memory accesses for the GPU 170 may be lowered and higher performance can be achieved. (*Id.*, 17:25-32.) Naffziger also discloses its GPU 170 can be a traditional dedicated graphics rendering device for a personal computer. (*Id.*, 16:57-60, 17:13-16.) Thus, a POSITA would have understood that Naffziger discloses typical GPU and CPUs within a processor that work together to execute software applications in a consumer-producer relationship. For example, the CPU can leverage the GPU as an accelerator, and hands-off the specialized workload for execution through completion on the specialized graphical processor. For example, when rendering a scene (e.g. windows of an operating system GUI, a game, CAD/CAM, or other visualization application) to be displayed on the computer screen, the CPU prepares the graphical primitives (e.g. screen co-ordinates, vertices, textures, images, icons ... etc.) and dispatches them to GPU. In the process of displaying each new frame of the scene, the CPU produces the changes

in the source primitives, which are sent to the GPU to be consumed by the GPU core(s), for rendering and presenting the visual end-result on the screen.

Conversely, the CPU can be the consumer when it is processing the output from the GPU cores, such as when the GPU is encoding a video stream for a CPU to merge with audio and other streams, from which the CPU broadcasts via the network and/or records on local storage.

94. Naffziger further teaches the power management unit 130 selects one discrete power-performance state (P-state) from several given discrete P-states for the cores based on their real-time power consumption (e.g., monitored activity level) when running an application or applications. (*Id.*, 6:33-54.) The selected P-state is between a maximum performance state and a minimum performance state, with the maximum performance state P_0 including a maximum operating frequency and the minimum power state P_4 including a minimum operating frequency, as shown in Fig. 2. (*Id.*, 6:39-43.) The P-states in between P_0 and P_4 include “given scaled values for a combination of the operating frequency and the operational voltage.” (*Id.*, 6:43-46.)



(*Id.*, Figure 2 (annotated).) Naffziger also discloses that the power management unit continuously receives power consumption estimates from the computation unites, and selects new corresponding P-state accordingly, including “determining whether to boost, throttle or maintain a currently selected P-states.” (*Id.*, 8:16-34.) Naffziger explains that “[a] ‘throttle’ of a P-state includes decrementing the currently selected P-state by one P-state to a lower power consumption P-state,” and “a ‘boost’ of a P-state includes incrementing the currently selected P-state by one P-state to a higher performance P-state.” (*Id.*, 7:65-8:2) As shown in Fig. 2, decrementing the currently selected P-state entails reducing the operating frequency of the computation unit, while incrementing the currently selected P-

state entails increasing the operating frequency of the computation unit. As an example, Naffziger discloses that when the activity level of a computation unit is determined to be higher, it receives power credits from a different inactive computation unit, and the power management unit 120 increments the P-state of the computation unit to a higher power consumption P-state. (*Id.*, 9:29-64.)

95. It would have been obvious to a POSITA that when the GPU core is to execute a producer thread and the CPU core is to execute a consumer thread, the GPU core is working with a higher activity level, while the CPU is relatively inactive. Thus, following the teachings in Naffziger, the power management unit 130 will increment the GPU core to a higher power consumption state, i.e., increasing its operating frequency, while decrement the CPU core to a lower power consumption state, i.e., decreasing its operating frequency. Thus, Naffziger renders Claim 13 obvious.

96. A POSITA would have been motivated to combine the teachings of Naffziger with Bose-Anderson. Like Bose, Naffziger is also directed to power management of an integrated circuit with multiple computation/processing units. Particularly, in the Naffziger's system, "the CUs measure a corresponding activity level and convey ... data indicative of the measured levels to a power manager," and the power manager allocates power credits among the CUs based on the corresponding activity levels. (Naffziger, 2:5-16.) When running an application or

applications, the power manager will select a discrete power state for each CU and adjust the power state of each CU based on its real-time power consumption, where each power state optimally affords a different range of operational frequencies and different levels of performance. (*Id.*, 6:33-46.) Bose also discloses that its global PMCU 110 chooses appropriate power mode for each core element based on real-time statistics, such as the monitored activity level of each component, and each power mode specifies operating voltage and frequency. *See supra* ¶¶50-56 (Claim 1[g]). Thus, both Naffziger and Bose disclose a similar power management method based on monitored activity level of each component within the highly integrated multi-core package. A POSITA would have thus understood both systems to be substantially the same except that Naffziger specifically discloses that its system can accommodate many different types of processing cores. And thus, a POSITA would have had a reasonable expectation of success to adjust the frequencies of a GPU core and a CPU core so as to adjust the respective power, as taught by Naffziger in the Bose system as modified by Anderson.

97. Moreover, a POSITA would have been explicitly motivated to include separate per-core control since the workload of each core varies drastically in real-world use according to application characteristics (e.g., graphical application, compute application, word processing application, etc.), whether in a multi-core

CPU or a multi-core CPU plus GPU integrated device. As a POSITA would have recognized, when performing rendering, AI or other highly-parallelized compute workload, most of the power needed for completing the workload must be directed to the GPU core; whereas, a CPU core (without the core control disclosed by White) which dispatches the work to the more-compute-efficient GPU core, will otherwise waste power in a “tight polling loop,” i.e., the CPU core is continuously checking for completion of the threads it dispatched to the GPU core, and thus negatively impacting useful progress of the GPU’s towards workload completion within the shared thermal budget of highly-integrated device. Thus, in a producer consumer relationship, control of the power-state for each core of a highly-integrated multi-core device has always been balanced towards independent control favoring the core needing the power budget, at the moment where it can be most usefully applied, according to the shift of the workload between the inner cores.

B. Grounds Based on White

1. Ground 5 – White Renders Obvious Claims 4, 7

a. Independent Claim 4

98. As viewed by a POSITA, White renders obvious Claim 4.

**4

```
:
```** A non-transitory storage medium including code, when executed, to cause a machine to perform the energy efficiency operations of:

99. White discloses Claim 4

```
 to the extent the preamble is limiting.
```

White discloses a power management logic 200 that is included in “a battery powered computing device, such as a laptop, handheld computer, or PDA,” including a multi-core integrated circuit 100 having logic cores 120 and a common bridge logic 110:




FIG. 2A

(*Id.*, Fig. 2A (annotated), 8:59-9:1.) The power management logic 200 is configured to optimize energy efficiency within the thermal envelope, as it adjusts operating characteristics for one or more of logic cores 120A, 120B and common bridge logic 110 “in order to adjust the overall power consumption for integrated circuit 100.”

