UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION
Petitioner

v.

PACT XPP SCHWEIZ AG
Patent Owner

Case IPR2020-00541
U.S. Patent No. 9,075,605

DECLARATION OF WILLIAM HENRY MANGIONE-SMITH, PH.D.

TABLE OF CONTENTS

Page

I, William Henry Mangione-Smith, declare as follows:

I. INTRODUCTION

1. I have been engaged by Quinn Emanuel Urquhart & Sullivan, LLP, counsel for PACT XPP Schweiz AG. ("Patent Owner"). I understand the Patent Trial & Appeal Board (PTAB) of the U.S. Patent & Trademark Office is evaluating a Petition for *Inter Partes* Review ("IPR") of the patentability of claims 1-6 of U.S. Patent No. 9,075,605 ("the '605 patent") filed by Intel Corporation ("Petitioner"). I understand that Patent Owner is filing a Response to that Petition. I have been asked to provide certain opinions relating to that Response. The following discussion provides my opinions. I am being compensated at my standard rate of $650 per hour. No part of my compensation is dependent on my opinions or the outcome of this proceeding, and I have no financial or beneficial interest in any of the parties to this request.

2. This declaration is based on the information that was available to me as of this declaration. I may revise, supplement, or expand my opinions based on further review and analysis of information and opinions provided to me subsequent to the filing of this declaration, including information and opinions submitted by the Petitioner.

3. I have considered information from various sources in forming my

opinions. I have reviewed and considered each of the materials listed in the Petition, as well as any other references mentioned in this declaration, in forming my opinions.

II. BACKGROUND & QUALIFICATIONS

4. My technical background covers most aspects of computer system design, including low level circuitry, computer architecture, computer networking, graphics, application software, client-server application, Web technology, and system software (e.g., operating systems and compilers). I am a member of the Institute of Electrical and Electronics Engineers and the Association for Computing Machinery, which are the two most significant professional organizations in my profession. I have been employed as a design engineer, research engineer, professor, and technical expert. Over my professional career I have been an active inventor with 121 issued U.S. patents, 209 published and pending U.S. patent applications, and many unpublished U.S. patent applications.

5. From 1984 until 1991 I attended the University of Michigan in Ann Arbor, Michigan. I was awarded the degrees of Bachelor of Science and Engineering, Master of Science and Engineering, and Doctor of Philosophy. My doctoral research focused on high-performance parallel computing systems including computer architecture, applications and operating system software, and compiler technology. One of my responsibilities during my graduate studies included teaching senior

undergraduate students who were about to enter the profession.

6. After graduating from the University of Michigan I was employed by Motorola in Schaumburg, Illinois. While at Motorola, I was part of a team designing and manufacturing the first commercial battery-powered product capable of delivering Internet email over a wireless (i.e., radio frequency) link and one of the first personal digital assistants. I also served as the lead architect on the second-generation of this device. Part of my responsibilities at Motorola involved the specification, design, and testing of system control Application-Specific Integrated Circuits ("ASICs"). I conducted the initial research and advanced design that resulted in the Motorola M*Core embedded microprocessor. M*Core was designed to provide the performance of desktop microprocessors with the low power of contemporaneous embedded processors. The M*Core received widespread use in a number of communications products including various telephonic handsets, advanced pagers, and embedded infrastructure. I also developed tools to evaluate the performance of a distributed parallel computing system. While at Motorola I was the sole inventor on one U.S. patent.

7. From 1995 until 2005 I was employed by the University of California at Los Angeles ("UCLA") as a professor of Electrical Engineering. I was the director of the laboratory for Compiler and Architecture Research in Embedded Systems

("CARES") and served as the field chair for Embedded Computing Systems. The CARES research team focused on research, engineering and design challenges in the context of battery-powered and multi-media mobile computing devices. One of the key developments of my lab was the Mediabench software tool, which is widely used to design and evaluate multi-media embedded devices. My primary responsibility, in addition to classroom teaching, involved directing the research and training of graduate students. I was a tenured member of the faculty and had responsibilities for teaching as well as scholarly research. While at UCLA I was a named inventor on three U.S. patent applications, one of which issued as a patent. My colleagues at UCLA were some of the leading scientists and engineers in the world with a long list of innovations from computer network security devices to the nicotine patch. The graduate student researchers in my laboratory came from a diverse set of backgrounds, all with undergraduate degrees in computer engineering, electrical engineering or computer science, many with multiple years of experience working as professional engineers in areas such as software development, computer system design, and ASIC circuit design.

8. From 2005 until 2009, I was employed at Intellectual Ventures in Bellevue, Washington. My responsibilities at Intellectual Ventures included business development, technology assessment, market forecasting, university

outreach, collaborative inventing, intellectual property licensing support, and intellectual property asset pricing. My colleagues and co-inventors at Intellectual Ventures included the former lead intellectual property strategist at Intel, Intel's lead IP counsel, Microsoft's chief software architect, the founder of Microsoft research, the designer of the Mach operating system, the architect of the U.S. Defense Department's Strategic Defense Initiative, the founder of Thinking Machines (a seminal parallel processing computer system), and Bill Gates. I had responsibility for hiring and managing over 15 staff members including multiple Ph.Ds. with degrees in electrical engineering and decades of experience in product design and engineering.

9. For further details regarding my employment and academic history, please refer to my curriculum vitae (Ex. 2025).

III. RELEVANT LAW

10. I am not an attorney and I have no formal legal training. For the purposes of this Declaration, I have been informed about certain aspects of the law that are relevant to my opinions. My understanding of these aspects is summarized below.

A. Claim Construction

11. I have been informed that the first step of a patentability analysis is to

construe the claims. I understand that, when construing claims in an IPR, the PTAB will apply the same standard applied in civil actions in federal courts. I have been informed that claim construction begins with the language of the claims and that the words of a claim are generally given their ordinary and customary meaning, which is the meaning that the term would have to a person of ordinary skill in the art ("POSITA") at the time of the invention, i.e., as of the effective filing date of the patent application.

12. I also understand that the specification is the single best guide to the meaning of a disputed term. The specification acts as a dictionary when it expressly defines terms used in the claims or when it defines terms by implication. I understand this is called "lexicography." The specification can also disavow or disclaim the otherwise full scope of a claim term.

13. I have been informed that the prosecution history is another source of intrinsic evidence that can inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it would otherwise be.

14. Finally, I am aware that extrinsic evidence, such as expert testimony and dictionaries, may be useful in educating the PTAB regarding the field of the invention or helping determine what a POSITA would have understood claim terms to mean.

B. Perspective of a person of ordinary skill in the art at the time of the claimed invention

15. I have been informed that claim construction and patentability are generally analyzed from the perspective of a hypothetical POSITA at the time of the invention.

16. I understand that Professor Sechen opined that a POSITA "would have a master's degree in electrical or computer engineering, or equivalent, and three years of experience in processor system design." Ex. 1001 (Sechen Decl.) ¶85.

17. I will apply Professor Sechen's definition for this declaration. Based on my education, experience, and research, I believe I am qualified to form opinions from the perspective of that a POSITA, and I have done so in this declaration.

C. Obviousness

18. I have been informed that 35 U.S.C. §103 governs the determination of obviousness. According to §103:

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

19. I have been further informed that the first three factors to be considered in an obviousness inquiry are: (1) the scope and content of the prior art, (2) the differences between the prior art and the claim at issue, and (3) the level of ordinary skill in the pertinent art.

20. I also have been informed that a combination of claimed elements in a patent claim is obvious when all of the claimed elements were known in the prior art and there was a reason for a POSITA to combine or modify the prior art to obtain the elements as claimed with no change in their respective functions, provided a POSITA would have had a reasonable expectation of successful modification or combination. Conversely, I am informed that a patent claim would not be obvious if a proposed modification would render the prior art being applied unsuitable for its

intended purpose or if the proposed modification would change the principle of operation of the prior art.

21. I have been informed that, for a dependent claim to be obvious, the combination of claimed elements in the dependent claim must be obvious in view of the prior art. Also, the combination of all of the limitations recited in the base claims from which the dependent claim depends, either directly or indirectly, must be obvious in view of the prior art.

