

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

QUALCOMM INCORPORATED,
Petitioner,

v.

DAEDALUS PRIME LLC,
Patent Owner.

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

**PATENT OWNER'S PRELIMINARY RESPONSE
PURSUANT TO 37 C.F.R. § 42.107(A)**

TABLE OF CONTENTS

I.	Introduction.....	1
A.	Overview of the '494 Patent.....	2
II.	Overview of Select Art	7
A.	Overview of Bose	7
B.	Overview of Naffziger	8
C.	Overview of White	9
D.	Overview of Simeral	9
III.	Ground 1 Fails	10
A.	Bose does not render obvious Claim 1[d]	10
B.	Bose does not render obvious Claim 1[g.ii].....	11
C.	Bose does not render obvious Claim 1[g.iii].....	13
IV.	Ground 2 Fails	17
V.	Ground 3 Fails	17
A.	A POSITA Would Not Have Been Motivated to Combine Bose with Naffziger.....	17
B.	Bose in view of Naffziger does not render obvious Claim 1[d] ..	20
C.	Bose in view of Naffziger does not render obvious Claim 1[g.ii]	21
D.	Bose in view of Naffziger does not render obvious Claim 1[g.iii]	22
VI.	Ground 4 Fails	22
A.	A POSITA Would Not Have Been Motivated to Combine White with Simeral	23
B.	White in view of Simeral does not render obvious Claim 1[d] ..	26
C.	White in view of Simeral does not render obvious Claim 1[e]...	27
D.	White in view of Simeral does not render obvious Claim 1[f] ...	28
E.	White in view of Simeral does not render obvious Claim 1[g.ii]	30
F.	White in view of Simeral does not render obvious Claim 1[g.iii]	32

Conclusion 33

TABLE OF AUTHORITIES

Page(s)

Federal Case

<i>Johns Manville Corp. v. Knauf Insulation, Inc.</i> , IPR2018-00827, Paper 9.....	19, 24
--	--------

Regulations

37 C.F.R. § 42.24(a).....	35
37 C.F.R. § 42.24(d)	35

I. Introduction

U.S. Patent No. 8,898,494 (Ex-1001, “the ’494 Patent”) improved power management techniques in integrated circuits, such as computer processors, by balancing power allocation between processing cores and communication buses. *See, e.g.*, Ex-1001, 1:20–57. Whereas previously, power management policies were implemented separately amongst system-on-chip components (*e.g.*, each component received a power budget and implemented its own power savings techniques), or merely in a uniform blanket fashion across the integrated circuit (*e.g.*, components were uniformly capped), the inventors of the ’494 Patent devised of an integrated circuit that improved system performance by balancing allocation between processing components such as a communication bus and a processor core, based on a power limit and a workload threshold comparison. *See* Ex-1001 at 1:45-56. Using the systems and methods described in the ’494 Patent, processors are able to increase overall performance without increasing power consumption.

The cited art disclosed by Qualcomm Incorporated (“Petitioner”¹) fails to render obvious the claims of the ’494 Patent, and the Petition should be denied in its entirety².

A. Overview of the ’494 Patent

The ’494 Patent describes a processor that efficiently balances performance and power between one or more cores and a communication bus based on measured workloads to increase system performance while operating within thermal limits. Ex-1001, Abstract, 15:4-21.

For example, the ’494 Patent describes that if the workload of one component in a processor system is too high, then that component’s frequency/power may be increased. Ex-1001, 17:8–11. The ’494 Patent explains that a component’s “workload” is its processing activity:

¹ As filed, the Petition identified Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc. (collectively “Samsung”) along with Qualcomm Incorporated as Petitioners. However, on June 2, 2023, Patent Owner and Petitioner Samsung jointly moved to terminate as to Samsung. Paper No. 10. On June 27, 2023, The Board granted the motion and terminated the proceedings as to Samsung, and noted that Qualcomm Incorporated would remain as Petitioner. Paper No. 11.

² Petitioner challenges independent claim 1 as allegedly obvious only in Grounds 1, 3, and 4. Thus, this Preliminary Response primarily addresses those Grounds; each of the challenged dependent claims is non-obvious at least because it depends from a non-obvious independent claim. Patent Owner reserves the right to address all arguments made in the Petition, including those made for other Grounds, in further briefing.

Determining a workload includes any known method for determining activity of hardware, a workload of software, a number of instructions executed or to execute, or any other metric for determining use of a processing element. As one example, a workload for core 505 includes an amount time core 505 is active during a period of time.