(*Id.*, 9:1-18.) This power management logic 200 can be implemented “in hardware,

firmware, *software*, or a combination thereof, that is configured to monitor and adjust the operating frequency and/or voltage of various portions of the integrated circuit 100, such as the logic cores 120 and common bridge logic 110.” (*Id.*, 14:32-38, 16:56-63.) Thus, the power management logic can be stored in a non-transitory storage medium in the computing device in the form of software code, and when the software code for the power management logic 200 is executed, it causes the integrated circuit 100 to perform energy efficiency operations, such as adjusting operating frequency and/or voltage of the logic cores so as to adjust overall power consumption of IC 100.

4[a]: [2] receiving a bus workload for a communication bus and a first processing element workload for a first processing element based on an integrated circuit reaching a limit, [1] the integrated circuit including the communication bus and the first processing element, wherein the communication bus is to connect the first processing element with at least one other device of the integrated circuit;

100. White discloses Claim 4[a]. First, White discloses a multi-core integrated circuit 100, which includes multiple logic cores, such as logic core 120A and logic core 120B:




FIG. 1

(White, Fig. 1 (annotated), 6:1-56.)

101. These logic cores 120A and 120B “may represent processor cores that executes instructions supplied to integrated circuit 100,” and thus are processing elements in the integrated circuit. (*Id.*, 6:41-56.) Integrated circuit 100 also includes a common bridge logic 110 (“communication bus”), which connects logic core 120A and logic core 120B to “other elements or peripherals either *internal* or external to integrated circuit 100, such as bus 150, peripheral 160 and/or memory subsystem 170.” (*Id.*, 7:1-7, Fig. 1.) The ’494 patent discloses that a “[c]ommunication bus [] includes any known bus or interconnect.” (’494, 15:38-39.) Thus, White discloses an integrated circuit including a communication bus

(e.g., the common bridge logic 110) and a first processing element (e.g., logic core 120A), where the common bridge logic 110 connects logic core 120A with at least one other device of the integrated circuit, such as bus 150, peripheral 160 and/or memory subsystem 170. (White, 6:1-56, 7:1-11, Fig. 1.)

102. Second, White discloses that the power management logic receives a bus workload of the common bridge logic 110 and receives a first processing element workload for a first processing element. As noted above, power management logic 200 “is configured to *monitor and adjust the operating frequency and/or voltage of various portions of the integrated circuit 100, such as the logic cores 120 and common bridge logic 110.*” (*Id.*, 14:32-38, 9:39-53, 9:56-62, 10:43-46.) White discloses that power management logic 200 is configured to “*monitor and adjust the power consumption and usage* for the system.” (*Id.*, 18:22-25.) For example, “power management logic 200 may detect *a low power state* and may attempt to reduce the operating [frequency/voltage] for one or more [components] ... to reduce the overall power consumption of the system.” (*Id.*, 15:49-54, 17:29-44; 18:28-39, 20:39-57.) White discloses that “power management logic 200, may determine that logic core 120A is consuming power in excess of that allocated to it,” and “alter its frequency of operation” in response to this determination. (*Id.*, 9:56-62.) White also discloses that the power management logic 200 adjusts the operating voltage/frequency of the logic cores

and common bridge logic 110 when determining that “logic cores 120A and 120B may consume significantly more power than common bridge logic 110. (*Id.*, 9:36-53.) White further discloses that in determining the power consumption of a component, power management logic 200 is configured to “***determine the current output voltage levels*** for a logic core 120 [or common bridge logic 110] by accessing data stored in configuration registers for voltage regulator 130.” (*Id.*, 10:43-46.) The power management logic 200 also similarly receives the operating frequency “in order to adjust the operating frequencies for the logic cores 120 and/or for common bridge logic 110.” (*Id.*, 10:11-24.)

103. The '494 Patent indicates that “workload” is “any [] metric for determining use of a [component].” (*See e.g.*, '494 Patent, 9:47-52.) Thus, a POSITA would have understood the power management logic 200 in White monitoring the operating voltages and frequencies of the logic core 120A and common logic bridge 110 is monitoring the workload of those components. It was well understood that operating frequency and voltage of a component reflect the activity and use, or workload, of that component. Further, a POSITA would have recognized that wherever possible, directly measuring the power consumption of a unit or sub-unit, is in fact the best and most preferable implementation to monitor component usage or workload, since the stated end-goal is component power management, and other indirect workload metrics (such as clock or unit activity)

do not as exactly account for the size of the unit/component (i.e. the number of transistors causing its actual active power consumption) nor it's leakage power when active or inactive.

104. White also discloses that the power management unit 200 receives the operating voltages and frequencies based on the integrated circuit reaching a limit. For example, White discloses that “a system including a multi-core processor that may be operating from battery power and power management logic 200 may be configured to detect when the system is operating on battery power and adjust the operating frequency of [the components].” (*Id.*, 15:54-59.) White further discloses that the power management logic 200 also detects “when the system transitions from battery power to wall power and adjust the operating frequency of [components] to a higher frequency to take advantage of the fact that the system is no longer on battery power.” (*Id.*, 15:59-64; *see also* 18:25-39, 20:6-12, 24:7-29.) Thus, White discloses receiving workload of the components based on the integrated circuit reaching the power supply limit, such as the battery power limit. A POSITA would have also understood that in the White system, the power management logic 200 receives workload information based on the integrated circuit reaching other limit, such as the maximum clock frequency for the given voltage of a lower power state (e.g. requiring the power state to be raised)

105. Additionally, White discloses that the power management logic 200 receives workload information based on the integrated circuit reaching a thermal design power limit. White discloses that the integrated circuit has “*a predetermined maximum operating temperature*”, above which the performance of the device may be erratic or damage may potentially occur. (White, 2:34-38.) White further discloses the operating characteristics of the multiple logic cores and the common bridge logic is adjusted to maintain the operating temperature of the integrated circuit within “*a maximum temperature.*” (*Id.*, 2:39-44, 12:59-13:6) For example, the operating characteristics for each of the multiple logic cores in the integrated circuit may be controlled by a temperature control logic 210 (*id.*, 11:61-64, Figs. 2B-2C), and “temperature control logic 210 and power management logic 200 may both be part of a single set of logic ... 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.” (*Id.*, 14:32-38.) A POSITA would have recognized that a rise in the temperature (of a sub-unit, core, or package) is in fact a direct result of insufficient cooling to remove the thermal power generated at said unit/core/package level. This is a direct result and indication of exceeding a thermal design power (TDP) limit.