IV. COPY OF CHALLENGED INDEPENDENT CLAIMS

22. Claim 1 recites:

> **[a]** A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing

data, the method comprising:

[b] the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

[c] the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

[d] the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

23. Claim 2 recites:

The method of claim 1 wherein: said clock frequency is determined in part by processor voltage.

24. Claim 3 recites:

The method of claim 1 wherein: said clock frequency is determined in part by available power

25. Claim 4 recites:

The method of claim 1 wherein: said clock frequency is determined in part by a user setting.

26. Claim 5 recites:

The method of claim 1 wherein: said clock frequency is determined in part by available power and a user setting.

27. Claim 6 recites:

The method of claim 1 wherein: said number of pending operations is determined at by the fill level of one or more buffers.

V. OVERVIEW OF THE '605 PATENT

28. The '605 patent is directed to systems and methods for separately setting the clock frequencies of different hardware modules within an integrated circuit. The '605 patent also discloses specific ways to assign clock frequencies to different modules in order to improve power management and increase the operating efficiency. Specifically, the '605 patent discloses changing the clock frequencies of portions of the multi-core system in a particular way to take advantage of the processing power of certain cores and in the meantime increase power efficiency. For example, clock frequencies of groups of processors such as those shown in Fig. 7 may be changed to achieve improved power efficiency:

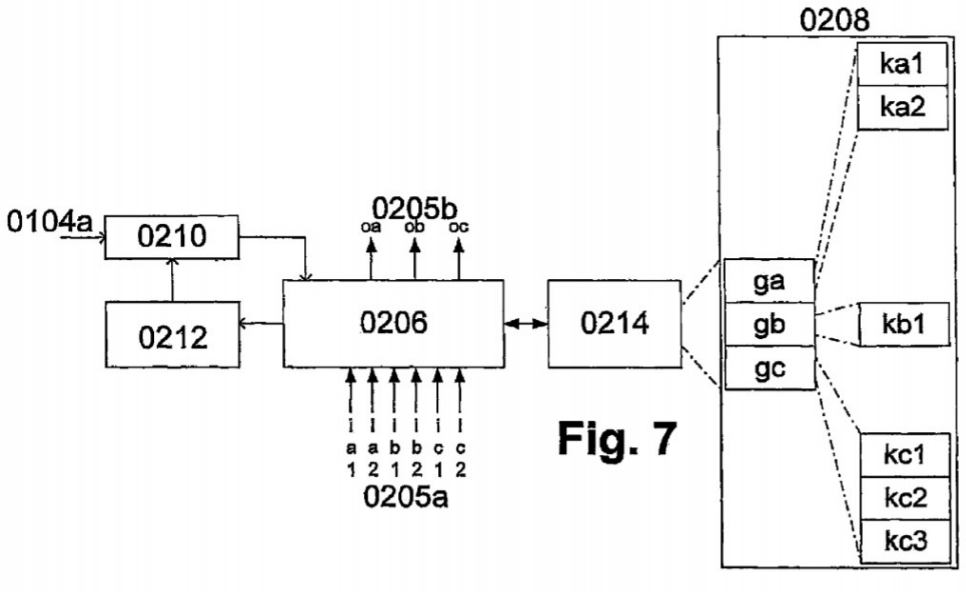

Fig. 7

VI. TECHNOLOGY BACKGROUND

A. Nicol

29. I have been asked to review U.S. Patent No. 6,141,762 which was granted to Christopher J. Nicol and Kanwar Jit Singh. I will refer to this work as Nicol. Nicol is directed to methods for reducing power consumption of a multi-processor chip system. Ex. 1008, Abstract; 1:5-7. Specifically, Nicol describes a system that controls the voltage applied to an array of processors and the processing load assigned to each processor. *Id*. at 2:18-22. Nicol describes two different embodiments of particular note, one shown in Figure 2 and another shown in Figure 4. For the Figure 2 embodiment, all of processors (or "processing elements") operate at the same clock frequency. For this embodiment, a controller "allocates tasks to

the individual processors to equalize processing load among the chips." *Id*. at 2:24. This function is performed using a scheme such as the one shown in Figure 2, which depicts "a block diagram of a multi-processor chip with supply voltage control" (*Id*. at 2:37-38) and is reproduced below:

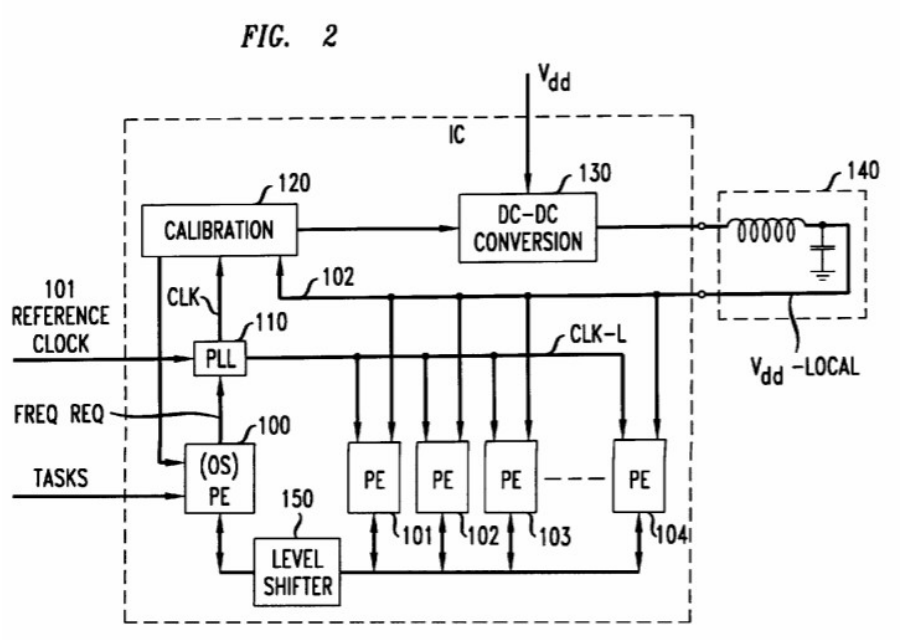

30. PE 100, upon which the operating system ("OS") is executed, "ascertain[s] the required completion time, divide[s] the collection of tasks as evenly as possible (in terms of needed processing time), consider[s] the PE with the tasks that require the most time to carry out, and adjust[s] the clock frequency to [e]nsure that the most heavily loaded PE carries out its assigned tasks within the required completion time." *Id*. at 3:67–4:5. After determining a clock frequency to ensure all the tasks will be completed in time, PE 100 adjusts the clock frequency by increasing

the voltage in a stepped manner as shown in Fig. 3, reproduced below. *Id*. at 5:29-32. Upon completion of the first collection of tasks, a new collection of tasks is created and divided among the processors. The frequency is then gradually increased or decreased, in the same stepped manner, from the frequency set for the first collection of tasks, to a new frequency that will ensure the most heavily loaded PE for the second set of tasks completes its assigned task in time. *Id.* at 4:54–5:2.

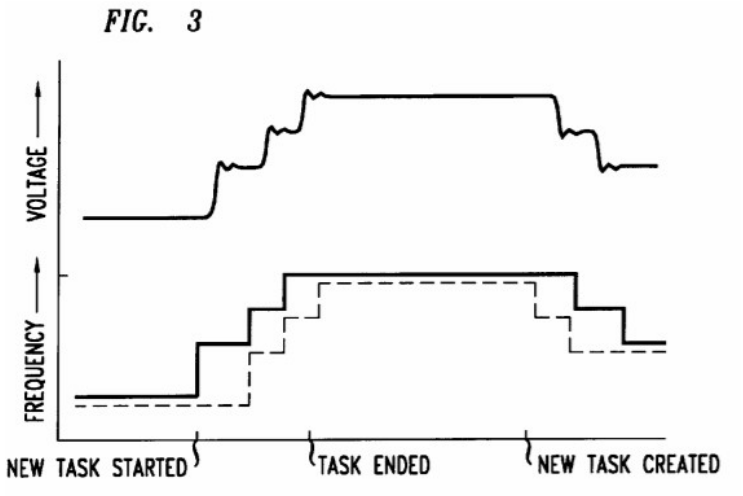

31. Nicol also discusses a separate embodiment where "each PE [has] its own supply voltage." *Id*. at 5:55-56. Nicol describes the need for this embodiment: "In some applications, however, one may encounter tasks that cannot be partitioned into concurrent evenly-loaded threads" *Id*. at 5:60-61. This embodiment, depicted in Figure 4, has "[a] separate power supply for each PE in a chip" and an "operating