Ex-1001, 9:47-57. The inventors of the '494 Patent thus determined that—contrary to prior power management approaches which would cap or cut power supplied to high-activity processor elements (because those are the elements consuming high power)—**increasing** power supplied to high-activity elements improves overall system functionality by reducing bottlenecking. However, because the amount of power available within an integrated circuit is limited, the increase in power supplied (e.g., by increasing frequency) to high-activity processor components must come from a limited power budget. The inventors of the '494 Patent therefore determined that the power allocation must be tuned between the components. The '494 Patent thus explains:

In one embodiment, power control module 160, which may include hardware, firmware, software, or a combination thereof, is able to balance and allocate power/frequency between devices of IC 100 to achieve maximum performance; even when IC 100 is under a power limit. Additionally, such allocation, in one scenario,

is based on actual workload of devices, instead of assumptions of workload or power consumption. For example, a power limit (either a maximum power limit or a reduced limit during power saving modes) is provided for IC 100. Previously, a maximum current (and as a corollary frequency) was reduced for all processing elements of IC 100 to meet the power limit. Instead, in this scenario, a device that is determined to be a bottleneck for performance, such as throughput performance as viewed from the perspective of an application, is allocated more current/frequency, while the other competing device is capped/limited to ensure the power limit is still met. In other words, ***a dynamic balance is continuously struck between devices to ensure the device that needs more resources for better overall performance is allocated such extra resources, while other devices are ‘balanced’ (reduced in performance) to meet a given power limit.***

Ex-1001, 7:40-61.

As described in the '494 Patent, “in one embodiment, power (e.g., current, voltage, frequency, or other metric affecting power consumption) is balanced between one or more cores and a communication bus.” Ex-1001, 15:18–21; 16:10–14. The '494 Patent describes that when the communication bus has high workload, increasing the communication bus frequency may need to be “counterbalanced in the CPU.” Ex-1001, 17:11–15. The '494 Patent describes that “any number of

balancing effects may be utilized,” such as by decreasing the frequency for the CPU cores. Ex-1001, 17:16–19. “Therefore, a power increase for communication bus 950 ... is dynamically balanced with a cap of CPU cores 905 ... to alleviate the bottlenecks in the system to increase performance, while not exceeding the current limit of [the integrated circuit] 900.” Ex-1001, at 17:35–41.

Figure 9 illustrates one embodiment:




FIG. 9

Ex-1001, Figure 9.

As described by the '494 Patent, an integrated circuit 900 comprises multiple cores (e.g., CPU cores 905), cache or memory 960, communication bus 950, and

power management service 920. *See, e.g.*, Ex-1001, Fig. 9, 2:21-24, 15:48-52, 16:10-41, 16:57-59, 17:11-15. Each core has an associated workload monitor. Ex-1001 at Fig. 9, 9:47-52, 16:14-56. The '494 Patent also describes a system in which communication bus 950 connects cores to shared resources such as cache or memory 960, and includes its own workload monitor. Ex-1001, Fig. 9, 15:38-67.

Power management service 920 balances power among the components based on their workloads. Ex-1001 at Fig. 9, 16:57-59. For example, the '494 Patent discloses increasing the communication bus's frequency/power if its workload is too high, alleviating a bottleneck between the cores and cache, and counterbalancing that increase by decreasing the frequency/power of the cores. Ex-1001, 17:4-41; *see also* Ex-1001 at Fig. 10, 18:35-19:46.




FIG. 10

Ex-1001, Fig. 10. In other examples, the core's power/frequency may be increased and be counterbalanced by a decrease in the power/frequency of another element to maintain operation of the integrated circuit below its thermal limit. Ex-1001 at 19:3-63.

II. Overview of Select Art

A. Overview of Bose

U.S. Pat. No. 7,421,601 (Ex-1005, "Bose") generally refers to "systems and methods for a power/performance monitoring and control unit." Ex-1005 at 1:16-22. Bose's system includes "a power-performance monitor and control unit (PMCU)," that "chooses appropriate power optimizations for each core elements 112-1, 112-2, 112-R, [cache hierarchy] 106, and [interconnect] 108." Ex-1005, at 7:8-11. Bose's

approach to increased performance is across-the-board increased power consumption:

When the main objective of the system is performance, the global PMCU 110 will choose the power modes and assign power for each core element 112-1, 112-2, 112-R, 106, and 108 such that all elements can run at the highest possible performance within the budget.

Ex-1005, 7:24-28;

If the actual total power is below or equal to the total power budget, the loop transitions through edge 308 to state 312, which commands all the cores to run in the highest power mode and divide the power budget equally among all the cores.

Ex-1005:9:6-10.

B. Overview of Naffziger

U.S. Pat. No. 8,510,582 (Ex-1007, “Naffziger”) generally refers to a system for “efficient die power management.” Ex-1007 at 1:8. The system includes a microprocessor with its cores being “an example of a computation unit 120.” Ex-1007, 5:49-50. A computation unit 120 determines its power consumption during a time interval. Ex-1007, 4:12-14. The system also includes a “power management unit” that, if the core power consumption is known, “select[s] one discrete power-performance state (P-state) from server given discrete P-states.” Ex-1007, 6:33-39.