4[b]: determining if the first processing element workload indicates the first processing element is a performance bottleneck;

106. White renders Claim 4[b] obvious. White discloses when it is detected that “logic cores 120A and 120B [are consuming] 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,” power management logic 200 supplies a higher voltage and frequency to common bridge logic 110 “to reduce any potential bottlenecking of common bridge logic 110.” (*Id.*, 9:36-49; *see also* 7:46-62.) The power management logic 200 is configured to “adjust the operating characteristics for one logic core to *a lower level* than the operating characteristics of another logic core depending upon the overall power consumption ... of the system.” (White, 9:25-30.)

107. A POSITA would have understood that the primary reason for the multi-core system disclosed in White to separate voltage/frequency states for each core, was to effect independent control of each core so as to optimally balance the system, as was the common architectural goal of the day. And further, that like the common bridge logic, each individual core may also potentially be a bottleneck to the overall performance as indicated by its respective power consumption, particularly when the multiple logic cores are interfaced with the same common bridge logic and sharing the same cache. For example, unproductive overactivity of a unit may overwhelm the common bus or common cache, and deprive other units

of timely access to memory data needed to feed their core and/or cause pipeline stalls which waste power, and also impedes further forward progress, such as may be necessary to exit a deadlock or stall condition. White is clearly identifying that any unit in an integrated multi-core processor chip can become a bottleneck to the overall chip performance, especially when the processing cores share common cache and memory sub-systems, even the common bus itself. Thus, determining whether a logic core is a performance bottleneck is an inevitable design consideration. As White discloses the power consumption of the common bridge logic 110 indicates it is a bottleneck, a POSITA would have found it obvious to determine whether the power consumption (i.e., workload) as indicated by the operating voltages/frequencies of the logic core 120A similarly indicates it is a bottleneck to the overall performance of the system, and hence power management logic 200 could adjust its operating characteristics accordingly to maintain optimal overall performance of the system. Doing so would have required the same determinations and adjustments described for when the common bridge logic is the bottleneck. Instead of supplying a higher voltage and frequency to common bridge logic 110 “to reduce any potential bottlenecking,” the same higher voltage and frequency could be supplied to any core that is determined to be a bottleneck. This would have simply been the application of the same known technique, in the same way, to a similar component. A POSITA would have had a reasonable expectation

of success in doing so, as White is clear that the purpose of the power management logic is to manage power of the system, and that it is already configured to adjust the voltage and frequency supplied to the cores and to the common bridge logic.

4[c]: determining if the communication bus is able to accommodate the bus workload;

108. White discloses Claim 4[c]. For example, White discloses that when it has detected that logic cores 120A and 120B are consuming significantly more power than common bridge logic 110 and “at the same time, common bridge logic 110 may be more of a *bottleneck to overall performance* than any individual logic core,” power management logic 200 supplies a higher voltage and frequency to common bridge logic 110 “to reduce any potential bottlenecking of common bridge logic 110.” (*Id.*, 9:36-49; *see also* 7:46-62.) In raising the voltage and frequency of common bridge logic 110 in response to detecting that logic cores 120A and 120B are consuming significantly more power than common bridge logic 110, power management logic 200 determines common bridge logic 110 (i.e., “the communication bus”) is not able to accommodate its workload.

109. White also discloses that when power management logic 200 changes the operating frequency of a logic core, it is configured to adjust the operating frequency of common bridge logic 110 for the new operating frequency of the logic core so as to “maintain the communication between the logic core and the common bridge logic.” (White, 17:29-44; *see also* 20:39-57.) Also, “[w]hen

adjusting common bridge logic 110 in response to a new operating frequency for a logic core, power management logic 200 may also take into account the operating frequencies of other logic cores with which common bridge logic 110 may be communicating.” (*Id.*, 17:66-18:3.) Adjusting the operating frequency and/or voltage of common bridge logic 110 in response to new operating voltage and/or frequency of logic cores requires determining if common bridge logic 110 (i.e., “communication bus”) is able to accommodate its workload and make forward progress.

4[d]: capping a bus frequency of at least the communication bus and a frequency of a second processing element in response to determining the first processing element workload indicates the first processing element is the performance bottleneck and determining the communication bus is able to accommodate the bus workload; and

110. White renders Claim 4[d] obvious. As discussed above, White discloses that when it is detected that “logic cores 120A and 120B [are consuming] 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,” power management logic 200 supplies 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.*, 9:36-49; *see*

also 7:46-62.) Thus, White discloses reducing (i.e., “capping”) the operating frequencies of two logic cores 120A (“first processing element”) and 120B (“second processing element”) in response to determining the power consumption (i.e., workload) as indicated by the operating frequency/voltage of common bridge logic 110 indicates that common bridge logic 110 is the performance bottleneck.