system to independently program the lowest operating frequency and corresponding lowest supply voltage for each PE." *Id*. at 5:66-6:2. Because the frequency for each PE can be set independently, the lowest operating frequency for each PE can be determined based on the load on that PE. Because each PE's frequency and voltage is set independently, "[e]ach PE in Fig. 4 needs an independent controller that performs the functions of PE 100 (except it does not divide tasks among PEs)." In Fig. 4, "all of the controllers are embodied in a single controller 200, which may be just another processing element of the integrated circuit that contains the other processing elements." *Id*. at 6:6-9. In the Figure 4 embodiment, because "the frequencies at which the individual PEs operate differ from one another and from other elements within the system where the multiprocessor chip is employed, there is an issue of synchronization that must be addressed." *Id*. at 6:14-18. "To effect such synchronization, each PE within the Fig. 4 arrangement is connection to an arrangement comprising elements 150 and 160. Level converter 150 converts the variable voltage swings of the PEs to a fixed level swing, and network 160 resolves the issue of different clock domains." *Id*. at 6:29-33. Nicol's Figure 2 embodiment does not include a network 160 because it does not have synchronization issues, since all of the PEs in Figure 2 embodiment operate at the same frequency.

B. Kling

32. I have been asked to review U.S. Patent No. 6,367,023 which was granted to Ralph Kling and Edward Grochowski. I will refer to this work as Kling. Kling is directed to "limiting the power consumed in a computer system by throttling the power consumed by an integrated circuit in response to a high power condition." Ex. 1010, 1:8-10. Specifically, Kling describes a method shown in Figure 4, reproduced below, whereby the power consumed by a processor is continually measured. *Id*. at 7:41-47. When the power consumed exceeds a threshold, a throttle signal is sent to one or more integrated circuit of a computer system to reduce power consumption. *Id*. at 2:43-47. The threshold is set at a value that is less than a power level that would trip a fuse or circuit breaker of a power outlet from which the computer system draws power.

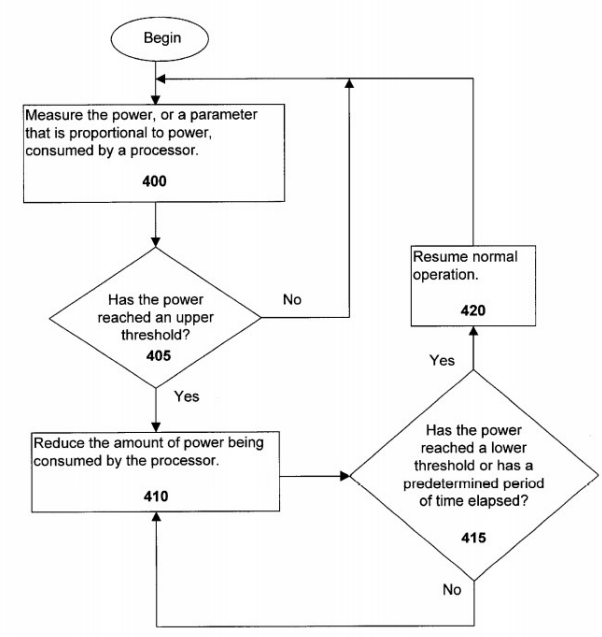

Figure 4

33. As shown in Figure 2A, power is decreased when an upper limit is reached. In addition, there is shown a lower limit. The lower limit acts to cause the integrated circuit to resume normal operations. *Id*. at 5:45-51. As a result, when activated, Kling's ICs will oscillate between the upper limit and lower limit and maintain a high processing power while avoiding tripping the protection circuit of a wall outlet. *Id*. at 5:31-37.

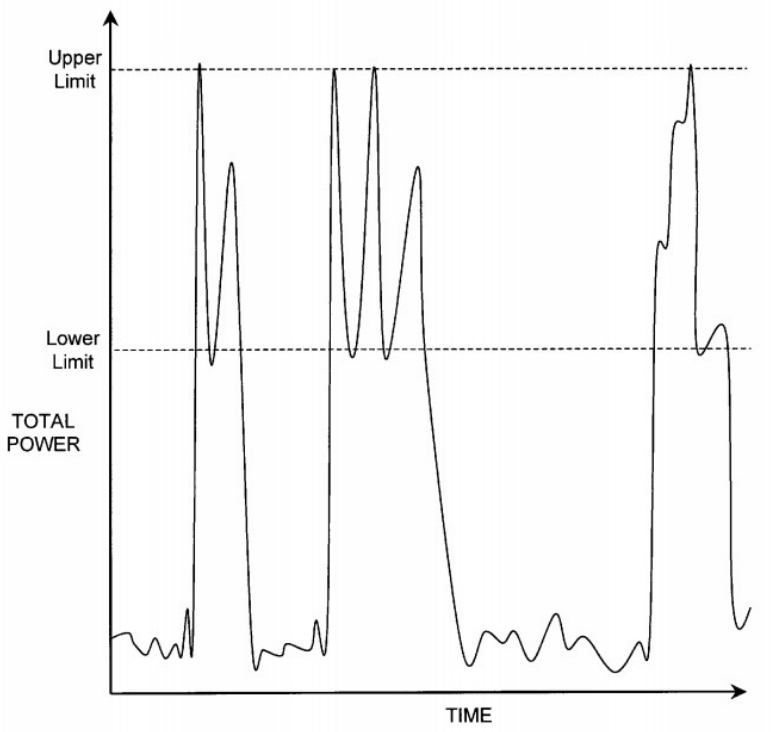

Figure 2A

C. DeHon

34. I have been asked to review U.S. Patent No. 9,052,773, which was granted to Andre DeHon, Michael Bolotski, and Thomas F. Knight, Jr. I will refer this patent as DeHon. DeHon is directed to a processor that includes a "gate array . . . configurable to adapt to the particular time-changing demands," which allows such processors to "operate at speeds only achievable if the device were specifically designed for the applications." Ex. 1010 at 2:14-18.

D. Bhatia

35. I have been asked to review U.S. Patent No. 6,535,798, which issued to

Rakesh Bhatia, Dennis Reinhardt, and Barnes Cooper. I will refer to this Patent as Bhatia. Bhatia is directed to a thermal management method that monitors a temperature of a computer and changes a clock frequency of a processor according to the monitored temperature. Ex. 1009, Abstract. Specifically, Bhatia describes a process, shown in reproduced Figure 3 below, by which a computer's processor is transitioned between a high performance state and a low performance state based on a determination of whether a temperature sample ("T_n") of a thermal zone within computer exceeds a temperature threshold ("T_t"). *Id*. at 2:38-44; 3:1-3:10; 3:32-46; 7:20-45. The processor clock frequency and voltage are higher in the high performance state than the processor clock frequency and voltage in the low performance state. *Id*. at 2:36-44.