C. Overview of White

U.S. Pat. No. 7,263,457 (Ex-1008, “White”), generally refers to “operation of and communication between elements of an integrated circuit operating at independent frequencies and/or voltages.” Ex-1008 at 1:9-12. The computer system of White may “detect a need to change the operating frequency for a logic core” based on, for example, when the system transitions to battery power, when the system transitions away from battery power, a determination to adjust the operating temperature of one or more cores. Ex-1008 at 15:44-49; 15:42-16:56. The system may then determine a new operating frequency for a core, adjust its voltage to be compatible if necessary, set the new operating frequency, and adjust common bridge logic for the new frequency. Ex-1008, 16:57-18:16; Fig. 4, 410-460. White’s approach to decreasing power consumption is across-the-board power supply reductions.

D. Overview of Simeral

U.S. Pat. Appl. Publ. No. 2009/0150689, (Ex-1009, “Simeral”) generally refers to “integrating a power management controller into a data path controller.” Ex-1009 at [0002]. The data path controller “monitor[s] the activity level of the bus (e.g., the data transfer rate of the bus), and then disabl[es] (e.g., shut[s] down) components that have very low (i.e., negligible) activity levels to save power.” Ex-1009.

III. Ground 1 Fails

The Petition asserts in Ground 1 that Claims 1 and 12 of the '494 Patent are rendered obvious by Bose.

A. Bose does not render obvious Claim 1[d]

Petitioner failed to demonstrate that Bose discloses or suggests a “communication bus to connect the first core and the cache memory” as claimed.

For example, Petitioner does not assert that Bose discloses a communication bus to connect the first core and the cache memory, instead alleging that “it would have been obvious to a person of ordinary skill in the art (“POSITA”) that Bose’s interconnect 108 couples and connects to each core and to L2/L3 cache hierarchy” Pet., p.25. However, Bose does not disclose that the cores are connected to the logic cache via the specific “interconnect 108” that Petitioner aligns with the claimed “communication bus.” To the contrary, FIG. 1 illustrates the interconnect 108 is ***specifically avoided*** by the connection 122 between the cores 118 and the cache (via the global PMCU 110). Ex-1005, Fig. 1.

Bose is thus distinguishable from the claimed system, in which a specifically claimed “communication bus” is required “to connect the first core and the cache memory,” have its workload monitored by the bus workload monitor, and have its power allocation dynamically tuned.

B. Bose does not render obvious Claim 1[g.ii]

Petitioner fails to demonstrate that Bose discloses or suggests “dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit” as claimed.

For example, Petitioner asserts that Bose discloses that “the global PMCU 110 will choose the power modes and assign power for each core element,” but fails to demonstrate that Bose discloses tuning power allocation “between” the components as claimed. For example, Petitioner merely asserts that “a POSITA would have understood such power optimizations as the global PMCU dynamically tuning power allocation” and that “the peak temperature limit of the chip” discloses the power limit for the integrated circuit. Pet., p.32. However, adjusting the cores, the L2/L3 cache, and the interconnect “such that all elements can run at the highest possible performance within the budget” is not “dynamically tun[ing] power allocation **between** the first core and the communication bus” as claimed. For example, tuning allocation between components involves adjusting the distribution amongst them, not merely adjusting them both the same way in the same amount. As the '494 Patent explains, “a dynamic balance is continuously struck between devices to ensure the device that needs more resources for better overall performance is allocated such extra resources, while other devices are ‘balanced’ (reduced in performance) to meet a given power limit.” Ex-1001, 7:40-61.

As one example, power module 520 is to increase frequency for core 505 and/or decrease frequency for core 510 in response to the dynamically determined workload for core 505 being greater the workload for core 510. Here, if one core is demonstrating a greater need of performance due to a higher workload, then its frequency is increased. And when an integrated circuit including cores 505 and 510 are up against a power limit, then the core demonstrating less need for performance requirements is decreased in frequency to counterbalance/compensate for the increase in frequency to the other core. In another example, power module 520 is to increase frequency for core 505 and/ or decrease frequency for core 510 in response to determining core 505 is a performance bottleneck. In other words, core 505 is limiting overall performance, because core 510 is relying on output from core 505 to continue further execution. But core 505 is operating to slowly. Then, core 505's frequency is increased to ensure overall performance is accelerated. Yet again, if a power limit is encountered, then core 510's frequency may be similarly decreased to ensure the power limit is enforced.

Ex-1001, 10:25-44.

Accordingly, Bose does not disclose or suggest “dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit” as claimed.

C. Bose does not render obvious Claim 1[g.iii]

Petitioner failed to demonstrate that Bose discloses or suggests “a bus workload threshold,” as claimed. Petitioner presents two conflicting assertions as to what allegedly satisfies the requirement for a “workload threshold”: (1) the “power budget,” and (2) “the predicted values included in the Trade-off Table 206.” Pet., p.33. Both fail.