111. A POSITA would have found it obvious to reduce or cap the operating frequency of common bridge logic 110 and the operating frequency of logic core 120B when the power consumption of logic core 120A (i.e., “the first processing element workload”) indicates that logic core 120A is the performance bottleneck. White discloses with respect to the common bridge logic 110 that when a component is determined to be a performance bottleneck to the rest of the system, a change in voltage and/or frequency needs to be applied to that component in order to mitigate the performance bottleneck. Thus, a POSITA would have understood that in response to determining the logic core 120A is the performance bottleneck, instead of the common bridge logic, White’s power management logic would apply a change in voltage and/or frequency to the logic core 120A in order to optimize overall performance of the system. At the same time, because the integrated circuit is operating within a power limit (whether a thermal-power limit and/or power-supply limit), when more resources are allocated to the logic core 120A, the operating frequency and/or voltage of the logic core

120B and the common bridge logic 110 may need to be reduced so as to reduce overall power consumption. White discloses that power management logic 200 is configured to “adjust the operating characteristics for one logic core to a lower level than the operating characteristics of another logic core depending upon the overall power consumption ... of the system.” (White, 9:25-30.) A POSITA would also have understood that when common bridge logic 110 is sufficiently able to accommodate its workload, it is not the performance bottleneck, and there may be room to further reduce its operating voltage and/or frequency.

4[e]: *increasing a previous cap on the bus frequency of at least the communication bus in response to determining the first processing element workload indicates the first processing element is the performance bottleneck and determining the communication bus is not able to accommodate the bus workload.*

112. White renders Claim 4[e] obvious. As discussed above, White renders obvious capping the operating frequency of the common bridge logic 110 in response to determining the workload of logic core 120A indicates logic core 120A is the performance bottleneck. *See supra* ¶¶110-111 (Claim 4[d]). Also, as discussed above, White discloses ***raising*** the operating voltage and frequency of common bridge logic 110 in determining that common bridge logic 110 is not able to accommodate its workload. *See supra* ¶¶108-109 (Claim 4[c]). White also discloses that when there is more power budget available, previously reduced operating characteristic, such as operating frequency may optimally be raised

again. (White, 23:12-22, 24:10-21.) A POSITA would have understood that when the common bridge logic is determined not able to accommodate its workload after the operating frequency of common bridge logic 110 is reduced, it means more power resources need to be allocate to common bridge logic 110, which can be done by raising its previously reduced operating frequency. Thus, White renders obvious increasing a previous cap of the operating frequency of the common bridge logic 110 (i.e., “communication bus”) in response to determining the workload or power-consumption of a logic core 120A (i.e., “first processing element”) indicates logic core 120A is the performance bottleneck and determining the common bridge logic 110 is not able to accommodate its workload, and responding appropriately to effect said optimizations.

b. Dependent Claim 7

7: The non-transitory storage medium of claim 4, wherein the limit includes a thermal design power limit.

113. White discloses Claim 7. *See supra* ¶¶100-105 (Claim 4[a]).

2. Ground 6 – White in View of Anderson Renders Obvious Claims 5-6

a. Dependent Claim 5

5: The non-transitory storage medium of claim 4, [2] wherein the first processing element includes one or more graphics cores, the second processing element includes one or more processor cores, the communication bus includes a ring

interconnect, and [1] wherein the code includes privileged level power service code.

114. As viewed by a POSITA, White in view of Anderson renders Claim 5 obvious.

115. First, White discloses the code includes privileged level power service code. Privileged level code typically refers to code that has higher level access to resources in a computer operating system. Code with higher privilege level interacts more directly with the physical hardware such as the control registers and I/O controllers of a computer. White discloses the power management logic 200 may be implemented in “hardware, firmware, and software or a combination thereof.” (White, 16:58-60.) A POSITA would have understood that hardware code usually has higher privilege level, and similarly firmware code is executed above the privilege level of ordinary software. Thus, White discloses that the power service code can have privileged level.

116. A POSITA would have recognized power management in a computer system requires robust and higher-privilege operation that is independent of the operating system and software. Power management must necessarily be implemented in the highest privileged level to ensure immediate robust control, especially where it is clearly anticipated that bottlenecking situations can occur such as that in White. Additionally, a POSITA would have recognized from White’s description of bottlenecks that any application of thermal power control

within non-privileged levels of software running within such bottlenecked core or busses, would itself have been stalled and thus unable to resolve the condition that was causing the bottleneck it was intended to mitigate.

117. Second, while White does not specifically disclose that the first processing element includes one or more graphics cores, the second processing element includes one or more processor cores, or the communication bus includes a ring interconnect, such would have been obvious over White in view of Anderson. As discussed above, Anderson, like White, is directed to power management of an integrated circuit with a multi-core processor. Anderson discloses the first processing element includes one or more graphics cores (i.e., GPU cores), and the second processing element includes one or more processing cores (i.e., CPU cores). *See supra* ¶¶84-87 (Claim 2) (addressing how Anderson discloses the processor including both CPU cores and GPU cores as well as ring interconnect). Anderson also discloses that the communication bus includes a ring interconnect. (*Id.*) It would have been obvious to a POSITA to apply these teachings to White to include both CPU cores and GPU cores. White and Anderson are both directed to thermal power management in multi-core processor. While White does not explicitly disclose a system with different types of processing cores, such as including both GPU cores and CPU cores, White's power management logic can easily accommodate such a system because it independently

adjust the operating characteristics of each component in adjusting the overall power consumption for an integrated circuit. The GPU core can be treated as other core elements, like CPU cores and the common bridge logic. For example, the power for the GPU cores may be supplied through circuitry that is isolated from which supplies power to the other logic cores within an integrated circuit. Indeed, even though White does not explicitly disclose a GPU core, a POSITA would have understood the general reference to “logic core” to encompass GPU core. Thus, a POSITA would have had a reasonable expectation of success in applying the White power management mechanism to an integrated circuit with both GPU core and CPU core as no substantial change is required.

118. A POSITA would have also been motivated to apply White’s power management technique in a processor with heterogeneous cores as disclosed in Anderson. As disclosed in Anderson, efficiently balancing thermal energy generation and system performance could be tricky with a system that has heterogeneous processing components “because the various processing components ... are not created equal.” (Anderson, 1:64-67.) Because in White’s performance management system, the operating characteristics of each component is independently adjustable based on its power consumption and overall optimal performance, the individual features and capability of each component would be

taken into consideration and accommodated, and including GPU cores would add known graphics functionality to White's system.