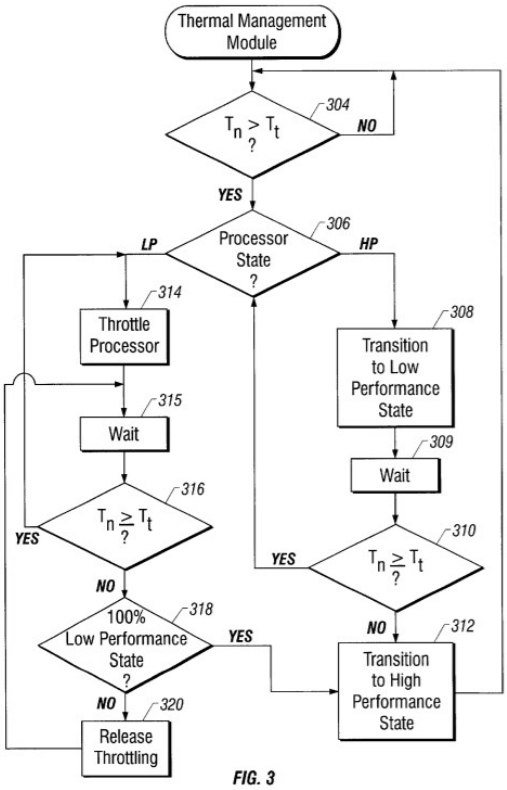
Thermal Management Module
304
$T_n \geq ? T_t$
NO
YES
306
Processor State ?
LP
HP
314
Throttle Processor
308
Transition to Low Performance State
315
Wait
309
Wait
316
$T_n \geq ? T_t$
YES
310
$T_n \geq ? T_t$
YES
NO
318
100% Low Performance State ?
YES
NO
312
Transition to High Performance State
NO
320
Release Throttling

FIG. 3

36. Bhatia's system aims to improve on existing thermal management systems for computers by maintaining system performance at a high rate through active management of clock frequencies according to temperature measurements. *Id*. at 1:6-49. As an aside to this main purpose, Bhatia also briefly suggests that transitions between high and low performance states may also be effectuated in response to an "environmental change" such as "docking or undocking of the computer system 10 and the coupling or discoupling of the external source port 58 to an external power source." *Id*. at 12:5-12:19. These events are unrelated to temperature, but Bhatia suggests that a user may nonetheless configure a processor to trigger a switch between performance states. *Id*. Bhatia does not specify whether or how these other triggers would be incorporated into Bhatia's thermal management system.

VII. LACK OF MOTIVATION TO COMBINE THE ASSERTED PRIOR ART

Lack of Motivation to Combine the Different Embodiments in Nicol

37. I understand the chip depicted in Figure 2 of Nicol and the chip in Figure 4 of Nicol to be distinct embodiments with different modes of operation. For the Figure 2 embodiment, Nicol describes a circuit which has only one voltage

supply because all of the PEs are designed to operate at the same clock frequency. Nicol describes how to balance the load "as evenly as possible" because the purpose is to find the allocation of tasks that will allow operation of the entire chip at the lowest clock frequency. In this case, it is my opinion that, it does not make sense to allocate more tasks to one PE as that would force running that PE, and hence the entire chip, at a higher frequency in order to finish within the required completion time. In contrast, for Figure 4 Nicol describes that "in some applications, however, one may encounter tasks that ***cannot be partitioned into concurrent evenly-loaded threads***, and therefore, some PE within the multiprocessor would require a higher operating frequency and a higher operating voltage." Ex. 1009 5:60-61; Ex. 1001, 33:4-16. This is the motivation Nicol provides in describing the Figure 4 embodiment, where each PE has its own voltage supply, so that each PE can be operated at its own lowest operating frequency. In this scenario, it is my opinion that there is no longer a need to allocate tasks as evenly as possible; indeed, Nicol states this explicitly. *See* Ex. 1009, 6:3-6 ("Each PE in Fig. 4 needs an independent controller that performs the functions of PE 100 (***except it does not divide tasks among PEs***)"). The frequency is not based on the most heavily loaded PE. Instead, each PE can be set independently, and a lowest operating frequency can be determined based on its own load. Because of this, it is my opinion that the Figure

4 embodiment cannot simply be combined with Figure 2 embodiment without modification. Certainly, the disclosure about the most heavily loaded PE that Petitioner relies on is not applicable for the Figure 4 embodiment. In my opinion, the features in Figures 2 and 4 are not interchangeable as Figure 2 and Figure 4 deal with different scenarios (when tasks can and cannot be evenly divided, and when frequencies for PEs cannot and can be set independently), and Petitioner has not described at all how they should be combined. Petitioner simply treats Figure 4 and Figure 2 as if the various features are interchangeable and a POSITA would be motivated to combine these different embodiments.

VIII. GROUND I FAILS

38. In my opinion, Nicol in view of Kling fails to disclose the features of independent claim 1 or the features of dependent claims 2-6.

A. PETITIONER HAS NOT SHOWN THAT NICOL IN VIEW OF KLING RENDERS OBVIOUS INDEPENDENT CLAIM 1

39. In my opinion, Nicol fails to disclose "setting a clock frequency, of at least a part of the multiprocessor system to a minimum," " increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor," "being adapted for sequentially processing data," "reducing the clock frequency of the at

least the part of the multiprocessor system in accordance with . . . an operating temperature threshold preventing over-temperature," and "reducing the clock frequency of the at least the part of the multiprocessor system in accordance with . . . a hysteresis characteristic." as recited in claim 1.

(a) **Petitioner has not Carried its Burden to Show That Nicol in view of Kling discloses "being adapted for sequentially processing data"**

40. I understand Petitioner argues that a POSITA would understand Nicol's inclusion of an asynchronous communication network ("ACN") to teach sequentially processing data based on Petitioner's three step analysis. Pet., 31. I understand Petitioner argues (a) ACN would not have been necessary if Nicol's PEs were not configured to pass results of their tasks to other processors, *id*. at 31-32; (b) systems designed to handle digital signal processing ("DSP") applications **typically** involve a pipeline process whereby PEs pass data from one to the next in a sequence, *id*. at 33; and (c) Nicol's system **could be** configured like a systolic architecture, *id*. at 34-35. Beyond Petitioner's admission that its analysis rests on mere possibilities (which do not support the contention that a POSITA would assume an ACN implies programmable processing sequences), each step of Petitioner's analysis is flawed.

(i) **Petitioner's Evidence Does Not Show that a POSITA Would Understand an ACN to Imply**

Passing of Data or Instructions Between Processors

41. As shown below, Nicol's Fig. 4 includes a level converter and asynchronous communication network connecting a controller and the PEs:

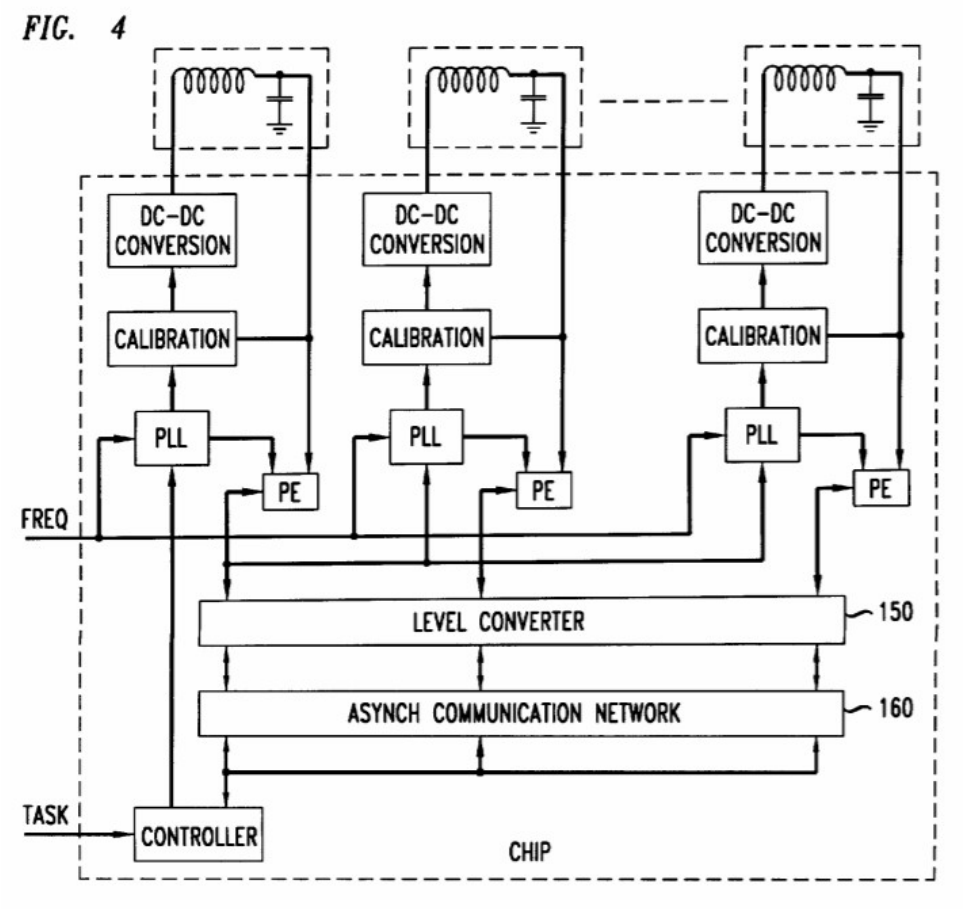

The level converter and ACN are used to effect synchronization required because the PEs operate at different voltages and frequencies. Ex. 1008 (Nicol) at 6:29-33. The level converter (also known as a level shifter or level translator) receives data signals from each of the PEs, which are supplied with different levels of voltage

