

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 IPR2023-01343  
Patent 8,898,494

---

**PETITION FOR *INTER PARTES* REVIEW OF  
CLAIMS 8, 10-18, AND 24-47 OF U.S. PATENT NO. 8,898,494  
UNDER 35 U.S.C. §§ 311-319 AND 37 C.F.R. §42.100 *et seq.***

---

Filed on behalf of Petitioner:

Helen Trac (Reg. No. 62,250)  
HOGAN LOVELLS US LLP  
Four Embarcadero, #3500  
San Francisco, CA 94111  
Telephone: 415-374-2300  
Facsimile: 415-374-2399

Celine Jimenez Crowson (Reg. No. 40,357)  
Joseph Raffetto (Reg. No. 66,218)  
Scott Hughes (Reg. No. 68,385)  
Ryan Stephenson (Reg. No. 76,608)  
Nicholas Rotz (Reg. No. 75,959)  
HOGAN LOVELLS US LLP  
555 13th Street N.W.  
Washington, D.C. 20004  
Telephone: 202-637-5600  
Facsimile: 202-637-5710

## TABLE OF CONTENTS

|              |                                                                                                                   |           |
|--------------|-------------------------------------------------------------------------------------------------------------------|-----------|
| <b>I.</b>    | <b>INTRODUCTION .....</b>                                                                                         | <b>1</b>  |
| <b>II.</b>   | <b>MANDATORY NOTICES .....</b>                                                                                    | <b>1</b>  |
|              | A. Real Parties-in-Interest.....                                                                                  | 1         |
|              | B. Related Matters.....                                                                                           | 1         |
|              | C. Counsel and Service Information.....                                                                           | 5         |
| <b>III.</b>  | <b>NOTICE OF FEES PAID .....</b>                                                                                  | <b>6</b>  |
| <b>IV.</b>   | <b>CERTIFICATION OF GROUNDS FOR STANDING .....</b>                                                                | <b>6</b>  |
| <b>V.</b>    | <b>PRECISE RELIEF REQUESTED .....</b>                                                                             | <b>6</b>  |
| <b>VI.</b>   | <b>OVERVIEW OF THE '494 .....</b>                                                                                 | <b>7</b>  |
|              | A. Subject Matter of the '494.....                                                                                | 7         |
|              | B. Prosecution History of the '494 .....                                                                          | 11        |
| <b>VII.</b>  | <b>CLAIM CONSTRUCTION AND POSITA DEFINITION.....</b>                                                              | <b>12</b> |
|              | A. Claim Construction.....                                                                                        | 12        |
|              | B. Definition of a Person of Ordinary Skill in the Art.....                                                       | 12        |
| <b>VIII.</b> | <b>ANALYSIS OF GROUNDS FOR UNPATENTABILITY .....</b>                                                              | <b>13</b> |
|              | A. Ground Based on Bose .....                                                                                     | 13        |
|              | 1. Ground 1 – Bose Discloses or Renders Obvious Claim 1 .....                                                     | 15        |
|              | 2. Ground 2 – Bose in view of White and Simeral Renders<br>Obvious Dependent Claims 3, 14-15 .....                | 27        |
|              | 3. Ground 3 – Bose in view of Anderson Renders Obvious<br>Dependent Claims 2, 13 .....                            | 39        |
|              | 4. Ground 4 – Bose (as Modified by Anderson) in view of<br>Naffziger Renders Obvious Dependent Claim 13 .....     | 45        |
|              | B. Grounds Based on White.....                                                                                    | 50        |
|              | 1. Ground 5 – White Renders Obvious Claims 4, 7 .....                                                             | 50        |
|              | 2. Ground 6 – White in View of Anderson Renders Obvious<br>Claims 5-6 .....                                       | 60        |
|              | 3. Ground 7 – White (as modified by Anderson) in view of<br>Naffziger Renders Obvious Claims 17-18.....           | 64        |
|              | 4. Ground 8 – White in view of Simeral Renders Obvious Claims<br>1, 3, 14 and 15 .....                            | 66        |
|              | 5. Ground 9 – White (as Modified by Simeral) in view of<br>Anderson Renders Obvious Claim 2 .....                 | 82        |
|              | 6. Ground 10 – White (as Modified by Simeral) in view of<br>Anderson and Naffziger Renders Obvious Claim 13 ..... | 84        |

|            |                                        |           |
|------------|----------------------------------------|-----------|
| <b>IX.</b> | <b>TRIAL SHOULD BE INSTITUTED.....</b> | <b>84</b> |
| A.         | <i>Fintiv</i> Factors.....             | 84        |
| B.         | <i>General Plastic</i> .....           | 85        |
| C.         | 35 U.S.C. § 325(d).....                | 87        |
| <b>X.</b>  | <b>CONCLUSION .....</b>                | <b>87</b> |

## **TABLE OF AUTHORITIES**

|                                                                                                               | <b>Page(s)</b> |
|---------------------------------------------------------------------------------------------------------------|----------------|
| <br><b>Cases</b>                                                                                              |                |
| <i>Apple, Inc. v. Fintiv, Inc.</i> ,<br>IPR2020-00019 (PTAB Mar 20, 2020) .....                               | 84             |
| <i>Boss Corp. v. Koss Corp.</i> ,<br>IPR2021-00680 (PTAB Oct. 13, 2021) .....                                 | 84             |
| <i>General Plastic Industrial Co. v. Canon Kabushiki Kaisha</i> ,<br>IPR2016-01357 (PTAB Sept. 6, 2017) ..... | 85             |
| <i>Sotera Wireless, Inc. v. Masimo Corp.</i> ,<br>IPR2020-01019 (PTAB Dec. 1, 2020) .....                     | 84             |
| <i>Valve Corp. v. Electronic Scripting Products, Inc.</i> ,<br>IPR2019-00062 (PTAB Apr. 2, 2019) .....        | 85, 86         |
| <i>Valve Corp. v. Electronic Scripting Products, Inc.</i> ,<br>IPR2019-00064 (PTAB May 1, 2019) .....         | 85             |
| <br><b>Statutes</b>                                                                                           |                |
| 35 U.S.C. 282(b) .....                                                                                        | 12             |
| 35 U.S.C. § 325(d) .....                                                                                      | 87             |

## **I. INTRODUCTION**

Petitioner Mercedes-Benz USA, LLC. (“Mercedes” or “Petitioner”) challenges Claims 1-7, 13-15, 17-18 of US8898494 (“**’494**”) (EX1001). The ’494 claims nothing more than what was well-known in the art of integrated circuitry at the time it was filed. The ’494 purports to invent the concept of efficiently balancing performance and power between processing elements and a communication bus based on measured workload. But this was not new at the time of the ’494. Dynamic power management based on the power limit of an integrated circuit was well known before this. For example, the primary references herein, US7421601 (“**Bose**”) (EX1005) and US7263457 (“**White**”) (EX1006), both disclose a near-identical ways of balancing power among multiple cores and a communication bus. Neither Bose nor White was considered during prosecution of the ’494, and alone and in combination with other prior art, they render Claims 1-7, 13-15, 17-18 of the ’494 invalid.

## **II. MANDATORY NOTICES**

### **A. Real Parties-in-Interest**

The real parties-in-interest are Petitioner Mercedes-Benz USA, LLC; Mercedes-Benz Group AG; Mercedes-Benz AG; and Mercedes-Benz Intellectual Property GmbH & Co. KG.

### **B. Related Matters**

Assignee of record Daedalus Prime LLC (“**Daedalus**”) asserted the ’494 in the following proceedings:

- *Daedalus Prime LLC v. Arrow Electronics, Inc et al.*, Case No. 1-22-cv-01107 (D. Del.), filed August 23, 2022 against Arrow Electronics, Inc. (“Arrow”); Avent, Inc. (“Avnet”); Digi-Key Electronics (“DigiKey”); Future Electronics Inc. (“Future”); Mazda Motor Corporation, Mazda North American Operations, and Mazda Motor of America, Inc. (collectively, “Mazda”); Mercedes-Benz Group AG, Mercedes-Benz AG, and Mercedes-Benz USA, LLC (collectively, “Mercedes”); Mouser Electronics, Inc. (“Mouser”); Newark; NXP Semiconductors N.V. and NXP USA, Inc. (collectively, “NXP”); Rochester Electronics, LLC (“Rochester”); and Visteon Corporation (“Visteon”). **This case was voluntarily dismissed by Daedalus on February 28, 2023.**
- *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, Case No. 1-22-cv-01108 (D. Del.), filed August 23, 2022 against Mazda Motor Corporation, Mazda North American Operations, and Mazda Motor of America, Inc. (collectively, “Mazda”), MediaTek Inc. and MediaTek USA Inc. (collectively, “MediaTek”), Mercedes-Benz Group AG, Mercedes-Benz AG, and Mercedes-Benz USA, LLC (collectively,

“Mercedes”), and Visteon Corporation (“Visteon”). **This case was voluntarily dismissed by Daedalus on November 14, 2023.**

- *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, Case No. 1-22-cv-01109 (D. Del.), filed August 23, 2022 against Mazda Motor Corporation, Mazda North American Operations, and Mazda Motor of America, Inc. (collectively, “Mazda”), Mercedes-Benz Group AG, Mercedes-Benz AG, and Mercedes-Benz USA, LLC (collectively, “Mercedes”), Qualcomm Inc. and Qualcomm Technologies, Inc. (collectively, “Qualcomm”), and Visteon Corporation (“Visteon”).