First, Petitioner does not explain how or why it asserts that Bose’s “power budget” serves as the claimed “workload threshold.” See Pet., pp.32-33. Instead, Petitioner merely recounts that as disclosed in Bose, “[t]he global PMCU is responsible for assigning the budget to each core element,” and then proclaims that a “POSITA would have thus understood Bose’s power budget as defining the workload threshold for a given workload operating in an associated power mode.” Pet., p.33. However, Petitioner does not explain how the power budget, a function of the physical restrictions of the overall system (e.g., thermal restrictions) allegedly discloses a workload threshold for each processing component that is used to determine whether to increase power supplied to that component. Petitioner is conflating workload with power, which as described above is incorrect in view of

the '494 Patent's descriptions, but is also in conflict with Bose's own disclosures. For example, Bose itself distinguishes between "power" and "workload": "One objective of PMCU is to address the challenge of reducing **power** by exploiting the variability within a **workload** and between different workloads." Ex-1005, 4:31-33.

Second, Petitioner does not explain how or why it asserts that Bose's "predicted values included in the Trade-off Table 206" serve as the claimed "workload threshold." Instead, Petitioner merely asserts that a change in power mode of a core "may be based on the defined power modes in Trade-off table 206, which would also have been considered another workload threshold." Pet., p.34. However, Petitioner does not explain how the trade-off table, which "includes the predicted power and performance for each power mode" serves as a workload threshold, rather than, for example, a predicted current performance. See Ex-1005, 8:36-44.

Petitioner has also failed to demonstrate that Bose discloses or suggests "a comparison between the bus workload and a bus workload threshold," as claimed. That is, even if Bose's power budget served as the claimed "workload threshold" (it does not), that would still not satisfy the requirement for "a comparison" to "the bus workload."

For example, Bose discloses that Bose's power budget and power modes are calculated for only for **the processor cores**, not **the interconnect** that Petitioner

aligns with the claimed “communication bus.” For example, Bose states that *each core* has a “power budget” and “power mode” calculated by the Global PMCU. Ex-1005, 8:35-40. Indeed, as Petitioner concedes, “Bose discloses that the global PMCU calculates the power budget and power mode for each core ***but does not explicitly state the same for the interconnect.***” Pet., p.35. Bose does not disclose that the global PMCU calculates the power budget and power mode for the interconnect by comparing interconnect’s local PMCU data against a pre-determined threshold for the interconnect. Instead, Bose’s disclosure is limited to calculating the power budget and power mode for *each core*. As required by claim 1 of the ’494 patent, there must be a comparison between the “bus workload” and the “bus workload threshold.” See Pet., 35; see also Ex-1002 ¶123.

Petitioner incorrectly concludes that despite the admitted deficiency in Bose’s disclosure, a POSITA would have understood Bose to have disclosed “comparison between the bus workload and a bus workload threshold” on the basis that the Phase History Table 213 receives history information, performance statistics, and performance events “[f]rom all Local PMCUs,” which includes local PMCU 116 on interconnect 108. Ex-1005, 6:56-60; Fig. 2. However, Petitioner does not explain why that is the case, given Bose did not disclose or provide a reason why data from local PMCU 116 for interconnect 108 would be included in a similar calculation despite having the information available. By contrast, in the ’494 Patent, the claimed

communication bus plays an integral role in the processing throughput of the overall system (and therefore for which it makes sense to monitor and evaluate its workload). In Bose, the interconnect 108 does not perform a similar function. For example, the '494 Patent explains the criticality of the workload capabilities of the communication bus:

Assume that communication bus 960 is operating at 100% workload...GPU...97% workload activity, and CPU...at a 80% workload activity...the biggest bottleneck in the system is the communication bus 950, as it's fully occupied. But GPU cores 910 also have a high workload activity in comparison to CPU cores 910. As a result, if IC 900 is up against a power, thermal, or other limit, CPU cores 905 may be capped or limited...while power to GPU cores 910 are slightly capped, maintained, or actually increased. And communication bus 950 is allocated even more power. Therefore, a power increase for communication bus 950...is dynamically balanced with a cap of CPU cores 905...to alleviate...bottlenecks...increase performance, while not exceeding the current limit for IC 900.

Ex-1001, 17:23-41. Bose contains no similar indication of the importance of the workload of the interconnect 108. Indeed, the interconnect 108 in Bose is not

illustrated as being a communication pathway between cores 112, such that it would be a workflow bottleneck. *See* Ex-1005, FIG. 1.

Accordingly, Bose does not disclose or suggest “comparison between the bus workload and a bus workload threshold” as claimed.

IV. Ground 2 Fails

The Petition asserts in Ground 3 that Claims 3, 12, 14, and 15 are rendered obvious by Bose in view of Fukatsu. Petitioner does not assert that the combination with Fukatsu affects its assertions with respect to independent claim 1, as set forth in Ground 1. This Ground, which challenges only claims which depend from claim 1, therefore fails for the same reasons set forth above with respect to Ground 1.

V. Ground 3 Fails

The Petition asserts in Ground 3 that Claims 1, 3, 12, 14, and 15 are rendered obvious by Bose and Naffziger.

A. A POSITA Would Not Have Been Motivated to Combine Bose with Naffziger

Petitioner asserts that a “POSITA would have been motivated to combine the teachings of Bose and Naffziger...because both relate to similar methods of managing power consumption and performance in a processor.” Pet., pp.49-50. Petitioner has not explained how the two references relate to “similar methods.”