from the various DC-DC converters, and converts the signals to signals with a common voltage level. Ex. 1009, 6:31-32 ("Level converter 150 converts the variable Voltage Swings of the PES to a fixed level Swing"). The level converter then transmits the modified data signals to the ACN which is configured to receive and coordinate data from the individual PEs which output data at different times. Ex. 1009, 6:33 ("[N]etwork 160 resolves the issue of different clock domains."). The PEs in Fig. 4 operate at different clock frequencies. Ex. 1009, 6:14-17. The ACN is needed to receive data from the PEs at different times, coordinate the data as part of the same set of tasks, and pass the coordinated data along to the controller. *Id*. The use of the ACN does not necessitate communication across the PEs, much less imply that sequential data processing is performed. *Id*. I understand that Professor Sechen argues to the contrary. Pet., 31-32. I disagree with Professor Sechen's conclusion and believe it is unfounded. I understand Petitioner concludes, quoting Nicol, that "the asynchronous communication network is provided, *solely*, in the situation where there is 'a need to communicate data between PEs (or with other system elements).'" *Id*. (emphasis original). Nicol, quoted in context, states only that "a synchronization schema must be implemented when there is a need to communicate data between PEs." Ex. 1008, 6:19–21. Nicol describes a situation when a synchronization schema is needed ("when there is a need to communicate data

between PEs"), not that the use of a synchronization schema necessarily implies that data is being communicated between PEs.

42. I understand that Petitioner nonetheless refers to the coprocessor configuration described in Ex. 1010, Hayes, John, Computer Architecture and Organization (third edition), McGraw-Hill International Edition: Computer Science Series (1998) ("Hayes Textbook") to argue that the ACN necessarily implies communication between PEs. Pet., 34. In my opinion, the comparison is unfounded for two reasons.

43. First, Nicol and Hayes Textbook describe two distinct systems. Nicol describes a multi-processor system where "each PE contains a central processing unit (CPU) and a local cache memory." Ex. 1008, 2:51-53. Hayes Textbook, on the other hand, describes a processor and floating point coprocessor configuration within a single CPU. Ex. 10 Fig. 3.14. Additionally, unlike the processors in Nicol, Hayes Textbook's coprocessor is a "specialized instruction execution unit" that does not include its own cache memory. Ex. 1010 at 159. The coprocessor is designed only to execute specific floating point operations which are sent directly by the microprocessor, which fetches and decodes instructions and transfers only the command portions to the coprocessor. *Id*. Finally, while the Nicol ACN exists specifically to address the problem of multiple processing elements with different

computing frequencies, that situation does not exist within the Hayes Textbook. Accordingly, in my opinion, Hayes Textbook's coprocessor configuration is not analogous to Nicol's multi-processor system, and a POSITA would not look to Hayes Textbook to understand Nicol's ACN.

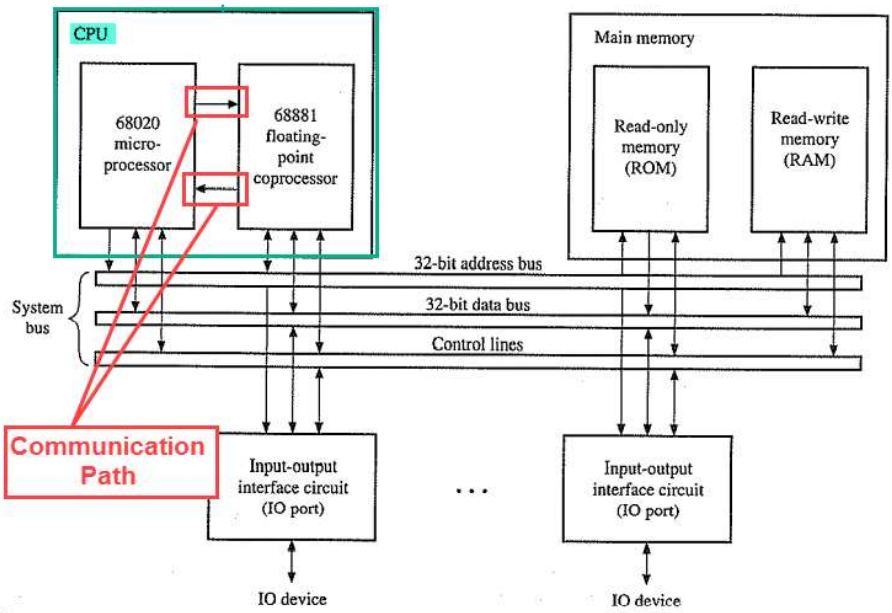

Figure 3.14

44. Second, even if a POSITA were to consider Hayes Textbook, it is my opinion that the disclosure Petitioner points to does not show that an ACN implies data communication between processors. I understand that Petitioner suggests that the microprocessor and coprocessor of Hayes Textbook's would communicate along the 32-bit data bus shown in Fig. 3.14. Pet., 32. However, Fig. 3.14 indicates that the microprocessor and coprocessor communicate directly within the CPU, and in my opinion, a POSITA would not expect processors within a CPU to communicate over

an external system bus due to the great decrease in efficiency. Direct communication between the processor and co-processor would be faster and less complex because the information passed between them would not have to be configured to be communicated across the bus. Because the microprocessor and co-processor communicate directly and are unlikely to communicate across the bus, it is my opinion that a POSITA would not assume that a bus connected to multiple processors (as in Nicol) implies it is used by the processors to share data between them.

(ii) Nicol's Digital Signal Processing does not Require a Pipeline Process with PEs Passing Data in a Sequence

45. I understand Petitioner to argue that because Nicol performs digital signal processing and digital signal processing is allegedly typically performed using a chain of processors, Nicol must be performed in the same manner as, for example, Exhibit 1018 ("Williams"). Pet., 27-28. First, in my opinion, it is inaccurate to say that digital signal processing is "typically" performed using a chain of processors – there are a multitude of architectures for performing digital signal processing. By far the most common signal processing tasks are finite impulse response (FIR) and infinite impulse response (IIR) filters, which are single algorithms running on individual processors. Next, in my opinion, Nicol does not suggest that it would process data according to the configuration shown in Williams. Williams is directed

to graphics processing. *See, e.g.,* Ex. 1018, Abstract. Graphics processing is not a form of digital signal processing, and the tasks are not analogous to a POSITA. Nicol provides no indication that it is designed to perform graphics processing. Williams's system includes a pixel operations unit, a rasterization unit, a per fragment operations unit, and a frame buffer. Ex. 1018, 10:56–11:19. Nicol does not mention any of these elements, the functions of any of these elements, or even the concept of graphics processing. Accordingly, it is my opinion that a POSITA would not read Williams to disclose a specific configuration that is designed to perform a specific type of digital signal processing that Nicol does not mention.

(iii) **A POSITA Would not Modify Nicol's Multiprocessor System to Operate as a Systolic Architecture.**

46. I understand that Petitioner argues that a POSITA would have known that the PEs in Nicol's multiprocessors system could be configured to operate as a systolic architecture. Pet., 34-35. I disagree. In my opinion, a POSITA would understand that Nicol would not and could not be operated as a systolic architecture.

47. Systolic architecture refers to a chain of processors hardwired in series using point-to-point communications links. They have a very narrow set of related applications and require complicated analysis and planning before data is transmitted to a systolic architecture. In contrast Nicol's PEs have no hardwired point-to-point

connections. It could not be used in Nicol's architecture which, as shown below in reproduced Fig. 4, includes processors connected in parallel. I understand that Professor Sechen further argues that in a systolic array, "[d]ata words flow *synchronously*." Ex. 1001, ¶143. However, I am unable to ascertain any explanation from either Petitioner or Professor Sechen as to how an asynchronous system, such as Nicol's Fig. 4, could be combined with a synchronous systolic architecture (much less why a POSITA would be motivated to combine them). It is also unclear whether and how such a combination would work in an asynchronous manner, which is how Nicol describes its ACN. To reiterate, in my opinion, Nicol could not integrate a systolic array and, if it did somehow integrate a systolic array, a POSITA would not understand that system to operate in an asynchronous manner.