**Daedalus voluntarily dismissed Defendants Mazda, Mercedes and Visteon from this case on March 20, 2023.**

- *In the Matter of Certain Semiconductors and Devices and Products Containing the Same, Including Printed Circuit Boards, Automotive Parts, and Automobiles*, Inv. No. 337-TA-1332 (U.S. International Trade Commission), filed August 23, 2022 against Arrow Electronics, Inc. (“Arrow”), Avnet, Inc. (“Avnet”), Digi-Key Electronics (“Digi-Key”), Future Electronics Inc. (“Future”), Mazda Motor Corporation, Mazda North American Operations, and Mazda Motor of America, Inc. (collectively, “Mazda”), MediaTek Inc. and MediaTek USA Inc. (collectively, “MediaTek”), Mercedes-Benz Group AG, Mercedes-

Benz AG, and Mercedes-Benz USA, LLC (collectively, “Mercedes”), Mouser Electronics, Inc. (“Mouser”), Newark (an Avnet Company) (“Newark”), NXP Semiconductors N.V. and NXP USA, Inc. (collectively, “NXP”), Rochester Electronics, LLC (“Rochester”), and Visteon Corporation (“Visteon”). **On January 20, 2023, Daedalus moved to withdraw the complaint as to all Respondents and the Administrative Law Judge granted the termination on January 30, 2023.**

- *Daedalus Prime LLC v. Samsung Electronics Co., Ltd. et al*, Case No. 2-22-cv-00352 (E.D. Tex.), filed September 12, 2022 against Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., (collectively “Samsung”), and Taiwan Semiconductor Manufacturing Company Limited (“TSMC”). **Daedalus and Samsung jointly filed a motion to dismiss the claims against Samsung in this lawsuit on June 6, 2023.**
- *In the Matter of Certain Integrated Circuits, Mobile Devices Containing the Same, and Components Thereof*, Inv. No. 337-TA-1335 (U.S. International Trade Commission), filed September 12, 2022 against Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc. (“Samsung”), and Qualcomm Incorporated

**(“Qualcomm”). On April 25, 2023, the Administrative Law Judge granted the joint motion of Daedalus and Samsung to immediately suspend the Procedural Schedule as to Samsung. On June 14, 2023, Daedalus and Samsung jointly moved to terminate investigation as to Samsung based on settlement.**

The ’494 is also the subject of IPR2023-00617 (challenging Claims 1, 3, 12, 14 and 15) filed by Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., (collectively, “Samsung”) and Qualcomm, Inc. Samsung and Daedalus jointly submitted a motion to terminate this IPR as to Samsung on June 2, 2023 based on settlement.

**C. Counsel and Service Information**

Lead counsel is Celine Crowson (Reg. No. 40,357). Backup counsel are Joe Raffetto (Reg. No. 66,218), Scott Hughes (Reg. No. 68,385), Ryan Stephenson (Reg. No. 76,608), Nicholas Rotz (Reg. No. 75,959), and Helen Y. Trac (Reg. No. 62,250). Service information is as follows:

|                               |                                                                                                                                                                            |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Post and Hand Delivery</b> | Hogan Lovells US LLP<br>555 13th Street, N.W.<br>Washington, D.C. 20004                                                                                                    |
| <b>Email</b>                  | celine.crowson@hoganlovells.com<br>joseph.raffetto@hoganlovells.com<br>ryan.stephenson@hoganlovells.com<br>scott.hughes@hoganlovells.com<br>nicholas.rotz@hoganlovells.com |

|                              |                             |
|------------------------------|-----------------------------|
|                              | helen.trac@hoganlovells.com |
| <b>Telephone / Facsimile</b> | 202.637.5600 / 202.637.5910 |

### III. NOTICE OF FEES PAID

Fees are submitted herewith. If additional fees are due during the proceeding, the undersigned authorizes the Office to charge them to Deposit Account No. 50-1349.

### IV. CERTIFICATION OF GROUNDS FOR STANDING

Petitioner certifies that the '494 is available for *inter partes* review, and that it is not barred or estopped from requesting this review.

### V. PRECISE RELIEF REQUESTED

The relief requested is cancellation of the challenged claims, as follows:

| Ground   | References                   | Claims <sup>1</sup> | Basis   |
|----------|------------------------------|---------------------|---------|
| <b>1</b> | Bose                         | <b>1</b>            | 102/103 |
| <b>2</b> | Bose / White / Simeral       | 3, 14-15            | 103     |
| <b>3</b> | Bose / Anderson              | 2                   | 103     |
| <b>4</b> | Bose/ Anderson / Naffziger   | 13                  | 103     |
| <b>5</b> | White                        | <b>4, 7</b>         | 103     |
| <b>6</b> | White / Anderson             | 5-6                 | 103     |
| <b>7</b> | White / Anderson / Naffziger | 17-18               | 103     |

---

<sup>1</sup> Independent claims in bold.

| Ground | References                             | Claims <sup>1</sup> | Basis |
|--------|----------------------------------------|---------------------|-------|
| 8      | White / Simeral                        | 1, 3, 14-15         | 103   |
| 9      | White / Simeral / Anderson             | 2                   | 103   |
| 10     | White / Simeral / Anderson / Naffziger | 13                  | 103   |

## VI. OVERVIEW OF THE '494

### A. Subject Matter of the '494

The '494 purports to provide a method for efficiently balancing performance and power among processing cores and/or a communication bus based on measured workloads, so as to achieve maximum overall performance under a power limit. ('494, Abstract, 7:40-61, 15:4-21.) The '494 provides 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 depicts an integrated circuit including multiple processing cores and a communication bus that implements the claimed method:




FIG. 9

(*Id.*, Fig. 9 (annotated).)

The integrated circuit (“IC”) 900 includes a bus 950 (yellow) which connects GPU cores 910 (purple), CPU cores 905 (dark blue), and a cache memory 960 (grey). (*Id.*, 15:59-63.) As shown in Figure 9, GPU cores 910 are associated with a GPU workload monitor 911 (orange), CPU cores 905 are associated with a CPU workload monitor 935 (light blue), while communication bus 950 is associated with a bus workload monitor 951 (green). (*Id.*, 16:14-19.) The workload

monitors measure workload of the respective components, determining each component's activity or "any other metric for determining use of a processing element." (*Id.*, 16:19-23.) The integrated circuit 900 further includes a power management service ("PMS") 920 to dynamically balance power between GPU cores 910, CPU cores 905 and communication bus 950 based on those measured workloads:




FIG. 10

(*Id.*, 16:57-17:3; Fig. 10.)

For example, the '494 discloses that if the workload of communication bus

950 is too high, such as over the a threshold, then communication bus frequency/power is increased. (*Id.*, 17:4-15.) But if IC 900 has already been operating at its power limit, then “increase in communication bus frequency is counterbalanced in the CPU and/or GPU.” (*Id.*) The ’494 also discloses that 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.)

As another example, the ’494 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 communication bus 950 is able to accommodate its workload, then frequency of both the communication bus 950 and 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 communication bus 950 is not able to accommodate its workload, the cap on communication bus frequency will be increased so that it can accommodate more work. (*Id.*, 19:26-30.) The ’494 explains that 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.)

## **B. Prosecution History of the '494**

The '494 was filed on February 16, 2012 as a continuation-in-part of US9304570, filed on December 15, 2011.<sup>2</sup> The '494 received two office actions. In a Non-Final Rejection, independent claim 1 was rejected over US2008/0104428 (“**Naffziger**”) in view of US5758133 (“**Evoy**”). (EX1002, '494 File History, Non-Final Rejection mailed October 4, 2013.) In response, Patent Owner amended the claim to state that power allocation between the first core and communication bus is based on a power limit for the integrated circuit and “a comparison between the bus workload and a bus workload threshold,” and “the power limit corresponding to a maximum thermal dissipation capacity for the integrated circuit.” (*Id.* Amendments to the Claims mailed January 3, 2014.) In a Final Rejection, amended claim 1 was rejected as anticipated by US2005/0060594 (“**Barr**”), and independent claim 4 was rejected as anticipated by US2011/0310413 (“**Nakao**”). (*Id.*, Final Rejection mailed on April 22, 2014.) In response, Patent Owner further amended independent claim 1 to recite “wherein a power consumption of one of the first

---

<sup>2</sup> Mercedes contends that the '494 is not entitled to the December 2011 priority date. For example, the claims only find support in the material newly added to the '494, including Fig. 9 and Fig. 10 and descriptions thereof. However, all references cited in this Petition are prior art even if the '494 is entitled to this priority date.