Petitioner asserts that both Bose and Naffziger describe “transferring or adjusting power allocation between cores or dies for efficiently managing or

controlling the processor's power consumption." Pet., p.50. However, this alleged "similarity" is acknowledged in both Bose and Naffziger to be one of the quintessential design considerations for all modern microprocessors. Ex-1005, 1:24-26 ("chip-level power dissipation limits constitute a fundamental design constraint"); *see also* Ex-1007, 1:10-12 ("The power consumption of modern integrated circuits has become an increasing design issues with each generation of semiconductor chips"). Allocating power between cores or dies for efficient management or control of a processor's power consumption is too remote a connection to demonstrate that these references "relate to similar method" as Petitioner asserts. Pet., 49-50.

Petitioner also asserts that "both Bose and Naffziger have similar components and structure." Pet., p.50. However, Petitioner goes on to merely identify components common to nearly all multi-core processors: a processor, cores, an interconnect or "crossbar switch," a shared cache, and a "PMCU" or "power management unit."³ Inclusion of such rudimentary micro-processor components is insufficient to demonstrate a motivation to combine; otherwise, every reference

³ Petitioner asserts these various components are "similar," so Patent Owner addresses them together here. However, Patent Owner takes no position here as to whether the components are indeed similar, e.g., whether Bose's "interconnect" is similar to Naffziger's "cross switch."

describing any processor would be combined, and the entire motivation to combine analysis would be mooted.

Additionally, even if Bose and Naffziger were directed to similar subject matter, Petitioner's alleged motivation is insufficient. Merely alleging that references are compatible and would have led to a "more efficient and better functioning processor" does not provide a motivation to combine them. *Johns Manville Corp. v. Knauf Insulation, Inc.*, IPR2018-00827, Paper 9 at 10 (Designated Informative) (citing *Securus Techs., Inc. v. Global Tel*Link Corp.*, 701 F. App'x 971, 977 (Fed. Cir. 2017) and *Personal Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 993 (Fed. Cir. 2017)).

Petitioner's additional support for a motivation to combine these references is grounded in impermissible hindsight. For example, Petitioner alleges that "Naffziger's use of thresholds...would provide for more efficient power allocation between components." Pet., pp.61-62. Petitioner states that a POSITA would have recognized "Naffziger's pre-defined P-states as a more effective way to predictably adjust power allocation." Pet., p.62. However, Petitioner is effectively arguing that if a POSITA were confronted with this combination they would recognize the benefit and be motivated to combine as a result. This is classic hindsight. Importantly, Petitioner does not identify any reason or justification *why* a POSITA would have

been motivated to modify Bose by combining it with Naffziger beyond the assertion that a POSITA would have recognized the alleged benefit of the combination.

Petitioner also fails to analyze the traditional considerations for a motivation to combine analysis. For example, Petitioner fails to allege that Naffziger provides a solution to a problem identified in Bose, that the proposed combination represents a substitution of any known elements from Naffziger for an element in Bose, or that the proposed combination uses a technique disclosed in Naffziger to improve Bose. Petitioner also cites zero market factors or other factors creating an incentive for a variation of Bose, and Petitioner provides no explanation as to how or why, in combining Bose with Naffziger, the other teachings of Bose would have been unchanged so as to not affect the features allegedly rendering obvious the claims from which the challenged claims of Ground 3 depend.

B. Bose in view of Naffziger does not render obvious Claim 1[d]

Petitioner does not assert that Naffziger discloses or suggests claim element 1[d], and instead relies exclusively on the Bose. Therefore, Ground 3 fails for at least the same reasons that Ground 1 does: Petitioner failed to demonstrate that Bose discloses or suggests Claim 1[d].

C. Bose in view of Naffziger does not render obvious Claim 1[g.ii]

Petitioner failed to demonstrate that Bose in view of Naffziger discloses or suggests “dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit,” as claimed.

Petitioner relies on Naffziger as allegedly disclosing this limitation, but fails to demonstrate that a POSITA would have incorporated the identified teachings with Bose. Pet., pp.54. For example, Petitioner asserts that a “POSITA would have understood” that Naffziger’s “allocating of power credits...and that the number of power credits...corresponds to a maximum amount of power a cooling system in a computer is able to dissipate” but does not explain why. Pet., p.79. There is no indication in the cited art, or articulated by Petitioner, that any aspects of Bose should or could be implemented according to known methods. For example, Bose does not even disclose dynamic tuning of power allocations between components, let alone that this could be done in combination with other methods.

Accordingly, Bose in view of Naffziger does not disclose or suggest “dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit.”

D. Bose in view of Naffziger does not render obvious Claim 1[g.iii]

Petitioner failed to demonstrate that Bose in view of Naffziger discloses or suggests “a comparison between the bus workload and a bus workload threshold,” as claimed.