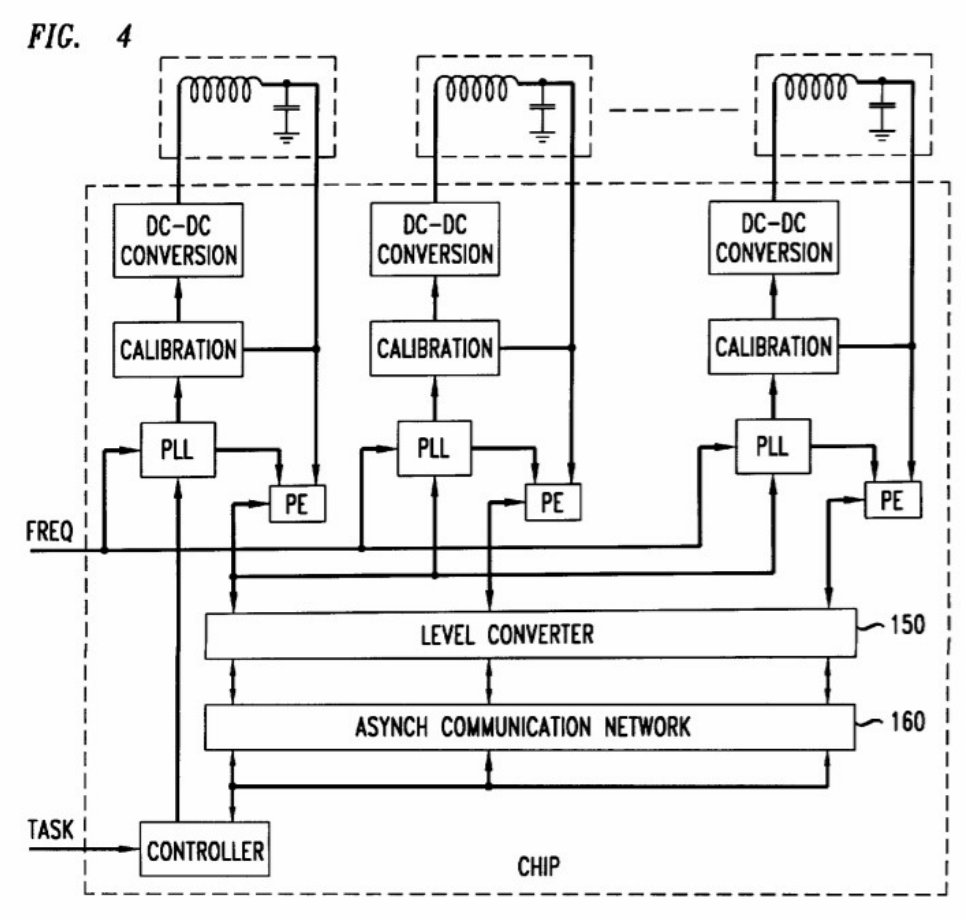

(b) **Nicol in view of Kling does not disclose "setting a clock frequency, of at least a part of the multiprocessor system to a minimum"**

(i) **Petitioner Mixes Embodiments Without Justification**

48. Petitioner points to features from different embodiments in Nicol without explaining how they are interchangeable or how they fit together into a single system. I understand Petitioner starts by showing an annotated picture of

Figure 4, but throughout the text points to passages that describe the Figure 2 embodiment. In my opinion, Petitioner needs to rely on Figure 4 for other limitations of claim 1. As described in Section VII above, in my opinion, a POSITA would not be motivated to combine the Figure 2 and Figure 4 embodiments of Nicol. Figure 2 discloses a system where all of the processing elements operate at the same voltage and frequency while Figure 4 discloses a system where all of the processing elements operate on their own, and potentially different, voltages and frequencies. These two designs are diametrically opposed. Not only would a POSITA not be motivated to combine them it is not obvious to me, today, how they could be usefully combined.

49. For example, I understand that Petitioner points to a disclosure from Nicol that the controller "allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage." Ex. 1009, 2:23-27. The equalizing of processor load describes the process followed in the context of Figure 2, where the entire chip (all of the processing elements) are set to the same frequency. In my opinion, that scenario is the motivation for equalizing the processing load, so that no PE takes (significantly) longer to complete its tasks. This scenario is different from Figure 4, where "the operating system [can] independently program the lowest operating

frequency and corresponding lowest supply voltage for each PE." Ex. 1009, 5:67-6:2.

50. I understand that Petitioner then relies on disclosure from Nicol that states: "If the number of instructions that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the tasks to the available processors in order to balance the computation." Ex. 1009 (Nicol) at 3:13-18. I understand the Petitioner to omit the very next sentence which states: "The intermediate goal, of course, is to maximize the parallelism and to evenly distribute the load presented to the Fig. 2 system among all of the PE's." *Id*. at 3:19-20. Again, it is my opinion that this description relates to the Figure 2 embodiment where the load is distributed as evenly as possible. This situation is the opposite to that which the Figure 4 embodiment addresses: "In some applications, however, one may encounter tasks that cannot be partitioned into concurrent evenly-loaded threads" *Id*. at 5:60-61. Thus, in my opinion, not only does Petitioner attempt to point to two separate embodiments in Nicol without providing any motivation to combine or explain how they may fit together – these two embodiments are not intended to be combined as they use different modes of operation to address different scenarios, and as I explained above I do not believe they can be profitably combined.

(ii) Does not show "in accordance with number of pending operations of a first processor"

51. I understand Petitioner to point to an arbitrary "PE" in the annotated Figure 4 as the First Processor (PE1), but in my opinion, Petitioner does not show that the clock frequency is set based on the number of pending operations of PE1. For Figure 4, Nicol describes that "[a] separate power supply for each PE in a chip overcomes this limitation by allowing the operating system to independently program the lowest operating frequency and corresponding lowest supply voltage for each PE… Each PE in Fig. 4 needs an independent controller that performs the functions of PE 100 (except it does not divide tasks among PEs)." Ex. 1009 (Nicol) at 5:66-6:6. That is, as I understand it, rather than using the technique described for Figure 2 (divide the tasks among the PEs as evenly as possible), each PE can set its own operating frequency independently and the controller does NOT divide the tasks among the PEs. Thus, in my opinion, PE1 does not set the frequency of the system, but only its own operating frequency.

(iii) Gated clock disclosure is misinterpreted

52. I understand Petitioner also points to a gated clock disclosure where certain clock frequencies are set to zero, and claims it is additional disclosure of the minimum frequency required by this claim limitation, but Petitioner misreads Nicol.

Petitioner points to the disclosure "the operating system employs gated clocking techniques to 'shut down' PEs that are not needed," (Nicol at 5:42-44) and simply assumes without considering the context that this is done in accordance with the number of pending operations. But read in context, this is not what Nicol describes. The fuller passage is below:

> Of course, V_{dd}-local can only be reduced so-far before the circuits start to fail, at which point the operating system employs gated clocking techniques to 'shut down' PEs that are not needed.

Ex. 1009 (Nicol) at 5:41-44. In my opinion, Nicol is here describing the need to shut down certain circuits because they fail due to low voltage. This passage does not describe how clock frequencies are set according to the number of pending operations (of any processor). Therefore, it is my opinion that this is not a disclosure of "a minimum in accordance with a number of pending operations of a first processor" as required by the claim.