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.” (*Id.*, Amendments to the Claims mailed July 18, 2014.) As to independent claim 4, Patent Owner did not make substantive amendments but argued that Nakao is directed to a copy machine that does not disclose elements of claim 4. (*Id.*, Remarks/Arguments mailed July 18, 2014.) Following this, the examiner issued a Notice of Allowance. (*Id.*, Notice of Allowance and Fees Due.)

## **VII. CLAIM CONSTRUCTION AND POSITA DEFINITION**

### **A. Claim Construction**

Claims in an *inter partes* review “shall be construed using the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.” 37 C.F.R. §42.100(b). Unless indicated otherwise, all claim terms herein are given such ordinary and customary meaning.

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

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. (Wyatt Decl., ¶21.)<sup>3</sup> The more education one has, the less experience needed to attain an ordinary level of skill. (*Id.*) Similarly, more experience in the field may serve as a substitute for formal education. (*Id.*)

## **VIII. ANALYSIS OF GROUNDS FOR UNPATENTABILITY**

### **A. Ground Based on Bose**

Bose, titled “Method and System for Controlling Power in a Chip through a Power-Performance Monitor and Control Unit,” is US7421601, which was filed February 17, 2006 and issued September 2, 2008. Bose is a prior art under at least pre-AIA § 102(b).

Bose discloses “a power/performance monitoring and control unit (PMCU)” that can improve 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

---

<sup>3</sup> David Wyatt is an expert in power management for integrated circuits, among other things. His Declaration is included as EX1003 (“Wyatt Decl.”).

and between different workloads.” (*Id.*, 4:30-46.) Fig. 1 depicts a high level overview of PMCU organization in a system 100 with a multi-core chip 104:




FIG. 1

(*Id.*, 5:61-67; Fig. 1.)

System 100 comprises cores 112-1, 112-2 and 112-R, a global PMCU 110, interconnect 108 and L2/L3 cache hierarchy 106 as well as local PMCUs 118-1, 118-2, 118-R, 114, 116. (*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 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.*, 18-31.)

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

Bose discloses or renders obvious Claim 1. (Wyatt Decl., ¶¶38-59.)

#### **1[pre]: A processor comprising:**

Bose discloses Claim 1[pre] to the extent limiting. (*Id.*, ¶39.) 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.) A microprocessor system 100, i.e., a processor that includes a multi-core chip 104, is illustrated in Fig. 1:




A processor

FIG. 1

(*Id.*, Fig. 1 (annotated), 4:21-24, 5:61-67.)

**1[a]: an integrated circuit including:**

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



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

**1[b]: a first core;**

Bose discloses Claim 1[b]. (Wyatt Decl., ¶41.) Bose discloses a multi-core chip 104, which includes core 112-1 (i.e., a first core), 112-2, 112-R. (Bose, 5:61-67, 6:18-19; Fig. 1.)




FIG. 1

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

**1[c]: a cache memory;**

Bose discloses Claim 1[c]. (Wyatt Decl., ¶42.) Bose discloses that multi-core chip 104 includes an “L2/L3 cache hierarchy” 106, which is a multi-level cache system including L2 and L3 cache memories. (Bose, 5:61-67, 6:18-19; Fig. 1.) Further, the system “may include at least one processor coupled directly or indirectly to memory elements,” including “*cache memories*.” (*Id.*, 5:16-24.)




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;

Bose discloses Claim 1[d]. (Wyatt Decl., ¶43.) Bose discloses that multi-core chip 104 includes “interconnect” 108. (Bose, 5:61-67, 6:56-60, 7:4-11, 7:18-23, Fig. 1.) Interconnect 108, first core 112-1 and L2/L3 cache hierarchy 106 are all “core element[s]” within multi-core chip 104 that share an overall power budget and are subject to the same power management and control system. (*Id.*, 5:64-67, 7:4-28.) Further, “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.) Because these components are interconnected in the power domain and are capable of simultaneous multithreading, a POSITA would have understood that the elements shown in Fig. 1 represents a physically interconnected system where interconnect 108 connects at least one core and the cache memories and performs the same interconnecting function as the “communication bus” in the ’494. (’494, 15:38-39.)





FIG. 1

(Bose, Fig. 1 (annotated).)

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

Bose discloses Claim 1[e]. (Wyatt Decl., ¶¶46-47.) Bose discloses that microprocessor system 100 includes a local PMCU 118-1 associated with core 112-1:



((Bose, Fig. 1 (annotated), 5:61-67, 6:18-20, 6:56-60.) Local PMCU 118-1 monitors “performance statistics such as cache miss rates, branch misprediction rates, instructions per cycle (IPCs) etc.” of core 112-1. (*Id.*, 6:56-62.) Thus, local

PMCU 118-1 is a core workload monitor as it monitors performance statistics reflecting workload of core 112-1. (Wyatt Decl., ¶46.) Further, local PMCU 118-1 reports monitored information to global PMCU (*id.*, 6:56-7:3), and thus determines a core workload for core 112-1 to report.

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

Bose discloses Claim 1[f]. (Wyatt Decl., ¶¶48-49.) Bose discloses that microprocessor system 100 includes local PMCU 116 associated with interconnect 108. (Bose, 5:61-67, 6:18-24, 6:56-60, 7:4-7, Fig. 1.)




FIG. 1

(*Id.*, Fig. 1 (annotated).) Local PMCU 116 performs the same functions as local PMCU 118-1 (i.e., the core workload monitor). (Wyatt Decl., ¶¶48-49.) Thus, Bose discloses a bus workload monitor that determines a bus workload for the communication bus. *See supra* Claim 1[e].

**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,***

Bose discloses Claim 1[g]. (Wyatt Decl., ¶¶50-56.) Bose discloses a chip-level global PMCU 110 functioning as “balancing control.”




FIG. 1

(Bose, Fig. 1 (annotated), 4:19-29, 5:61-67, 6:18-37, 6:56-7:31; *see also* 8:11-18, Fig. 2; Wyatt Decl., ¶50.)

First, Bose discloses PMCU 110 (“balancing control”) receives the bus workload from local PMCU 116 (“bus workload monitor”), as PMCU 116 reports “a summary of monitored results,” (i.e., workload of interconnect 118), upwards to global PMCU 110 (“balancing control”). (Bose, 6:18-34.)

Second, Bose discloses dynamic tuning of power allocation between core 112-1 and interconnect 108 based on a power limit, which corresponds to a maximum thermal dissipation capacity for the integrated circuit. Global PMCU

110 “dynamically manage[s] chip power between the components ... to meet the power limit with the expectation of increased system performance,” where the components include “core element 112-1, 112-2, 112-R, L2/L3 cache hierarchy 106 and interconnect 108.” (*Id.*, 4:42-45, 6:56-7:23; *see also id.*, 7:9-11, 7:18-21, 7:25-27, 3:40-44, 7:65-8:10; Wyatt Decl., ¶52.) Further, the power limit is “due to packaging/cooling constraints” (Bose, 4:34-35), which determines “peak temperature limit of the chip.” (*Id.*, 2:41-43.) A POSITA would have recognized this is expressing the same “power limit” in the ’494, i.e., “thermal dissipation power” or “thermal design power.” (Wyatt Decl., ¶52.)

Third, Bose discloses dynamic tuning of power allocation based on a comparison between bus workload and bus workload threshold. (*Id.*, ¶53.) Global PMCU “choose[s] the power modes and assign power [budget] for” core 112-1 and interconnect 108 based on workload information provided by local PMCUs 118-1, 116. (Bose, 7:4-31, *see also id.*, 6:56-60, 7:41-61, 8:11-12, 9:3-24.) Each power mode represents current or expected workload of a component operating in this particular work mode. (*Id.*, 7:4-31.) Bose further discloses a Power/Performance Trade-off Table 206 that “includes the *predicted power and performance* for each power mode,” which is used by global PMCU to “assist in calculating the power modes and power budges for each core.” (*Id.*, 7:41-45, 8:35-43.) The predicted

power and performance represents a workload threshold for a given component working in the associated “power mode.” (Wyatt Decl., ¶53.)