119. A POSITA would have also been motivated to implement a ring interconnect as disclosed in Anderson in White's system. As discussed above, the common bridge logic 108 in White's system interconnects the key elements in the integrated circuit, which would typically be located at separate locations inside the chip. As discussed above, the progression of more highly-integrated multi-core semiconductor devices naturally tends to favor the use of a ring bus topology for interconnecting, apart from being one of the simplest bus topologies to implement; to do otherwise would likely be wasting useful semiconductor die area. *See supra* ¶¶90 (Claim 2) (addressing why a POSITA would have been motivated to implement ring interconnect). A POSITA would also have had a reasonable expectation of success in implementing a ring interconnect in the White system as Anderson discloses such implementation in the same multi-core processor system as that in White.

b. Dependent Claim 6

120. As viewed by a POSITA, White in view of Anderson renders Claim 6 obvious.

6[a]: The non-transitory storage medium of claim 5, wherein the code, when executed, further causes the machine to perform the operations of: increasing a

previous cap on a frequency for the one or more processor cores in response to determining the one or more graphics cores is not a performance bottleneck;

121. White (as modified by Anderson) renders Claim 6[a] obvious. For example, White discloses that when “logic cores 120A and 120B may consume significantly more power than common bridge logic 110,” and “common bridge logic 110 may be more of a ***bottleneck to overall performance*** than any individual logic core,” the power management logic 200 supplies a higher voltage and frequency to common bridge logic 110 “to reduce any potential bottlenecking of common bridge logic 110.” (White, 9:36-53.) At the same time, the power management logic 200 supplies lower voltages and frequencies to the individual logic cores 120 to reduce overall power consumption. (*Id.*) Doing so, the power management logic 200 helps to “maintain a high overall performance.” (*Id.*) Thus, White teaches when the system is operating under a power limit and a component is determined to be a performance bottleneck, higher voltage and/or frequency is supplied to the bottleneck component, in order to mitigate potential bottlenecking, and the voltages and/or frequencies of other components need to be reduced so as to counterbalance the increase. *See also supra* ¶¶110-112 (Claim 4[d], 4[e].) A POSITA would have understood that this bottleneck component could be either logic core or common bridge logic (“communication bus”).

122. White also discloses operating frequencies of a logic core can be “adjusted upward” (*i.e.*, increasing a “cap”) “in response to greater power

availability.” (*Id.*, 23:12-22; *see also* 24:15-29.) Thus, Whites discloses increasing the operating frequency of a first core in response to determining a second logic core is not a performance bottleneck, because a determination that a second logic core is not a performance bottleneck indicates greater power budget availability, and more power budget can be allocated to other components in the system. As discussed above, White’s power management system, in view of the teachings of Anderson, is applicable to both processor cores and graphics cores. *See also supra* ¶¶114-119 (Claim 5). Thus, a POSITA would have understood it could be graphics core that is determined not a performance bottleneck, with the previous cap on the frequency of the processor cores being increased.

6[b]: capping the frequency for the one or more processor cores in response to determining the first processing element workload indicates the first processing element is the performance bottleneck and determining the communication bus is able to accommodate the bus workload.

123. White (as modified by Anderson) renders Claim 6[b] obvious. As discussed above, White teaches when the system is operating under a power limit and a component, such as a graphic core is determined to be a performance bottleneck, higher frequency is supplied to the bottleneck component, in order to mitigate potential bottlenecking, while the frequency of other components, such as a processor core is reduced (i.e., capped) so as to offset the increased power. *See supra* ¶¶110-112 (Claim 4[d], 4[e]), ¶¶121-122 (Claim 6[a]). Thus, a POSITA would have understood that when the GPU (i.e., “first processing element”) is a

performance bottleneck, frequency of other component such as the processor cores are decreased (i.e., capped) so that the excess power can be allocated to the bottleneck component to improve the overall system performance. Also, when the communication bus is able to accommodate its workload, a POSITA would have understood that it means the communication bus is not a performance bottleneck, and does not need additional power resources and will not be affected by capping the frequency of the processor core.

3. Ground 7 – White (as modified by Anderson) in view of Naffziger Renders Obvious Claims 17-18

a. Dependent Claim 17

17: The non-transitory storage medium of claim 5, wherein the code, when executed, further causes the machine to perform the operations of increasing the frequency of the second processing element and reducing the frequency of the first processing element when the second processing element is to execute a producer thread and the first processing element is to execute a consumer thread to consume information produced by the producer thread.

124. As viewed by a POSITA, White (as modified by Anderson) in view of Naffziger renders Claim 17 obvious.

125. As discussed above, Naffziger discloses a system with consumer-producer relationships between elements of a processor *See supra* ¶¶92-95 (Claim 13) (discussing how Naffziger discloses a system with consumer-producer relationships between elements in the processor). It would have been obvious for a POSITA to combine the teachings of Naffziger with White-Anderson. Like White,

Naffziger is also directed to power management of an integrated circuit with multiple computation/processing units. Particularly, in the Naffziger system, “the CUs measure a corresponding activity level and convey ... data indicative of the measured levels to a power manager,” and the power manager allocates power credits among the CUs based on the corresponding activity levels. (Naffziger, 2:5-16.) When running an application or applications, the power manager will select one discrete power state for each CU and adjust the power state of each CU based on its real-time power consumption, where each power state has different level of performance requirement and different operational frequency. (*Id.*, 6:33-46.) White also discloses that its power management logic 210 adjusts operating voltage/frequency for each core element based on real-time statistics, such as the monitored power consumption of each component. Thus, both Naffziger and White disclose a similar power management method based on real time statistics related to power consumption of each component. Furthermore, White discloses “power management logic 200 may be configured to supply higher voltages and/or frequencies to common bridge logic 110 than to the logic cores” when the common bridge logic is determined to be a performance bottleneck. (White, 9:36-53.) A POSITA would have understood that in this scenario, there is also a consumer/producer relationship between the common bridge logic and the logic cores with the common bridge logic being the producer. As a “producer” in the

system, the common bridge logic manages the source of data and instructions to be consumed by the cores. In view of these similarities, a POSITA would also have reasonable expectation of success in implementing the teachings of Naffziger to the White system as modified by Anderson to include both GPU cores and CPU cores.