(iv) **ACPI is not disclosed in either Nicol or Kling, and Bhatia is not part of the ground**

53. I understand Petitioner further relies on states disclosed in the Advanced Configuration and Power Interface (ACPI) Specification for this limitation. But ACPI is mentioned in neither Nicol nor Kling. Instead, I understand Petitioner and its expert Professor Sechen to simply assert, in conclusory fashion,

that ACPI "was commonly used and referred to by a POSITA" (Petition at 38; Sechen Decl. at ¶ 170) without any explanation. In order to overcome the fact that neither of the references, Nicol or Kling, that are part of the ground mention ACPI, I understand that Petitioner introduces Bhatia (Ex. 1011) and the ACPI Specification itself (Ex. 1020). I understand that neither of these are part of the ground, and in my opinion, Petitioner doesn't provide a motivation to combine these references with Nicol or Kling. Furthermore, I understand Petitioner asserts that a POSITA would know to use a particular low power mode C3 without any explanation other than that it preserves memory contents. I disagree. Nicol and Kling do not use preservation of memory contents to determine the clock frequency. Indeed, in my opinion, if this is the reason a POSITA would look to the C3 state, this runs counter to the claim itself. The claim requires that the clock frequency is set in accordance with a number of pending operations of a first processor. The claim does not state that the preservation of memory contents is a consideration in setting the clock frequency. Therefore, in my opinion, Petitioner never ties the C3 state to the number of pending operations of a first processor, and only mentions a vague reference to low-power mode.

(v) **Inconsistent interpretation of "minimum" is fatal to the motivation to combine.**

54. For this limitation, I understand Petitioner points to three separate examples of a "minimum": 70 MHz disclosed in Figure 3 of Nicol, clock gating with a frequency of zero, and the ACPI C3 state. As discussed above, it is my opinion that these three disclosures are described for different purposes, and a POSITA would not be motivated to combine them. But further, in my opinion, Petitioner provides no motivation why a POSITA would look to use multiple different modes of operation with different minimum frequencies in a single combined embodiment. The plain meaning of the word "minimum" is to take the lowest, and in my opinion, a POSITA would not try to combine these three different disclosures if the purpose were to try to find a minimum, and instead focus on using the lowest one if power savings were the motivation. In my opinion, the disclosure of three different "minima" does not suggest that a POSITA would be motivation to combine the different references or disclosures, and in fact makes it worse as it highlights and contrasts the different purposes of each.

(c) **Petitioner has not Carried its Burden to Show That Nicol in view of Kling discloses "increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance**

with a number of pending operations of a second processor"

(i) **Petitioner Mixes Embodiments Without Justification**

55. As with the limitation immediately above, I understand Petitioner to mix embodiments by citing to an annotated version of Figure 4 while quoting portions of Nicol describing the Figure 2 embodiment. It is therefore my opinion that Petitioner's showing on this limitation therefore fails for the same reasons discussed above in Section VII.

(ii) **Petitioner Fails to Establish a "second processor" and that a clock frequency is set "in accordance with number of pending operations of a second processor"**

56. I understand that this limitation requires "increasing the clock frequency of the at least ***the*** part of the multiprocessor system to a maximum," where "the" part of the multiprocessor system must be the same part of the multiprocessor system that was identified for the previous limitation, which describes "a part of the multiprocessor system." I further understand that Petitioner identifies specific PE's for the first processor and the second processor in its annotated Figure 4:

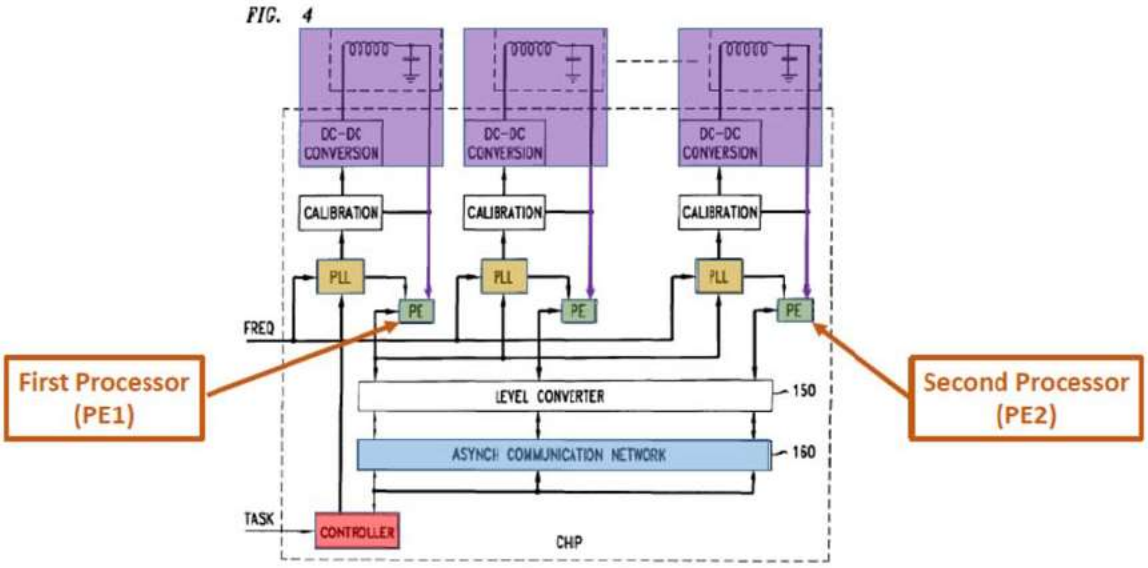

Petition at 60.

57. Nicol's Figure 4 embodiment has "[a] separate power supply for each PE…allowing the operating system to independently program the lowest operating frequency and corresponding lowest supply voltage for each PE." Nicol, 5:66-6:2. Thus, in my opinion, each PE's operating frequency is set independently based on its own load. PE1's operating frequency is set according to PE1's load, PE2's operating frequency is set according to PE2's load, etc. In the Figure 4 embodiment, it is my understanding that the clock frequency is not set according to what Petitioner describes ("Specifically the controller takes into account the number of pending instructions for each PE – including PE1 ("first processor") and PE2 ("second processor") – because it sets the clock frequency based on identifying the processing

requirement of 'the most heavily loaded PE.'" Pet., 40; Sechen Decl. at ¶176. As described in Section VII, the reference to "most heavily loaded PE" occurs in the context of the Figure 2 embodiment, which is not applicable here for Figure 4.

58. For the previous limitation, I understand Petitioner pointed to PE1 as the "first processor" and thus "***a*** part of the multiprocessor system" referred to PE1 (the only part of the system whose frequency is determined by the load on PE1), while I understand that here Petitioner points to PE2 as the "second processor," whose load only determines the frequency of PE2, which must map to "***the*** part of the multiprocessor system." If Petitioner is asserting that PE2 is "***the*** part of the multiprocessor system," it is my opinion that it is different from "***a*** part of the multiprocessor system" from the previous limitation, and thus Petitioner has failed to show how this limitation is met by the disclosures of Nicol. Alternatively, if Petitioner is asserting that "***the*** part of the multiprocessor system" is PE1, it is my opinion that the clock frequency of PE1 is not determined "in accordance with the number of pending operations of a second processor (PE2)."

(iii) **Petitioner has failed to establish "in accordance with number of pending operations of a second processor" for this limitation in any embodiment**

59. If I were to assume that the Figure 2 disclosure in Nicol is applicable

for the Figure 4 embodiment which Petitioner identifies in its annotated figure, it is my opinion that Petitioner has still failed to identify the clock frequency is set "in accordance with the number of pending operations of a second processor." Petitioner (and its expert states) that "it [the controller] sets the clock frequency based on identifying the processing requirement of 'the most heavily loaded PE.'" Pet., 40; Sechen Decl. at ¶176. In my opinion, this is the same way clock frequency was determined for PE1, as Petitioner also pointed to disclosures describing Figure 2, where the system attempted to "equalize processing load among the chips." It is further my opinion that Petitioner presents no reason why the description of the same technique for determining clock frequency in the Figure 2 embodiment of Nicol would consider a different PE, when the disclosure of Nicol is clear that the "most heavily loaded PE" is used to set the clock frequency in that embodiment. Therefore, in my opinion, Nicol does not disclose setting a clock frequency in accordance with the number of pending operations of a second processor – in the Figure 2 embodiment, the clock frequency always relies on the load of the most heavily loaded PE.