A POSITA would have understood global PMCU compares a workload threshold as exemplified by the predicted power and performance of each power mode and the monitored workload to select and adjust the power mode of a component. (Wyatt Decl., ¶¶54-55.) Bose explains power mode is adjusted based on “monitored power and performance levels” of the components, and adjustment of power mode means decreasing or increasing power level “with corresponding decrease or increase in net performance” of a component. (Bose, 12:30-41, Fig. 7.) Further, “the amount of such change in ... performance” is based on the predicted power and performance for each power mode. (*Id.*, 8:35-47, 12:42-45.) Thus, a POSITA would have understood that in Bose, selecting or adjusting power mode for a component entails comparing a workload threshold, such as predicted power and performance associated with each power mode with real time statistics, such as the monitored workload of a component. (Wyatt Decl., ¶¶55-56.)

***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.***

Bose discloses Claim 1[h]. (*Id.*, ¶¶57-59.)

First, Bose discloses reducing power consumption of a component, such as core 112-1 (“first core”) or interconnect 108 (“communication bus”). (*Id.*, ¶57.)

Bose discloses that global PMCU selects and adjusts power mode for core 112-1 and interconnect 108. (Bose, 7:4-17.) Bose also discloses the local PMCU for each component “will inform the global PMCU to *assign a lower power mode* to this [component]” “if the [assigned] 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.)

Second, Bose discloses operations of components will be maintained over a low limit of being powered off. (Wyatt Decl., ¶¶58-59.) Bose’s system is optimized for power with clock gating, which means alternatively declining to feed a clock signal to portions of the integrated circuit while they are *powered-on*. (*Id.*, ¶58.) Thus, the components in the Bose system will always remain powered, i.e., their operations will be maintained, but they may be intermittently clock gated for purpose of power conservation. (*Id.*) Bose also states its invention is meant to achieve performance goals, which could not be achieved with components being turned off. (*Id.*, ¶59.)

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

### **a. Dependent Claim 3**

Bose in view of White and Simeral renders obvious Claim 3. (Wyatt Decl., ¶¶60-73.)

**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:**

Bose discloses Claim 3[a]. (*See supra* 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**

Bose in view of White and Simeral renders Claim 3[b] obvious. (Wyatt Decl., ¶¶62-71.) 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.<sup>4</sup> (White, 2:15-44.) White discloses a multi-core integrated circuit 100 including a common bridge logic 110, which is a communication bus that communicates with multiple logic cores of the integrated circuit, such as logic cores 120A, 120B. (*Id.*, Abstract, 2:22-44, 6:34-40, Fig. 2A.) In White, operating voltage and/or frequency of each logic core and common bridge logic can be independently adjusted for power management and temperature control. (*Id.*)

---

<sup>4</sup> White issued August 28, 2007 and is a prior art at least under § 102(b).




FIG. 2A

(*Id.*, Fig. 2A (annotated).) White discloses both increasing frequency for common bridge logic (“communication bus”), and capping (i.e., limiting) frequency for one logic core (“first core”) in response to the common bridge logic being a performance bottleneck. (*Id.*, 9:36-53; Wyatt Decl., ¶¶62-64.)

Simeral further teaches adjusting operational characteristics of a bus to modify its power consumption based on a comparison of its activity level with a threshold. (Simeral, Abstract, [0007], [0021], [0024], [0026]; Wyatt Decl., ¶¶65-66.) 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.<sup>5</sup> (Simeral, Abstract.) Simeral discloses “a data path controller” in a computing device that

<sup>5</sup> Simeral published June 11, 2009, and is a prior art at least under pre-AIA §102(b).

provides 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 "activity level" can represent the amount of load 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 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], [0026].)




FIG. 1A

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

Simeral also teaches that, in response to the monitored activity level being above the threshold activity level, frequency of the component can be increased while frequency of other components can be reduced to meet the overall power limit. (Wyatt Decl., ¶¶67-68.) Specifically, if it is determined the monitored activity level “falls out of compliance,” power management controller 106 adjusts its operational characteristics. (Simeral, [0020], [0024], [0026].) Further, “the power management controller can recalibrate activity levels of other less essential

components to compensate for 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 also teaches that “activity level” can be expressed in terms of *frequency*. (*Id.*)

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

Similarly, White also contemplates adjusting operating voltages and/or frequencies of individual components so as to maintain total power consumption of all components under a power limit. (Wyatt Decl., ¶70.) 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. (*Id.*) Thus, A POSITA would have been motivated to include 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 taught by White and Simeral, in the system of Bose. (*Id.*) 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. (*Id.*)

Further, a POSITA would have had a reasonable expectation of success in applying the teachings in White and Simeral to Bose. (Wyatt Decl., ¶71.) 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 like Simeral, and also contemplated comparing real time data, such as monitored workload with a workload threshold in order to adjust power mode. (*Id.*) 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. (*Id.*) 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. (*Id.*)

**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.**

Bose in view of White and Simeral renders Claim 3[c] obvious. (Wyatt Decl., ¶72.) 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* Claim 3[b]. Simeral further discloses that in such system, data path controller 100 monitors data transfer rate of the bus as its activity level, and decreases bus width “if the data rate is below some threshold.” (Simeral, [0021].) Further, when the monitored activity level is not compliant with a threshold, other operating characteristics such as frequencies can be adjusted so the component can comply with its threshold activity level. (*Id.*, [0020], [0024], [0026].) Thus, a POSITA would have understood when the monitored activity level of a bus, like that in Bose, is below a threshold, bus frequency as an alternative of bus width can be decreased. (Wyatt Decl., ¶72.) Also as discussed above, Simeral discloses the activity level of some components could be reduced so as to compensate the increased activity level of other components. *See supra* Claim 3[b]. A POSITA would have understood that when activity level of the bus is reduced, the additional power could be used to compensate the activity level of

another component, such as a processing core, by increasing the frequency of the processing core. (Wyatt Decl., ¶72.)<sup>6</sup>

**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.***

Bose in view of White and Simeral renders obvious Claim 14. (Wyatt Decl., ¶75.) Bose as modified by White (and Simeral) discloses a system with balancing control to increase the communication bus frequency and decrease first core frequency. *See supra* Claim 1, 3[b]. White further discloses that in doing so, the operating 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-

---

<sup>6</sup> Bose also discloses Claim 3[c], in describing that a higher power mode is assigned when there is excess power as compared to the power budget. (Wyatt Decl., ¶73.) A POSITA would have understood bus workload being lower than workload threshold means there is excess power that could be allocated to other components, such as the first core and by increasing its operating frequency, and the power mode can be reduced for the bus to maintain overall power consumption without affecting system performance. (*Id.*)

44, 6:34-40.) White also 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 frequency of the common bridge logic is offset by corresponding decreases to the operating frequencies of *multiple cores*, a POSITA would have understood that frequency of common bridge logic could be increased by a different amount than the decrease of the voltage and frequency of *a single core*. (Wyatt Decl., ¶75.)

**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.***

Bose in view of White and Simeral renders obvious Claim 15. (Wyatt Decl., ¶¶76-80.) As discussed below, 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 first core 112-1, second core 112-2, and interconnect 108 based on system power limit and workload of each component. *See infra* VIII.A.3(Claim 2). Bose also discloses increasing or decreasing the power mode of a component

based on its power usage. (Bose, 9:14-30.) Each power mode specifies operating voltage and frequency of a component. (*Id.*, 11:3-7.) Thus, when the power mode of each component is adjusted, the frequency of that component will be adjusted accordingly. A POSITA would have understood that a higher power mode means more power budget is allocated to the component, and thus its operating frequency is increased. (Wyatt Decl., ¶76.) Bose discloses this power distribution scheme is equally and independently applicable to every core. (Bose, 7:4-31, 9:3:30) Thus, Bose discloses global PMCU (“balancing control”) further increases the frequency for a second core 112-2. (Wyatt Decl., ¶76.)

Bose does not explicitly disclose the second core and communication bus having power planes tied together. White discloses a multicore processor system in which a second logic core and common bridge logic (“communication bus”) have power planes tied together. For example, White discloses operating voltages/frequencies of common bridge logic and logic cores can change together in the same way at the same time. (*Id.*, 9:32-36 (disclosing adjusting the voltages/frequencies “in step”).) This indicates the logic cores and the communication bus are intended to be logically co-controlled in sync, as if the power planes were physically tied together. (Wyatt Decl., ¶78.) Also, White discloses the same voltage regulator 130 supplies voltages to both the logic cores and common bridge logic:




FIG. 1

(White, Fig. 1 (annotated), 7:23-31, 8:32-34.) This indicates these components have power planes tied together. (Wyatt Decl., ¶78.) Additionally, White discloses the same frequency regulator 140 supplies “clock frequency signals” to both logic cores and common bridge logic 100. (White, 7:31-38, Fig. 1; *see also* 8:37-40.) 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. (Wyatt Decl., ¶79.) 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. (*Id.*) 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.