126. A POSITA would have been motivated to combine the teaching of Naffziger to the White-Anderson system because increasing the frequency of the producer core and reducing the frequency of the consumer core was a necessary part of creating a reasonably interactive and efficient system with both CPU core and GPU core in a highly-integrated arrangement. A POSITA would have recognized that even small delays between input and output are extremely important in a system employing CPU and GPU for visually interactive environments (e.g. a Windows or MacOS Operating System). For example, between the movement of a mouse (as read by the CPU) in a GUI or visualization application and the production of the changed scene consumed and rendered by the GPU, the GPU must produce a visual result within milliseconds for the satisfaction of even an average user. Thus, strategically adjusting the operating voltage and/or frequency of GPU cores and CPU cores as taught in Naffziger is important in increasing the overall efficiency of the integrated circuit device. *See also supra*

¶¶96-97 Claim 13 (discussing motivation to combine the teachings of Naffziger to similar system).

b. Dependent Claim 18

18: The non-transitory storage medium of claim 17, wherein the code, when executed, further causes the machine to perform the operations of increasing the communication bus frequency by a different amount than an amount by which the frequency of the first processing element is decreased.

127. As viewed by a POSITA, White (as modified by Anderson) in view of Naffziger renders Claim 18 obvious. *See supra* ¶¶74-75 (Claim 14) (addressing how White discloses increasing the communication bus frequency by a different amount than an amount by which the frequency of the first processing element is decreased)

4. Ground 8 – White in view of Simeral Renders Obvious Claims 1, 3, 14 and 15

a. Independent Claim 1

128. As viewed by a POSITA, White in view of Simeral renders obvious Claim 1.

1[pre]: A processor comprising:

129. White discloses Claim 1[pre] to the extent the preamble is limiting. White discloses a computer including “a microprocessor1100” which is a processor and includes multiple logic cores. (White, 2:15-18; *see also id.*, Abstract, 6:23-25, 6:56-63, 24:30-41, Fig. 1., Fig. 8.)




FIG. 8

1[a]: an integrated circuit including:

130. White discloses Claim 1[a]. For example, White discloses the microprocessor 1100 includes an integrated circuit 100 as shown in Fig. 1. (White, 24:37-41, Fig. 1) White further discloses “integrated circuit 100 may represent a multi-core processor and logic cores may represent processor cores that execute instructions supplied to integrated circuit 100.” (*Id.*, 6:53-56.)




FIG. 1

(*Id.*, Fig. 1 (annotated)).

1[b]: a first core;

131. White discloses Claim 1[b]. For example, White discloses the integrated circuit 100 (i.e., multi-core processor) includes multiple logic cores, such as logic core 120A (“a first core”) and logic core 120B, as illustrated in Figure 1. (White, 6:41-56, Fig. 1.)





FIG. 1


(*Id.*, Fig. 1 (annotated).)

1[c]: a cache memory;

132. White discloses Claim 1[c]. For example, White discloses an L2 cache 1128. (White, 25:10-20, 41-49, Fig. 8) Further, White discloses that “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.*”



(*Id.*, 25:45-49, Fig. 8 (annotated).) White also discloses that the integrated circuit 100 includes an internal memory subsystem 170.



(*Id.*, 7:1-7, Fig. 1.) A POSITA would have understood when the L2 cache is integrated into processor 1100, it will be part of the memory subsystem internal to the integrated circuit 100.

1[d]: a communication bus coupled to the first core, the communication bus to connect the first core and the cache memory;

133. White discloses Claim 1[d]. For example, White discloses the multi-core processor 100 includes “a common bridge or interface unit, such as common bridge logic 110.” (White, 7:1-7, Fig. 1.) A POSITA would have recognized that a shared communication bus is going to include the logic necessary to effectively bridge and coordinate multiple access from multiple cores. Particularly, White discloses that the common bridge logic 110 is “configured to *communicate* with the logic cores as well as with other elements or peripherals either internal or external to integrated circuit 100, such as bus 150, peripheral 160 and/or memory subsystem 170.” (*Id.*, 7:3-7.) The ’494 Patent discloses that “[c]ommunication bus 950 includes any known bus or interconnect.” (’494, 15:38-39.) Because White discloses that common bridge logic 100 connects the logic cores 120A and 120B with other elements, it is a communication bus. (White, 7:1-7, Fig. 1) As discussed above, White discloses that the memory subsystem 170 can include the L2 cache 1128. *See supra* ¶132 (Claim 1[c]). White also specifically discloses that common bridge logic 110 connects logic core 120A and logic core 120B with memory subsystem 170 internal to the multi-core processor. Thus, White discloses common

bridge logic 110 (“communication bus”) connects logic core 120A (“first core”) and the L2 cache 1128 (“cache memory”). (White, 7:1-7, Fig. 1; *see also id.*, 2:45-67.)




FIG. 1

(*Id.*, Fig. 1 (annotated).)

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

134. White discloses Claim 1[e]. As discussed above, White discloses power management logic 200 monitors and determines the operating voltages and frequencies of components such as the logic cores and common bridge logic 110. *See supra* ¶¶100-105 (Claim 4[a]). A POSITA would have understood White monitoring the operating voltages and frequencies of the logic core 120A and

common logic bridge 110 is monitoring the workload of those components. *Id.*

Thus, the power management logic 200 is a core workload monitor that determines a core workload for the first core.




FIG. 2A

(*Id.*, Fig. 2A (annotated))

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

135. White discloses Claim 1[f]. As discussed above, the power management logic 200 is also a bus workload monitor that monitors the operating voltage and operating frequency (“bus workload”) of common bridge logic 110 (“communication bus”). *See supra* ¶134 (Claim 1[e]), ¶¶100-105 (Claim 4[a]). Specifically, White discloses that when determining 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,” power management logic 200 is configured to “supply a higher voltage and frequency to common bridge logic 110 to reduce any potential bottlenecking of common bridge logic 110, ... while supplying lower voltages and frequencies to the individual logic cores 120.” (*Id.*, 9:36-53.) In deciding the specific course of action to resolve bus bottlenecking, as described in White, it is clearly necessary to separately evaluate and compare the power consumption of the logic cores and common bridge logic 110, which includes power management logic 200 separately determining the operating voltage/frequency of the common bridge logic 110.