Petitioner relies on Naffziger as allegedly disclosing this limitation, but fails to demonstrate that a POSITA would have incorporated the identified teachings with Bose. Pet., pp.54-55. For example, Petitioner asserts that a “POSITA would have readily applied” Naffziger’s “threshold comparison to Bose’s similar power management schedule” because “Bose’s tables for calculating power budgets and power modes includes performance statistics from all local PMCUs,” but does not explain why a POSITA would have done so. Pet., p.55. There is no indication in the cited art, or articulated by Petitioner, that any aspects of Bose should or could be implemented according to known methods. For example, Bose does not even disclose a “bus workload” or a “bus workload threshold” as claimed, let alone that this could be done in combination with other methods.

Accordingly, Bose in view of Naffziger does not disclose or suggest “a comparison between the bus workload and a bus workload threshold,” as claimed.

VI. Ground 4 Fails

The Petition asserts in Ground 4 that 1, 3, 14, and 15 are rendered obvious by White in view of Simeral.

A. A POSITA Would Not Have Been Motivated to Combine White with Simeral

Petitioner asserts that a “POSITA would have been motivated to combine the teachings of White and Simeral ... because both relate to similar methods of managing power in circuits by adjusting power consumption of circuit components.” Pet., p.62. Petitioner has not explained how the two references relate to “similar methods”: Petitioner asserts that White describes adjusting the voltage and frequency of logic cores and a bridge logic, whereas Petitioner asserts that Simeral describes adjusting the “the operational characteristics of the bus in comparison to a threshold activity level.” Pet., p.62. Nothing in either description overlaps as would make these references “relate to similar methods” as Petitioner asserts. Pet., p.62. For example, Petitioner does not assert that both references adjust the voltage or frequency of the bridge logic, or that both adjust operational characteristics in comparison to a threshold; the two references simply describe distinct approaches to powering computing hardware. Petitioner also asserts that “both White and Simeral teach similar components and structures.” Pet., p.63. However, Petitioner goes on to merely identify components common to virtually every multi-core processor: a processor or integrated circuit, cores or “various devices,” and “management logic”

or “controllers.”⁴ Inclusion of such generic processor components does not provide a motivation to combine; otherwise, every reference describing any processor would be combined, and the entire motivation to combine analysis would be mooted. Additionally, even if White and Simeral were directed to similar subject matter, Petitioner’s alleged motivation is insufficient: merely because references could be compatible or relate to the same subject matter does not provide a motivation to combine them. *Johns Manville Corp. v. Knauf Insulation, Inc.*, IPR2018-00827, Paper 9 at 10 (Designated Informative) (citing *Securus Techs., Inc. v. Global Tel*Link Corp.*, 701 F. App’x 971, 977 (Fed. Cir. 2017) and *Personal Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 993 (Fed. Cir. 2017)).

Petitioner also alleges that “Simeral’s teaching would have improved White’s ability to quickly respond to power-related events, such as spikes in bus activity, thus increasing the number and duration of power-saving opportunities.” Pet., p.62; *see also* Pet., p.64 (“would have improved White’s power distribution efficiency”). However, Petitioner’s assertions amount to alleging that if confronted the combination, a POSITA would have recognized a benefit, and thus are merely rooted in impermissible hindsight. For example, Petitioner does not identify any reason or

⁴ Petitioner asserts these various components are “similar,” so Patent Owner addresses them together here. However, Patent Owner takes no position here as to whether the components are indeed similar, e.g., whether White’s “core[s]” are similar to Simeral’s “Various Devices.”

justification why a POSITA would have been motivated to modify White by combining it with Simeral, just that doing so could be beneficial. However, that is the very benefit identified by the '494 Patent from including the claimed bus workload monitor: “communication bus 950 is associated with a workload monitor 951... workload monitor 951 determines activity of communication bus... Once the workloads are determined, then power is balanced among the selected devices (e.g. cores 910, cores 905, and bus 960) based on those measured workloads.” Ex-1001, 16:18-59.

Petitioner’s alleged motivations to combine are merely rooted in impermissible hindsight. Petitioner nowhere analyzes the traditional considerations of a motivation to combine analysis. For example, Petitioner does not allege that Simeral provides a solution to a problem identified in White. Nor does Petitioner allege the proposed combination represents a substitution of any known elements from Simeral for an element in White. Nor does Petitioner allege the proposed combination is the use of a technique disclosed in Simeral to improve White. Nor does Petitioner identify any market or other factors creating incentive for a variation of White.

B. White in view of Simeral does not render obvious Claim 1[d]

Petitioner failed to demonstrate that White in view of Simeral discloses or suggests a “communication bus to connect the first core and the cache memory” as claimed.