(*Id.*)

**3. Ground 3 – Bose in view of Anderson Renders Obvious Dependent Claims 2, 13**

***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.***

Bose in view of Anderson renders obvious Claim 2. (Wyatt Decl., ¶¶80-90.)

First, Bose discloses a second core 112-2 and a local PMCU 118-2 (“second core workload monitor”) associated with core 112-2. (Bose, 5:61-67, 6:18-20, 6:56-60; Fig. 1.) *See supra* Claim 1[e].



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

Second, Bose discloses that global PMCU (“balancing control”) dynamically adjusts power allocation between first core 112-1, second core 112-2, and interconnect 108 (“communication bus”) based on the power limit, and workload of interconnect 108, core 112-1, and core 112-2. (Wyatt Decl., ¶¶82-83.) *See also supra* Claim 1[g]. Local PMCU 118-2 (“second core workload monitor”) functions the same as local PMCU 118-2 (“core workload monitor”) and local PMCU 116 (“bus workload monitor”), all of which report workload of the respective component upwards to global PMCU. (Bose, 6:18-37, 6:56-62.) As discussed, global PMCU “dynamically manage[s] chip power between the components ... to

meet the *power limit*’ (*id.*, 4:42-45) based on information collected from local PMCUs for each component. (*Id.*, 6:56-7:23; *see also* 7:18-21, 7:65-8:10.)

Third, while Bose does not specifically disclose 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. (Wyatt Decl., ¶¶84-90.) Anderson, titled “System and Method for Managing Thermal Energy Generation in a Heterogeneous Multi-core Processor,” was filed September 29, 2011, and issued December 3, 2013, and is prior art under pre-AIA §102(e). Anderson, like Bose, is directed to power management of an integrated circuit with a multi-core processor. (*Id.*, ¶85.) Specifically, Anderson discloses a thermal power management system for a *heterogeneous multi-core processor*. (Anderson, Abstract.) Anderson first discloses a chip 102 that includes an applications central processing unit (“CPU”) 110 with multiple CPU cores, such as core 0, core 1 and core N, each of which is a host processing core:<sup>7</sup>

---

<sup>7</sup> In the ’494, the terms central processor and host processor are used interchangeably, and refer to a processor other than the graphics processor. (’494, 1:50-52, 3:15-21.)




**FIG. 5A**

(*Id.*, Fig. 5A (annotated), 9:66-10:7.) A POSITA would have understood CPU cores in Fig. 5A are host processing cores as they are central processing components of a typical processor chip for a portable computing device, and run the host operating system 207 collectively. (Wyatt Decl., ¶85.) Anderson also discloses the same chip 102 comprises GPU cores, shown as graphics processors 135A-D. (*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 of that, a graphics processor, or a graphics processing unit (“GPU”) correlates to graphics core when integrated into a combined processor. (Wyatt Decl., ¶86.) Anderson further discloses CPU cores in

a multicore processor can be connected with a communication bus network that is a “ring interconnect.”




**FIG. 5B**

(*Id.*, 11:29-36, Fig. 5B; Wyatt Decl., ¶87.)

It would have been obvious to apply these teachings and include both CPU and GPU cores as well as a ring interconnect in Bose. (Wyatt Decl., ¶88.) Bose and Anderson are in the same field of art, as both are directed to thermal power

management in a multi-core processor. (*Id.*) 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 a multi-core processor like in Bose would often have a GPU core if graphics processing is required, as disclosed in Anderson. (*Id.*) A POSITA would have understood Bose's PMCU organization can accommodate 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 with the global PMCU continuing to function in the same way. (*Id.*) Thus, a POSITA would have had a reasonable expectation of success in applying the Bose power management mechanism to an integrated circuit with both a GPU core and a CPU core, as no substantial change is required and the components would continue to perform their same functions. (*Id.*)

A POSITA would have been motivated to apply Bose's power management technique in a processor with heterogeneous cores, including CPU and GPU cores, as disclosed in Anderson. (*Id.*, ¶89.) Anderson points out efficiently balancing thermal energy generation and system performance could be tricky with different types of processing elements. (Anderson, 1:64-67.) Because in Bose's system, each local component is associated with a local PMCU to monitor workload and adjust operating characteristics based on the power mode selected by global

PMCU, the individual features and capability of each component would be taken into consideration and accommodated. (Wyatt Decl., ¶89.)

A POSITA would have also been motivated to implement a ring interconnect as disclosed in Anderson in Bose's system. (*Id.*, ¶90.) A ring interconnect connects all the processing units in a circular arrangement and allows them to exchange data under certain communication protocols. 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 reduces the semiconductor die area needed for bus routing. (*Id.*) A POSITA would have recognized that a ring topology bus is not only more efficient in terms of physical space and ease of routing, it becomes necessary as more CPU and GPU cores and other sub-units are integrated into a compact package, such as the multi-core chip in Bose. (*Id.*) A POSITA would have had a reasonable expectation of success in implementing a ring interconnect in Bose system, as Anderson discloses such implementation in a similar multi-core processor system. (*Id.*)

**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.***

Bose (as modified by Anderson) in view of Naffziger renders obvious Claim 13. (Wyatt Decl., ¶¶91-97.) Naffziger, titled “Managing Current and Power in a Computing System,” was filed July 21, 2010 and is prior art at least under pre-AIA §102(e).

Naffziger is directed to a system and method for efficient power transfer on a die including two or more computing units (CUs). (Naffziger, Abstract.) Like Anderson, Naffziger discloses a processor with both GPU cores 172 and CPU cores 112. (*Id.*, 16:45-58, 17:9-13, Fig. 8; *see also id.*, 3:21-44.)





FIG. 8

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

A POSITA would have understood the CPU cores and GPU cores as disclosed in Naffziger are the same CPU cores and GPU cores described in the '494, which can form a consumer-producer relationship with either GPU cores being the producer and CPU cores being the consumer, or CPU cores being the producer and GPU cores being the producer. (Wyatt Decl., ¶93; '494, 7:62-8:17, 9:32-39, 15:8-11.) Naffziger discloses the GPU cores and CPU cores communicate via a crossbar switch and share cache memory subsystem. (Naffziger, 17:25-59,

Fig. 8.) 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 in a typical consumer-producer relationship, such as the GPU processing graphical primitives prepared by CPU and CPU processing the graphics output from the GPU. (Wyatt Decl., ¶93.)

Naffziger further teaches power management unit 130 selects one discrete power-performance state (P-state) for the cores based on their real-time power consumption (e.g., monitored activity level) when running an application or applications. (Naffziger, 6:33-54.) Each P-state specifies an operating frequency with P<sub>0</sub> including a maximum operating frequency and P<sub>4</sub> including a minimum operating frequency:



(*Id.*, Fig. 2 (annotated), 6:39-46.) Naffziger also discloses the power management unit continuously increments the P-state of a CU if its activity level is higher than what is expected for the current P-state, or decrements the P-state if its activity level is lower. (*Id.*, 7:65-8:2, 8:16-34, 9:29-64.) As shown above, decrementing the currently selected P-state entails reducing operating frequency of the CU, while incrementing the currently selected P-state entails increasing operating frequency of the CU. (*Id.*, Fig. 2.) 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, and the CPU is relatively inactive. (Wyatt Decl., ¶95.) Thus, a POSITA would have

understood that Naffziger teaches incrementing the GPU core to a higher P-state, i.e., increasing its operating frequency, while decrementing the CPU core to a lower P-state, i.e., decreasing its operating frequency. (*Id.*)

A POSITA would have been motivated to combine the teachings of Naffziger with Bose-Anderson. (*Id.*, ¶96.) Like Bose, Naffziger is directed to power management of an integrated circuit with multiple computation/processing units. (*Id.*) Both Naffziger and Bose disclose selecting and adjusting power mode based on monitored activity level of each component within a highly integrated multi-core package. (*Id.*) A POSITA would have understood both systems to be substantially the same, except that Naffziger specifically discloses its system can accommodate many different types of processing cores. (*Id.*) Accordingly, 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. (*Id.*)

## **B. Grounds Based on White**

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

#### **a. Independent Claim 4**

White renders obvious Claim 4. (Wyatt Decl., ¶¶98-112.)