136. Alternatively, Simeral discloses a bus workload monitor. Like White, Simeral also discloses adjusting operational characteristics of components in a computing device so as to modify their power consumption to comply with a performance goal. Simeral discloses “a data path controller with embedded power management for computing devices” in a computing device that can provide power management functions. (*Id.*, [0007], [0019], Fig. 1.) Simeral’s data path controller includes a general-purpose interface signal (“GPIS”) port (“k”) 191 which “can monitor activity associated with bus 130.” (*Id.*, [0019].) Simeral also explains that the term “activity level” means “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.*, [0020].) Thus, the GPIS port k is a bus workload monitor.




FIG. 1A

(*Id.*, Fig. 1A (annotated).)

137. It would have been obvious to a POSITA to apply these teachings to White to include the GPIS port in White’s system to monitor the workload of the common bridge logic 110. Both White and Simeral are directed to power management for a processing system that includes similar components, both can dynamically control supply of power to the similar components of the processor

based on workload data of each component in the system. For example, the data path controller 100 of Simeral includes “a power management controller 106” and “a power-monitoring interface 101” which includes GPIS ports that “monitor and quantify activity levels of a processor or a peripheral device for power management purposes.” (Simeral, [0019].) The power management controller 106 “determines the activity level at which a component ... is operating” and “compares that monitored activity level to a threshold activity level set forth in the performance profile.” (*Id.*, [0024].) Simeral further discloses that if the monitored activity level “falls out of compliance,” then actions will be taken to ensure that “an overall power consumption of the computing device compliant with a global activity threshold or watermark.” (*Id.*) Such actions include increasing or reducing operating voltages or frequencies of a component. (*Id.*, [0026].) This is similar to what is disclosed in White where the power management logic 200 monitors the power consumption or workload of each component and adjust the operating voltages or frequencies accordingly to adjust the overall power consumption.

138. A POSITA would have been motivated to combine the teaching of Simeral and White. For example, both White and Simeral teaches adjusting the operational characteristics (such as voltage, frequency, bandwidth and data throughput) based on the bus workload to adjust its power consumption. Simeral teaches comparing the monitored workload with a threshold in order to make the

adjustment, which is not explicitly disclosed in White but a POSITA would have understood that it is necessary for the power management logic 200 in White to make certain comparison so as to adjust the operating characteristics of the bus for purpose of meeting a performance goal while conserving power consumption.

139. A POSITA would also have reasonable expectations of success in combining the teachings of White and Simeral. White already discloses adjusting the operating characteristics of the common bridge logic 110 based on its power consumption so as to maintain a high overall performance within the power limit. (See e.g., White, 9:36-53.) Simeral discloses providing a dedicated workload monitor for the communication bus so as to more accurately monitor its workload and make real-time decisions for power management. A POSITA would have thus understood the systems to be complementary. Moreover, White discloses that common bridge logic 110 can be a Northbridge, which “is a common component of many processors.” (White, 2:48-49, 7:46-49.) Simeral discloses that the data path controller that includes the bus workload monitor can provide “Northbridge functionalities.” (Simeral, [0042].) Thus, a POSITA would have understood how to implement Simeral’s bus workload monitor, to White’s common bridge logic, and have a reasonable expectation of success.

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,

140. White in view of Simeral renders Claim 1[g] obvious.

141. First, White discloses a power management logic 200 (“balancing control”) that ***monitors*** the temperatures and power consumption of each component (e.g. logic cores and common bridge logic 110) and adjusts the operating voltages/frequencies of each component in real time. (White, 9:39-53, 9:56-62, 10:43-46, Fig. 2A.) Specifically, White discloses that “power management logic 200 may detect a low power state and may attempt to reduce the operating frequency for one or more [components] of a multi-core processor to ***reduce the overall power consumption of the system.***” (*Id.*, 15:49-54, 9:25-53.) White also discloses that “power management logic 200 may be configured to detect when the system transitions from battery power to wall power and adjust the operating frequency of one or more [components] to a higher frequency to take advantage of the fact that the system is no longer on battery power.” (*Id.*, 15:59-64, 9:25-53.) Thus, White discloses the power management logic 200 is configured to “dynamically tune power allocation between the first core and the communication bus.” As discussed above, White also discloses dynamically tuning power allocation based on a power limit for the integrated circuit, and the power limit can

correspond to the maximum thermal dissipation capacity for the integrated circuit.

See supra ¶¶100-105 (Claim 4[a]).

142. Second, Simeral discloses data path controller 100 receives the bus workload from the bus workload monitor (i.e., GPIS port k), and dynamically tunes power allocation. Simeral discloses data path controller 100 operates to “manage power consumption for any component,” (Simeral, [0020]) and “monitor and react to power-related events in real-time or (nearly so).” (*Id.*, [0023].) Specifically, the power management controller 106 within data path controller 100 “selectively monitor[s] and adjust[s] operational characteristics of various components ... during management of power.” (*Id.*, [0030]; see also *id.*, [0020].) Simeral explains that “operational characteristics” can be “voltage, frequency, temperature, bandwidth or any other characteristic that is measurable and can be altered or adjusted to influence power consumption.” (*Id.*, [0020], [0026], [0031].) Such real-time adjustment of each component’s power consumption counts as dynamically tuning power allocation.

143. Simeral also discloses that data path controller 100 dynamically tunes power allocation based on a comparison between the bus workload and a bus workload threshold. For example, Simeral discloses its power management controller 106 “determines the activity level at which a component ... is operating” and “compares that monitored activity level to a threshold activity level set forth in

the performance profile.” (*Id.*, [0024].) Simeral further discloses that if the monitored activity level “falls out of compliance,” then actions will be taken to ensure that “an overall power consumption of the computing device compliant with a global activity threshold or watermark.” (*Id.*) Such actions include increasing or reducing operating voltages or frequencies of a component. (*Id.*, [0026].)