For example, Petitioner does not assert that White discloses a communication bus to connect the first core and the cache memory, instead alleging (1) that “it would have been obvious to a POSITA to connect the cores to logic cache 1128,” and (2) that this would have been done with “the common bridge logic.” Pet., p.71. However, White nowhere discloses that the cores are connected to the logic cache with the specific “common bridge logic 110” that Petitioner aligns with the claimed “communication bus.” To the contrary, FIG. 8 illustrates the connection between the L2 cache and the processor core 1100 as separate from the “CPU Bus” 1124 connection between the processor core 1100 and the bridge 1102. White is thus distinguishable from the claimed system, in which a specifically claimed “communication bus” is required “to connect the first core and the cache memory,” have its workload monitored by the bus workload monitor, and have its power allocation dynamically tuned.

C. White in view of Simeral does not render obvious Claim 1[e]

Petitioner failed to demonstrate that White in view of Simeral discloses or suggests “a core workload monitor configured to determine a core workload for the first core” as claimed.

For example, Petitioner does not assert that White discloses any component that monitors the **workload** of the core, instead alleging that the “power management logic 200 [is] configured to monitor and adjust the operating frequency and/or voltage.” Pet., p.72. However, as the '494 Patent's descriptions make clear, that workload is not the same as “performance metric[s]” such as frequency and/or voltage:

if a power limit is reached, then GPU 605's performance metric, such as frequency, may be reduced, since it doesn't need full performance to complete its workload that was below an activity threshold.

Ex-1001, 12:62-65.

workload module 515 is to determine a workload of core 505, core 510, or both. Determining a workload includes any known method for determining activity of hardware, a workload of software, a number of instructions executed or to execute, or any other metric for determining use of a processing element. As one example, a workload for core 505 includes an amount time

core 505 is active during a period of time. Activity may be determined as when core 505 is in an active performance state versus a sleep or low power state or a number of cycles executing or performing operations.

Ex-1001, 9:47-57.

For example, GPU cores 510 workload information is determined, as discussed above. And an executing application provides a software hint of whether a future workload (i.e. a portion of the application) is to include a larger or smaller graphics workload in comparison to a previous portion of the application that resulted in the measured workload.

Ex-1001, 11:27-32.

D. White in view of Simeral does not render obvious Claim 1[f]

Petitioner failed to demonstrate that White in view of Simeral discloses or suggests “a bus workload monitor configured to determine a bus workload for the communication bus” as claimed.

For example, Petitioner does not assert that White discloses any component that monitors the **workload** of the bus, instead alleging that the “power management logic 200 [is] configured to monitor and adjust the operating frequency and/or voltage.” Pet., p.73. However, as discussed above with respect to Claim 1[e], the ’494 Patent’s description makes clear that a workload is not the same as

“performance metric[s]” such as frequency and/or voltage. Ex-1001, 12:62-65; 9:47-57; 11:27-32.

Petitioner also asserts that “White teaches recognizing whether the common bridge logic is “a bottleneck to overall performance,” as alleged support for the bus workload monitor. Pet., p.74 (citing Ex-1008, 9:39-53). However, critically, White does not teach recognizing a “bottleneck” at all—it merely describes it that way to the reader; nothing indicates the system recognizes it—let alone by monitoring a “workload” as claimed. See Ex-1008, 9:39-53 (describing evaluating power, *i.e.*, voltage, consumption).

Petitioner asserts that “Simeral also teaches a bus workload monitor that determines workload for the communication bus,” relying on disclosures allegedly related to Simeral’s “bus 130.” Pet., p.74. Even if Simeral disclosed monitoring workload of its bus 130, Petitioner has not demonstrated that Simeral’s “bus 130” serves as the claimed “communication bus”—for example that it is “to connect the first core and the cache memory” as claimed. Thus, Petitioner failed to demonstrate that Simeral discloses monitoring the workload of a bus to connect a first core and a cache memory as required by claim 1. Petitioner instead relies on the alleged integration of Simeral’s teachings into White’s system, but does not explain why a POSITA would have been motivated to combine them in the manner claimed. For example, Petitioner asserts that the “combined White-Simeral teachings would have

been implemented by applying Simeral's technique of measuring the bus's activity, corresponding to the bus's power consumption, in White's power management logic 200," as opposed to, for example, the opposite approach.

E. White in view of Simeral does not render obvious Claim 1[g.ii]

Petitioner failed to demonstrate that White in view of Simeral discloses or suggests "dynamically tune power allocation between the first core and the communication bus based on a power limit for the integrated circuit" as claimed.

For example, Petitioner asserts that White discloses that power management logic can control the voltage and/or frequency of the cores or the bridge logic, but fails to demonstrate that White discloses tuning power allocation "between" them as claimed. For example, Petitioner merely asserts that "temperature control logic 210 may be configured to reduce the operating frequency for one or more of the logic cores" and that "such adjustments may apply to the common bridge logic because adjusting voltages and frequencies used by common bridge logic may be included as part of changing the operating characteristics (i.e., voltage and frequency) of the logic cores." Pet., p.78. However, adjusting the cores and bridge logic downward together is not the same as "dynamically tun[ing] power allocation **between** the first core and the communication bus" as claimed. For example, tuning allocation between components involves adjusting the distribution amongst them, not merely adjusting them both the same way in the same amount. As the '494 Patent explains:

As one example, power module **520** is to increase frequency for core **505** and/or decrease frequency for core **510** in response to the dynamically determined workload for core **505** being greater the workload for core **510**. Here, if one core is demonstrating a greater need of performance due to a higher workload, then its frequency is increased. And when an integrated circuit including cores **505** and **510** are up against a power limit, then the core demonstrating less need for performance requirements is decreased in frequency to counterbalance/compensate for the increase in frequency to the other core. In another example, power module **520** is to increase frequency for core **505** and/ or decrease frequency for core **510** in response to determining core **505** is a performance bottleneck. In other words, core **505** is limiting overall performance, because core **510** is relying on output from core **505** to continue further execution. But core **505** is operating to slowly. Then, core **505's** frequency is increased to ensure overall performance is accelerated. Yet again, if a power limit is encountered, then core **510's** frequency may be similarly decreased to ensure the power limit is enforced.

Ex-1001, 10:25-44.

F. White in view of Simeral does not render obvious Claim 1[g.iii]

Petitioner failed to demonstrate that White in view of Simeral discloses or suggests “dynamically tune power allocation between the first core and the communication bus based on ... a comparison between the bus workload and a bus workload threshold” as claimed.

For example, Petitioner asserts “White teaches dynamically tuning power allocation between the cores and common bridge to avoid bottlenecking at the bridge,” but fails to demonstrate that White disclosed a system that could monitor and account for “bottlenecking.” Pet., p.79. That is, White discloses the impact its system allegedly has—reduced bottlenecking—but does not disclose that its system **monitors** for bottlenecking of performance specifically. Instead, White’s system merely opts to operate the common bridge at higher voltage and/or frequency than the cores, regardless of whether the common bridge is actually acting as a bottleneck.

Petitioner also relies on Simeral as allegedly disclosing this limitation, but fails to demonstrate that a POSITA would have incorporated the identified teachings with White’s. Pet., pp.79-80. For example, Petitioner asserts that a “POSITA implementing White’s workload-related teachings would have turned to additional references for further implementation details of dynamic tuning of power allocations, like Simeral,” but does not explain why. Pet., p.79. For example, there is no indication in the cited art, or articulated by Petitioner, that any aspects of White

should or could be implemented according to known methods. For example, White does not even disclose dynamic tuning of power allocations between components, as discussed above, let alone that this could be done in combination with other methods.

CONCLUSION

For at least the foregoing reasons, the Petition does not meet its burden of demonstrating a reasonable likelihood that Petitioner would prevail with respect to any of the challenged claims. Accordingly, Patent Owner respectfully requests that the Board deny institution.

Dated: August 14, 2023

Respectfully submitted,

/Peter F. Snell/

Peter F. Snell (Reg. No. 52,235)
MINTZ, LEVIN, COHN, FERRIS,
GLOVSKY AND POPEO, P.C.

919 Third Avenue

New York, New York 10022

Telephone: 212-935-3000

Facsimile: 212-983-3115

E-mails:

PFSnell@mintz.com

CERTIFICATE OF WORD COUNT

The undersigned certifies pursuant to 37 C.F.R. § 42.24(d) that the foregoing Patent Owner Preliminary Response, excluding any table of contents, table of authorities, certificates of service or word count, or appendix of exhibits or claim listing, contains 7,044 words according to the word-processing program used to prepare this paper (Microsoft Word). Patent Owner certifies that this Patent Owner Preliminary Response does not exceed the applicable type-volume limit of 37 C.F.R. § 42.24(a).

Dated: August 14, 2023

/Peter F. Snell/

Peter F. Snell (Reg. No. 52,235)

CERTIFICATE OF SERVICE

I certify that a copy of Patent Owner's Preliminary Response is being served by electronic mail on the following counsel of record:

Lead Counsel

William M. Fink, (Reg. No. 72,332)
O'MELVENY & MYERS LLP.
1625 Eye Street, NW
Washington, DC 20006
Telephone: (202) 383-5300
Facsimile: (202) 383-5414
Email: tfink@omm.com

Backup Counsel

Daniel Leventhal (Reg. No. 59,576)
Richard Zembek (Reg. No. 43,306)
Darren Smith (Reg No. 64,261)
NORTON ROSE FULBRIGHT
1301 McKinney, Suite 5100
Houston, TX 77010
Telephone: (713) 651-5151
Facsimile: (713) 651-5246
Emails:
daniel.leventhal@nortonrosefulbright.com
richard.zembek@nortonrosefulbright.com
Darren.smith@nortonrosefulbright.com

Eagle H. Robinson (Reg No. 61,361)
NORTON ROSE FULBRIGHT
98 San Jacinto Blvd., Suite 1100
Austin, TX 78701
Telephone: (512) 474-5201
Facsimile: (512) 536-4598
Email:
eagle.robinson@nortonrosefulbright.com

*Counsel for Petitioner Qualcomm
Incorporated*

Dated: August 14, 2023

/Peter F. Snell/

Peter F. Snell (Reg. No. 52,235)