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

White discloses Claim 4[pre] to the extent limiting. (Wyatt Decl., ¶99.) White

discloses power management logic 200 in a multi-core integrated circuit 100:




FIG. 2A

(White, Fig. 2A (annotated), 8:59-9:1.) Power management logic 200 is configured to perform energy efficiency operations, as it adjusts operating characteristics for 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.) Power management logic 200 can be implemented in software. (*Id.*, 14:32-38, 16:56-63.) Thus, power management logic can be software code stored in a non-transitory storage medium, and when executed, cause the integrated circuit 100 to perform energy efficiency operations, such as adjusting operating frequency and/or voltage of logic cores so as to adjust overall power consumption of IC 100. (*Id.*)

**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;*

White discloses Claim 4[a]. (Wyatt Decl., ¶¶100-105.) First, White discloses a multi-core integrated circuit 100, which includes a first processing element, such as logic core 120A which “represent[s] processor cores that executes instructions supplied to integrated circuit 100”:




**FIG. 1**

(White, Fig. 1 (annotated), 6:1-56.) Integrated circuit 100 also includes a communication bus, such as common bridge logic 110, which connects logic cores 120A, 120B to at least one other device of integrated circuit 100, such as “bus 150, peripheral 160 and/or memory subsystem 170.” (*Id.*, 7:1-7, Fig. 1.)

Second, White discloses power management logic receives workloads for common bridge logic 110 and logic core 120A, such as their operating voltages and frequencies. As noted, 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.) Power management logic 200 may increase or reduce operating frequency/voltage for a component to increase or reduce its power consumption. (*Id.*, 9:36-53, 9:56-62, 15:49-54, 17:29-44; 18:28-39, 20:39-57.) In determining power consumption of a component, power management logic 200 “*determine[s] the current output voltage levels ...* by accessing data stored in configuration registers for voltage regulator 130.” (*Id.*, 10:43-46.) Power management logic 200 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.) A POSITA would have understood that operating voltage/frequency reflect the activity and use, or workload, of a component. (Wyatt Decl., ¶103.)

White also discloses that power management unit 200 receives operating voltages/frequencies based on the integrated circuit reaching a limit. (*Id.*, ¶104.) For example, White discloses receiving workload of the components based on the integrated circuit reaching the power supply limit, such as battery power limit.

(White, 15:54-64; *see also* 18:25-39, 20:6-12, 24:7-29.) Additionally, White discloses that power management logic 200 receives workload information based on the integrated circuit reaching a thermal design power limit, such as “*a predetermined maximum operating temperature*” above which the performance of the device may be erratic or damage may potentially occur. (*Id.*, 2:34-38.) White discloses the operating characteristics of the logic cores and common bridge logic are adjusted to maintain the operating temperature of the integrated circuit within “*a maximum temperature.*” (*Id.*, 2:39-44, 12:59-13:6.)

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

White renders Claim 4[b] obvious. (Wyatt Decl., ¶¶106-107.) White discloses that when it is detected “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.” (White, 9:36-49; *see also* 7:46-62, 9:25-30.)

A POSITA would have understood that, like the common bridge logic, each individual core can potentially be a bottleneck to overall performance, particularly when multiple logic cores are interfaced with one common bridge logic and share

the same cache memory sub-system. (Wyatt Decl., ¶107.) Thus, determining whether a logic core is a performance bottleneck would be a routine consideration in such a system. (*Id.*) Following White's teachings, a POSITA would have understood to determine whether the power consumption (i.e., workload) of logic core 120A similarly indicates it is a bottleneck to overall system performance, and power management logic 200 could adjust its operating characteristics accordingly to maintain optimal overall system performance. (*Id.*) Doing so would have required the same determinations and adjustments for when common bridge logic is the bottleneck. (*Id.*) 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 determined to be a bottleneck. (*Id.*) This would have simply been the application of a 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 the purpose of power management logic is managing system power, and White is already configured to adjust the voltage and frequency supplied to cores and to common bridge logic. (*Id.*)

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

White discloses Claim 4[c]. (Wyatt Decl., ¶¶108-109.) White discloses that when 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.” (White, 9:36-49, 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, power management logic 200 determines common bridge logic 110 (“communication bus”) is not able to accommodate its workload. (Wyatt Decl., ¶108.) White also discloses adjusting operating frequency/voltage of common bridge logic 110 in response to new operating voltage/frequency of logic cores, which also requires determining if common bridge logic 110 is able to accommodate its workload. (Wyatt Decl., ¶109; White, 17:29-44; *see also* 17:66-18:3, 20:39-57.)

**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**

White renders Claim 4[d] obvious. (Wyatt Decl., ¶¶110-111.) White discloses when “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.” (White, 9:36-49; *see also* 7:46-62.) Thus, White discloses reducing (“capping”) 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 common bridge logic 110 is the performance bottleneck. (Wyatt Decl., ¶110.)

A POSITA would have understood to reduce or cap operating frequency of common bridge logic 110 and operating frequency of logic core 120B when the power consumption of logic core 120A (“first processing element workload”) indicates that logic core 120A is the performance bottleneck. (Wyatt Decl., ¶111.) White discloses, with respect to common bridge logic 110, that when a component is determined to be a performance bottleneck, a change in voltage and/or frequency needs to be applied to that component to mitigate the bottleneck. (*Id.*) Thus, a POSITA would have understood that, in response to determining the logic core 120A is the bottleneck instead of the common bridge logic, White’s power management logic would apply a change in voltage and/or frequency to logic core

120A to optimize system performance. (*Id.*) And, 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 logic core 120A, the operating frequency and/or voltage of logic core 120B and common bridge logic 110 may need to be reduced so as to reduce overall power consumption. (*Id.*) Further, 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 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. (Wyatt Decl., ¶111.)

**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.**

White renders Claim 4[e] obvious. (Wyatt Decl., ¶112.) As discussed, White renders obvious capping operating frequency of common bridge logic 110 in response to determining the workload of logic core 120A indicates logic core 120A is the performance bottleneck. (*See supra* Claim 4[d].) Further, White discloses **raising** the operating voltage and frequency of common bridge logic 110 in

response to determining common bridge logic is not able to accommodate its workload. (*See supra* Claim 4[c].) White also discloses that when there is more power budget available, previously reduced operating characteristics, such as operating frequency, may optimally be raised again. (White, 23:12-22, 24:10-21.) A POSITA would have understood that when common bridge logic is determined not able to accommodate its workload after its operating frequency is reduced, more power resources need to be allocated to common bridge logic 110, such as by raising its previously reduced operating frequency. (Wyatt Decl., ¶112.) Thus, White renders obvious increasing a previous cap of the operating frequency of common bridge logic 110 (“communication bus”) in response to determining the workload or power-consumption of a logic core 120A (“first processing element”) indicates logic core 120A is the performance bottleneck and determining common bridge logic 110 is not able to accommodate its workload, and responding appropriately to effect optimizations. (*Id.*)

**b. Dependent Claim 7**

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

White discloses Claim 7. (Wyatt Decl., ¶113.) *See supra* 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.***

White in view of Anderson renders Claim 5 obvious. (Wyatt Decl., ¶¶114-119.)

First, White discloses code including privileged level power service code. (*Id.*, ¶¶115-116.) White discloses power management logic 200 (“code”) may be implemented in “hardware, firmware, and software or a combination thereof.” (White, 16:58-60.) A POSITA would have understood hardware code has higher privilege level, and similarly firmware code has a privilege level to be executed above ordinary software. Thus, White’s power service code has a privileged level. (Wyatt Decl., ¶115.) Power management in a computer system requires robust and higher-privilege operation that is independent of the operating system and software. Power management must be implemented in the highest privileged level to ensure immediate robust control, especially where it is anticipated bottlenecking situations can occur such as in White. (*Id.*, ¶116.)

Second, while White does not specifically disclose 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's multi-core processor includes GPU cores, CPU cores and a ring interconnect. *See supra* Claim 2 (addressing how Anderson discloses the processor including both CPU cores and GPU cores as well as ring interconnect). It would have been obvious to a POSITA to apply these teachings of Anderson to White to include both CPU cores and GPU cores. (Wyatt Decl., ¶117.) 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, White's power management logic can accommodate such a system because it independently adjusts operating characteristics of each component in adjusting the overall power consumption. (*Id.*) The GPU core disclosed in Anderson would be treated as other core elements, like CPU cores and the common bridge logic. (*Id.*) A POSITA would also have had a reasonable expectation of success in applying the White power management mechanism to an integrated circuit with both GPU and CPU cores, as no substantial change is required. (*Id.*) Further, a POSITA would have been motivated to apply White's power management technique in a processor with heterogeneous cores as disclosed in Anderson, because the individual features and capability of

each component would be accommodated with White's power management system, and including GPU cores would add known graphics functionality to White's system. (*Id.*, at ¶118.)