144. It would have been obvious for a POSITA to apply Simeral’s dedicated bus workload monitor and comparison of bus workload with a threshold to White’s system, and a POSITA would have a reasonable expectation to succeed. *See supra* ¶¶135-139 (Claim 1[f]). A POSITA would also have been motivated to do so because White already contemplates adjusting operating characteristics of common bridge logic based on real-time statistics, and comparing the real-time data with a threshold is the most common and effective way to facilitate decision making for purpose of maximize overall system performance (such as alleviating bottlenecking situation) while maintain overall power consumption. *See also supra* ¶¶135-139 (Claim 1[f]).

1[h]: 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.

145. White discloses Claim 1[g]. For example, White discloses “rather than simply shutting down a processor that is operating too warmly (as is generally done in the prior art), one or more of the logic cores in the processor may be

adjusted to operate at a lower frequency and/or voltage so that the processor may generate less heat, and therefore may cool down to an acceptable operating temperature.” (White, 2:29-34; *see also* 22:15-22 (discussing reducing the operating frequency and/or voltage of logic cores “until reaching some *minimal level*”).) Thus, White discloses reducing the power consumption of either a logic core or the common bridge logic, to remain above a low limit, so as to maintain its operation.

b. Dependent Claim 3

3[a]: The processor of claim 1, wherein the balancing control, in response to the integrated circuit operating at the power limit, being adapted to:

3[b]: 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

3[c]: 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.

146. As viewed by a POSITA, White in view of Simeral renders obvious Claim 3. White in view of Simeral renders obvious Claim 3[a]. *See supra* ¶¶140-144 (Claim 1[g]). Further, as discussed above, White in view of Simeral discloses Claim 3[b] and Claim 3[c]. *See supra* ¶¶62-73 (Claim 3[b], 3[c]) (discussing how White and Simeral teach increasing and decreasing communication bus and first core frequency in response to bus workload).

c. Dependent Claim 14

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.

147. As viewed by a POSITA, White in view of Simeral renders obvious Claim 14. *See supra* ¶¶74-75 (Claim 14) (addressing how White discloses the balancing control is to increase the communication bus frequency by a different amount than the first core frequency decrease).

d. Dependent Claim 15

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.

148. As viewed by a POSITA, As discussed above, White discloses a second logic core 120B in addition to the first logic core 120A, and the power management logic 200 (“balancing control”) adjusts power allocation between logic cores 120A, 120B and common bridge logic 110 based on the power limit of the system and the workload of each component. *See infra* ¶149 (Claim 2). White also discloses power management logic 200 increases the operating frequency of “one or more of the logic cores” based on the power availability. (White, 9:1-12; see also *id.*, 11:14, 15:59-64, 23:12-22, 24:24-29.) Thus, White discloses the power management logic (“balancing control”) further increases the frequency for a second core 120B. White also discloses the second logic core 120B and common

bridge logic 110 (“communication bus”) have power planes tied together. (*Id.*) *See supra* ¶¶77-79 (Claim 15) (addressing how White discloses the second core and the communication bus having power planes tied together.).

5. Ground 9 – White (as Modified by Simeral) in view of Anderson Renders Obvious Claim 2

2: The processor of claim 1, further comprising [1] a second core and a second core workload monitor configured to determine a second core workload for the second core, [3] wherein the first core includes a host processing core, the second core includes a graphics processing core, and the communication bus includes a ring interconnect, and [2] wherein the balancing control is further adapted to dynamically tune power allocation between the first core, the second core, and the communication bus based on the power limit and at least the bus workload, the core workload, and the second core workload.

149. As viewed by a POSITA, White (as modified by Simeral) in view of Anderson renders obvious Claim 2. First, White discloses a second logic core 120B. (White, 6:41-49.) White further discloses power management logic 200 monitors and determines the operating voltages and frequencies of components such as both logic cores 120A, 120B, and thus is a core workload configured to determine a first core workload for the first core. *See supra* ¶134 (Claim 1[e]). A POSITA would have understood power management logic 200 equally monitors and determines the operating voltages and frequencies of the second core 120B, it is also a second core workload monitor that determines a second core workload for the second core. Second, White discloses power management logic 200 (“balancing control”) dynamically adjusts power allocation between first core

120A, second core 120B, and common bridge logic 110 (“communication bus”) based on their workload and the power limit. *See supra* ¶¶140-144 (Claim 1[g]), ¶¶98-112 (Claim 4). Third, while White (as modified by Simeral) does not specifically disclose that the first core includes a host processing core, the second core includes a graphics processing core, and the communication bus includes a ring interconnect, this would have been obvious over White (as modified by Simeral) in view of Anderson. *See supra* ¶¶114-119 (Claim 5).


6. Ground 10 – White (as Modified by Simeral) in view of Anderson and Naffziger Renders Obvious Claim 13

13: The processor of claim 2, wherein the balancing control is to increase a frequency of the second core and to reduce a frequency of the first core when the second core is to execute a producer thread and the first core is to execute a consumer thread to consume information produced by the producer thread.

150. As viewed by a POSITA, White (as modified by Simeral) in view of Anderson and Naffziger renders obvious Claim 13. *See supra* ¶¶124-126 (Claim 17).

VI. CERTIFICATION

I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true and, further, that these statements were made with the knowledge that willful false statements and the like are punishable by fine, imprisonment, or both, under Section 1001 of Title 18 of the U.S. Code.

A handwritten signature in dark ink, appearing to read "D. Wyatt", is written over a horizontal line.

David Wyatt

Dated: August 23rd, 2023

APPENDIX A

Materials Considered

| Exhibit No. | Description |
|-------------|---|
| 1001 | U.S. Patent No. 8,898,494 (“ ’494 Patent ”) |
| 1002 | File History for the ’494 Patent |
| 1003 | |
| 1004 | Curriculum Vitae of David Wyatt |
| 1005 | U.S. Patent No. 7,421,601 to Bose (“ Bose ”) |
| 1006 | U.S. Patent No. 7,263,457 to White (“ White ”) |
| 1007 | U.S. Patent No. 8,601,300 to Anderson (“ Anderson ”) |
| 1008 | U.S. Patent Application No. 2009/0150689 (“ Simeral ”) |
| 1009 | U.S. Patent No. 8,510,582 to Naffziger (“ Naffziger ”) |