A POSITA would have been motivated to implement a ring interconnect as disclosed in Anderson in White's system. (*Id.*, at ¶119.) As discussed, White's common bridge logic 108 interconnects the key elements in the integrated circuit, which would typically be located at separate locations inside the chip. The progression of more highly-integrated multi-core semiconductor devices favors use of a ring bus topology for interconnecting, as doing otherwise would waste useful semiconductor die area. *See supra* 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 White's system, as Anderson discloses such implementation in a similar multi-core processor system. (Wyatt Decl., ¶119.)

**b. Dependent Claim 6**

***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;***

White (as modified by Anderson) renders Claim 6[a] obvious. (Wyatt Decl., ¶¶120-1222.) White discloses that when its 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 to mitigate bottlenecking, and the voltages and/or frequencies of other components need to be reduced so as to counterbalance the increase. (White, 9:36-53.) *See also supra* Claim 4[d], 4[e]. 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 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. (Wyatt Decl., ¶122.) As explained above, a POSITA would have also understood that in White’s system, as modified by Anderson, the first core could be a processor core, and the second core could be a graphics core. (*Id.*)

***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.***

White (as modified by Anderson) renders Claim 6[b] obvious. (Wyatt Decl., ¶123.) White discloses 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 to mitigate 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* Claim 4[d], 4[e], 6[a]. Thus, a POSITA would have understood in White's system as modified by Anderson, when GPU (first processing element") is a bottleneck, frequency of other component such as the processor cores can be decreased (i.e., capped) so the excess power is allocated to the bottleneck component to improve overall system performance. (Wyatt Decl., ¶123.) Also, a POSITA would have understood that when the communication bus is able to accommodate its workload, the communication bus is not a bottleneck, does not need additional power resources, and will not be affected by capping frequency of the processor core. (*Id.*)

**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.***

White (as modified by Anderson) in view of Naffziger renders Claim 17 obvious. (Wyatt Decl., ¶¶124-126.) As discussed, Naffziger discloses consumer-producer relationships between elements in a processor. *See supra* Claim 13 (discussing how Naffziger discloses a system with consumer-producer relationships between processor elements). It would have been obvious for a

POSITA to combine the teachings of Naffziger with White-Anderson. (Wyatt Decl., ¶125.) Like White, Naffziger is directed to power management of an integrated circuit with multiple processing units. Both Naffziger and White disclose adjusting operating characteristics of components based on monitored activity level of each component within a highly integrated multi-core package. (*Id.*) 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, a consumer/producer relationship exists between common bridge logic and the logic cores with common bridge logic being the producer based on the teachings of Naffziger. As a “producer,” common bridge logic manages the source of data and instructions to be consumed by the cores. (Wyatt Decl., ¶125.) A POSITA would have had reasonable expectation of success in implementing the teachings of Naffziger to White’s system (as modified by Anderson to include both GPU cores and CPU cores). (*Id.*) A POSITA would have been motivated to do so because increasing the frequency of the producer core and reducing the frequency of the consumer core would create an interactive and efficient system with integrated CPU and GPU cores. (*Id.*, ¶126) *See also supra* 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.***

White (as modified by Anderson) in view of Naffziger renders claim 18 obvious. (Wyatt Decl., ¶127.) *See supra* Claim 14 (addressing how White discloses increasing communication bus frequency by a different amount than which 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**

White in view of Simeral renders obvious Claim 1. (Wyatt Decl., ¶¶128-145.)

**1[pre]: A processor comprising:**

White discloses Claim 1[pre] to the extent limiting. (*Id.*, ¶129.) White discloses a computer including a “a microprocessor 1100” (a “processor”):



(White, Fig. 8 (annotated), 2:15-18, Abstract, 6:23-25, 6:56-63, 24:30-41, Fig. 1.)

**1[a]: an integrated circuit including:**

White discloses Claim 1[a]. (Wyatt Decl., ¶130.) For example, White discloses microprocessor 1100 includes an integrated circuit 100 as shown in Fig. 1:




FIG. 1

(*Id.*, Fig. 1 (annotated), 24:37-41).

**1[b]: a first core;**

White discloses Claim 1[b]. (Wyatt Decl., ¶131.) White discloses the integrated circuit 100 includes a first core, such as logic core 120A:





FIG. 1


(White, Fig. 1 (annotated), 6:41-56.)

**1[c]: a cache memory;**

White discloses Claim 1[c]. (Wyatt Decl., ¶132.) White discloses an L2 cache 1128. (*Id.*, 25:10-20, 41-49, Fig. 8) White discloses that L2 cache 1128 can be “*integrated onto a semiconductor substrate with microprocessor 1100.*”



(White, Fig. 8 (annotated), 25:45-49.) Further, the integrated circuit 100 includes an internal memory subsystem 170:



(*Id.*, Fig. 1 (annotated), 7:1-7.) 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. (Wyatt Decl., ¶132.)

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

White discloses Claim 1[d]. (Wyatt Decl., ¶133.) 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.) The common bridge logic is an interconnect that effectively bridges and coordinates multiple access from multiple cores. (*Id.*, 7:1-7; Wyatt Decl., ¶134.) Because common bridge logic 100 interconnects the logic cores 120A and 120B with other elements, it is a communication bus as defined in the ’494. (’494, 15:38-39.)

White further discloses that common bridge logic 110 connects logic core 120A (“first core”) with memory subsystem 170. (White, 7:1-7.) As discussed above, memory subsystem 170 can be internal to the integrated circuit 100 and includes L2 cache. *See supra* Claim 1[c]. Thus, White discloses that common bridge logic 110 (“communication bus”) connects logic core 120A (“first core”) and a cache memory:




FIG. 1

(*Id.*, Fig. 1 (annotated), 7:1-7, 2:45-67; Wyatt Decl., ¶133.)

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

White discloses Claim 1[e]. (Wyatt Decl., ¶134.) 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* 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. (Wyatt Decl., ¶134.)




**FIG. 2A**

(White, Fig. 2A (annotated))

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

White discloses Claim 1[f]. (Wyatt Decl., ¶135.) The power management logic 200 is also a bus workload monitor that monitors the operating voltage and frequency (“bus workload”) of common bridge logic 110 (“communication bus”).

(*Id.*) See also *supra* Claim 1[e], 4[a].

Alternatively, to the extent White is determined not to disclose such a bus workload monitor, White in view of Simeral renders this obvious. 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 GPIS port k 191 as shown in Fig. 1A below

is a bus workload monitor. (Simeral, [0020], Fig. 1A.) Simeral discloses a data path controller that provides power management functions, and includes a general-purpose interface signal (“GPIS”) port (“k”) 191 that “can monitor activity associated with bus 130.” (*Id.*, [0007], [0019], Fig. 1.) Simeral also explains the term “activity level” means “a quantifiable **load** associated with a component.”




FIG. 1A

(*Id.*, [0020], Fig. 1A (annotated).)

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. (Wyatt Decl., ¶137.) Both White and Simeral are directed to power management for processing systems including similar components, and both can dynamically control supply of power to the similar components based on workload data of each component. (*Id.*) 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 overall power consumption. (Wyatt Decl., ¶137.)

A POSITA would have been motivated to combine the teaching of Simeral and White. (*Id.*, ¶138.) 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. (*Id.*) Simeral teaches comparing the monitored workload with a threshold in order to make the adjustment, and a POSITA would have understood that it is necessary for the power management logic 200 in White to make these same comparisons so as to adjust the operating characteristics of the bus for purpose of meeting a performance goal while conserving power consumption. (*Id.*)

A POSITA would also have reasonable expectations of success in combining the teachings of White and Simeral. (*Id.*, ¶139.) 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. (Wyatt Decl., ¶139.) A POSITA would have thus understood the systems to be complementary. (*Id.*) 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. (Wyatt Decl., ¶139.)

***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,***

White in view of Simeral renders Claim 1[g] obvious. (Wyatt Decl., ¶¶140-145.)

First, White discloses a power management logic 200 (“balancing control”) that ***monitors*** the 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.” (Wyatt Decl., ¶141.) 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* Claim 4[a].

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. (Wyatt Decl., ¶142.)

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. (Wyatt Decl., ¶143.) 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].)

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. (Wyatt Decl., ¶144.) *See also supra* 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 the purpose of maximizing overall system performance (such as alleviating bottlenecking situation) while maintaining overall power consumption. *Id.*

**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.*

White discloses Claim 1[g]. (Wyatt Decl., ¶145.) 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 common bridge logic, to remain above a low limit, so as to maintain its operation. (Wyatt Decl., ¶145.)

**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.*

White in view of Simeral renders obvious Claim 3. (Wyatt Decl., ¶146.)

White in view of Simeral renders obvious Claim 3[a]. *See supra* VIII.B.4.a(Claim 1[g]). Further, as discussed above, White in view of Simeral discloses Claim 3[b] and Claim 3[c]. *See supra* Section VIII.A.2.a(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.*

White in view of Simeral renders obvious Claim 14. (Wyatt Decl., ¶147.)

*See supra* Section VIII.A.2.b(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.***

White in view of Simeral renders obvious Claim 15. (Wyatt Decl., ¶148.) As discussed, White discloses a second logic core 120B in addition to the first logic core 120A, and that the power management logic 200 (“balancing control”) adjusts power allocation between logic cores 120A, 120B and common bridge logic 110 based on system power limit and the workload of each component. *See infra* VIII.B.5(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. (Wyatt Decl., ¶148.) White also discloses that second logic core 120B and common bridge logic 110 (“communication bus”) have power planes tied together. (*Id.*) *See supra* Section VIII.A.2.c(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.***

White (as modified by Simeral) in view of Anderson renders obvious Claim 2. (Wyatt Decl., ¶149.) First, White discloses a second logic core 120B. (White, 6:41-49.) Power management logic 200 monitors and determines the operating voltages and frequencies of components including both logic cores 120A, 120B, and is thus a core workload configured to determine a first core workload for the first core. *See supra* VIII.B.4.a(Claim 1[e]). A POSITA would have understood power management logic 200 equally monitors and determines the operating voltages and frequencies of second core 120B, and so is also a second core workload monitor that determines a second core workload. (Wyatt Decl., ¶149.) 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. (Wyatt Decl., ¶149.) *See supra* VIII.B.4.a(Claim 1[g]), 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. (Wyatt Decl., ¶149.) *See supra* 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.***

White (as modified by Simeral) in view of Anderson and Naffziger renders obvious Claim 13. (Wyatt Decl., ¶150.) *See supra* Claim 17.

**IX. TRIAL SHOULD BE INSTITUTED**

**A. *Fintiv* Factors**

The Board should institute review because Petitioner has been diligent in filing the Petition, and the circumstances surrounding the parallel district court cases and International Trade Commission investigations do not warrant discretionary denial under § 314(a). *Sotera Wireless, Inc. v. Masimo Corp.*, IPR2020-01019, Paper 12 at 11-21 (PTAB Dec. 1, 2020) (precedential) (citing *Apple, Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (PTAB Mar 20, 2020) (precedential)). A *Fintiv* analysis shows that discretionary denial is not applicable here, as Petitioner has been dismissed from any pending litigation in district court and terminated from any pending ITC investigations. *See, e.g., Boss Corp. v. Koss Corp.*, IPR2021-00680, Paper 15 at 14 (PTAB Oct. 13, 2021). Moreover, any

parallel district court and ITC proceedings were dismissed early-on, before *Markman* or any substantive decisions. The remaining district court and ITC actions do not include Petitioner, and it is purely speculative that any tribunal will address the distinct invalidity challenges raised in this Petition.

**B. *General Plastic***

The *General Plastic* factors do not weigh in favor of denying institution. *Gen. Plastic Indus. Co. v. Canon Kabushiki Kaisha*, IPR2016-01357, Paper 19 (PTAB Sept. 6, 2017) (Section II.B.4.i precedential).

**Factor 1** examines whether the same petitioner or a related petitioner previously filed a petition directed to the same claims of the same patent. *General Plastic at 16*; see also *Valve Corp. v. Elec. Scripting Prods., Inc.*, IPR2019-00062, Paper 11 (PTAB Apr. 2, 2019) (precedential) (“*Valve I*”); *Valve Corp. v. Elec. Scripting Prods., Inc.*, IPR2019-00064, Paper 10 (PTAB May 1, 2019) (precedential) (“*Valve II*”). Factor 1 strongly favors institution. Petitioner here has not filed any previous petitions challenging the ’494. While both this petition and the Qualcomm IPR challenge Claims 1, 3, 14 and 15, this Petition *also* challenges Claims 2, 4-7, 13, and 17-18, which are *not* at issue in the Qualcomm IPR. Moreover, Petitioner has no relationship with Qualcomm and Samsung (the original petitioners in the Qualcomm IPR), and Petitioner played no role in the preparation or filing of the Qualcomm IPR. Petitioner was not a party to the

pending parallel ITC case involving Qualcomm (and in fact is not a party to any currently pending case), and Petitioner was sued separately on distinct products in ITC Inv. No. 337-TA-1332 as compared to the products at issue in the pending ITC case against Qualcomm. Thus, unlike the facts in *Valve*, there is no evidence that there is a “significant relationship” between Petitioner and Qualcomm.

**Factors 2, 4, and 5** examine Petitioner’s knowledge of the prior art, and the elapse of time between filings of the petitions. These factors weigh in favor of institution. This is Petitioner’s first petition against the ’494, and thus the issue of whether Petitioner could have raised the prior art at issue here in the earlier Qualcomm IPR is not relevant (Petitioner, having no relationship with Qualcomm or the Qualcomm IPR, could not have done so). Further, this Petition relies on distinct combinations of references as compared to those used in the Qualcomm IPR. This Petition relies on different combinations of references to challenge the claims as compared to the Qualcomm IPR. For example, this Petition challenges: i) claim 2 over the combination of Bose and Anderson; ii) claim 13 over the combination of Bose, Anderson and Naffziger; iii) claims 4, 7 over White, iii) claims 5-6 over the combination of White and Anderson, iv) claims 17-18 over the combination of White, Anderson and Naffziger, v) claim 2 over the combination of White, Simeral and Anderson, and vi) claim 13 over the combination of White, Simeral, Anderson and Naffziger, which are grounds not at issue in the Qualcomm IPR. Moreover, there

is no unreasonable delay in the filing of this petition, as Petitioner is not a party to any pending litigation involving the '494.

**Factor 3** examines “roadmapping” concerns, which are not present here. This Petition challenges claims not at issue in the Qualcomm IPR, involves different combinations of references than the Qualcomm IPR, Petitioner had no involvement in the Qualcomm IPR, and no Institution Decision has issued in the Qualcomm IPR.

**Factors 6 and 7**, addressing the finite resources of the Board, favor institution or at most neutral. This is Petitioner’s first Petition challenging the '494, and the claims and grounds at issue in the Petition provide a basis for timely resolution of this review.

**C. 35 U.S.C. § 325(d)**

The Board should decline to exercise discretion under § 325(d). None of the combinations of references presented herein were considered by the Office during Examination. And, as discussed above, this Petition presents distinct challenges to claims of the '494 from those at issue in the Qualcomm IPR.

Should Patent Owner raise any discretionary denial challenge to the Petition, Petitioner intends to respond to any counter-arguments Patent Owner may raise.

**X. CONCLUSION**

Petitioner respectfully requests that a Trial be instituted. Claims 1-7, 13-15, and 17-18 of the '494 should be cancelled as anticipated or obvious over the prior art.

Dated: August 23, 2023

Respectfully submitted,

/s/ Celine J. Crowson

---

Celine Jimenez Crowson (Reg. No. 40,357)

Joseph J. Raffetto (Reg. No. 66,218)

Scott Hughes (Reg. 68,385)

Ryan Stephenson (Reg. 76,608)

Nicholas Rotz (Reg. No. 75,959)

HOGAN LOVELLS US LLP

555 13th Street N.W.

Washington, D.C. 20004

Telephone: 202-637-5600

Facsimile: 202-637-5910

Helen Trac (Reg. No. 62,250)

HOGAN LOVELLS US LLP

Four Embarcadero, #3500

San Francisco, CA 94111

Telephone: 415-374-2300

Facsimile: 415-374-2399

*Counsel for Petitioner Mercedes-Benz USA, LLC*

### Table of Exhibits

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

### **Certificate of Word Count**

I certify that, in accordance with 37 C.F.R. § 42.24(d), the word count for the foregoing Petition totals 13,993. This is less than the permitted total under the 37 C.F.R. 42.24(a)(1)(i).

/s/ Celine J. Crowson

Celine J. Crowson (Reg. No. 40,357)

*One of the attorneys for Petitioner Mercedes-Benz  
USA, LLC*

### **Certificate of Service**

The undersigned certifies that, in accordance with 37 C.F.R. § 42.6(e) and 37 C.F.R. § 42.105, the foregoing Petition for Inter Partes Review, and Exhibits 1001 through 1009 were served on August 23, 2023, by Overnight Courier at the following address of record and addresses known to Petitioner as likely to effect service and the address of record for the '494 Patent:

164 - KINNEY & LANGE, P.A.  
333 S. 7TH ST., SUITE 2700  
MINNEAPOLIS, MN 55402-2438

/s/ Joseph J. Raffetto  
Joseph J. Raffetto (Reg. No. 66,218)  
*One of the attorneys for Petitioner*  
*Mercedes-Benz USA, LLC*