(iv) **Nicol Does Not Teach Setting Both PEs to 140 MHz**

60. I understand Petitioner identifies 140 MHz in its annotated Figure 3 of

Nicol as the maximum frequency, and suggests "the clock frequency of the PEs" is set to this value in Nicol. I understand Petitioner cites to 3:31-33 as support for this, but a closer look at Nicol suggests this misreads Nicol. Nicol describes an example where "a given application that is executed on 1 PE requires operating the PE at 140 MHz" (Nicol at 3:31-33). Nicol further suggests that "the application is divided into two concurrent tasks and assigned to two PEs that are designed to operate at 140 MHz from a 2.7 V supply, then the PEs can be operated at 70 MHz and at a supply voltage of 1.8 V." (Nicol at 3:34-37). Nicol explains: "It should be understood that in the above example, when two PEs are employed and their operating frequency can be reduced to 70 MHz, the indicated reduction presumes that it is desired to perform the given tasks as if there was a single PE that operates at 140 MHz." (Nicol at 3:43-47). That is, in my opinion, Nicol teaches dividing the task so that it runs on two separate PEs each running at 70 MHz. Nicol does not teach running both PEs at 140 MHz, and suggests instead running at lower frequencies to save power. Petitioner's suggestion that the system is set to 140 MHz for all of the PEs in the system is not what Nicol teaches. I understand that Professor Sechen agrees that Nicol does not suggest setting both PEs to 140 MHz in this example. *See* Ex. 2027, 16:14-17:11. In my opinion, Petitioner provides no explanation how the example described in Nicol with respect to Figure 2 supports setting all PE's of the system

140 MHz in Figure 3 "in accordance with the number of pending operations of a second processor."

(v) **Nicol's purpose is to reduce power**

61. I understand that Petitioner points to Nicol's disclosure of "[f]or a maximally utilized system, all of the available processors ought to be operating at full speed when satisfying the maximum load encountered by the system." (Nicol at 2:65-67). Nicol goes on to say that "typically, computers spend 99% of their time waiting for a user to press a key. This presents a great opportunity to drastically reduce the average power consumption." Ex. 1008, 3:4-7. Indeed, Nicol's title is "Power Reduction" and the abstract discusses "minimiz[ing] overall power consumption" and the summary of the invention also suggests "minimiz[ing] overall power consumption." Ex. 1008, 2:22). I understand Petitioner also characterizes Nicol in such a way, describing Nicol as "directed to reducing power consumption of a multiprocessor system or chip" Pet., 11; Sechen Decl. at ¶91). I understand this, in other words, to mean Nicol's teaching is about how to avoid running the system at full speed as much as possible, and its descriptions focus on how to avoid running the system at the maximum frequency. I understand Professor Sechen has characterized Nicol as follows in a declaration in an *inter partes* review for the '047 patent which shares the same specification and for which Petitioner has asserted the

identical reference Nicol: "Nicol generally discloses power minimization of a multiprocessor DSP chip." (Sechen '047 Decl at ¶215). Under this characterization, it is my opinion that Nicol does not suggest running the system at a maximum frequency, and instead describes the maximal load as a condition to avoid. In addition, Nicol never connects this maximal load to the number of pending operations of PE2 that Petitioner identified. It is my opinion that Petitioner also never connects this disclosure in Nicol to the embodiments in Figure 2 (which teaches away from using 2 PEs at 140 MHz) or Figure 4, where each PE is set according to its own load.

(d) **Petitioner has not Carried its Burden to Show That Nicol in view of Kling discloses "reducing the clock frequency of the at least the part of the multiprocessor system in accordance with . . . an operating temperature threshold preventing over-temperature"**

62. I understand Petitioner admits that Nicol does not discuss overtemperature. Pet., 24. I understand that Petitioner therefore relies entirely on Kling. *Id*. at 43. For at least the reason discussed in section VII.G., it is my opinion that Petitioner fails to provide sufficient motivation to suggest Nicol and Kling can be combined, and thus fail to establish Nicol and Kling teach this element. In addition, Nicol discusses voltage scaling based on temperature variations, but not

clock frequency changes. Therefore, it is my opinion that a POSITA would understand Nicol to teach away from reducing a clock frequency due to temperature because doing so would prevent Nicol from completing tasks on time.

(e) **Petitioner has not Carried its Burden to Show That Nicol in view of Kling discloses "reducing the clock frequency of the at least the part of the multiprocessor system in accordance with . . . a hysteresis characteristic"**

63. I understand that Petitioner relies entirely on Kling to teach this claim element and only on a single passage at that: "In accordance with this embodiment, the lower threshold may be set to a predetermined value selected to provide hysteresis to reduce the occurrence of power oscillation between the upper and lower thresholds." Pet., 44, *quoting* Ex. 1010, 5:51-54. I understand that Petitioner argues that "a POSITA would have found it obvious to apply the well-known two threshold method (providing hysteresis) from Kling to implement an over-temperature monitoring system in Nicol . . . to reduce unwanted oscillations, and thus to improve the reliability and performance of Nicol's power management system." *Id*. at 45.

64. Once again, it is my opinion that Petitioner confuses benefit and motivation, relying on hindsight rather than focusing on what a POSITA would have been motivated to do at the time of the invention. Kling describes that hysteresis could be used to reduce power oscillation. Ex. 1010, 54-59. In my opinion, Kling

provides no explanation as to how hysteresis would be implemented (power limits do not constitute a particular form of hysteresis – they are merely input conditions to a hysteretic system) or how it could be implemented in alternative systems such as Nicol. Nicol, for its part, makes no mention of hysteresis and provides no insight as to how the undescribed hysteresis method from Kling could be introduced. Accordingly, in my opinion Petitioner relies on benefit alone and provides no motivation to add the use of hysteresis to Nicol.

65. Nonetheless, I understand Petitioner argues, based on Kling's single mention of the term hysteresis and Ex. 1021 (Kennedy, M. & Chua, L., Hysteresis in Electronic Circuits: A Circuit Theorist's Perspective, International Journal of Circuit Theory and Applications, Vol. 19, (1991)("Kennedy")) that Kling's undescribed implementation of hysteresis was so well known that a POSITA would be motivated to add it to Nicol to reduce unwanted oscillation. I disagree for three reasons: (1) Petitioner does not point to any unwanted oscillation in Nicol or a need to reduce such unwanted oscillation through hysteresis; (2) a single mention of hysteresis in Kling is insufficient to show that the undescribed implementation of hysteresis was so well known; and (3) Petitioner interprets Kennedy to suggest that hysteresis was "'indispensable' and commonly used 'in noisy environments and to guard against false triggers,'" Pet., 44, but Kennedy makes no such claim. I believe

the first two reasons are self-explanatory. With respect to the third, I understand Petitioner to quote Kennedy out of context. Kennedy does not state that hysteresis is indispensable in noisy environments; Kennedy recites "[hysteresis] elements have become indispensable building blocks ***in a variety of noise rejection and oscillator circuits***." Ex. 1021, 471 (emphasis added). That is, when using certain noise rejection and oscillator circuits, hysteresis elements are indispensable. Not for a full system such as Nicol. Not for all circuits. Not for all types of noise rejection and oscillator circuits. I understand Kennedy to teach only that Hysteresis is indispensable for certain noise rejection and oscillator circuits. In my opinion, Petitioner points to neither a noise rejection nor an oscillator circuit in Kling and Nicol. Indeed, neither reference discusses specific circuit level architecture. Accordingly, in my opinion, it is inaccurate, based on Kennedy, to say that hysteresis was indispensable in all systems that include unwanted noise (which is every computer system that I am aware), and in my opinion, Petitioner present no evidence that hysteresis would be indispensable in Nicol. Indeed, Nicol describes an alternative method for dealing with noise. Ex. 1009, 5:29-31 ("[T]he frequency is gradually programmed into the system . . . [to] prevent[] excessive noise.").

B. Petitioner has not Carried its Burden to Show That Nicol in view of Kling Renders Obvious Dependent Claim 2 ([SL])

66. In my opinion, Petitioner has not shown that Nicol in view of Kling renders obvious the "clock frequency [being] determined in part by processor voltage."

67. I understand that Petitioner does not argue that Nicol alone discloses this element. In Nicol, the clock frequency is determined by the processing load and the completion time, not by the voltage. As stated in the Abstract, "[the] controller in [Nicol's] multi-processor chip allocates tasks to the individual processors to **equalize processing load** among the chips, **then** the controller **lowers the clock frequency on the chip to as low a level as possible while assuring proper operation**, and **finally reduces the supply voltage**." Ex. 1009, Abstract (emphasis added). Specifically, "the operating system of PE 100 needs to **ascertain the required completion time**, divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and **adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time**." Ex. 1009, 3:66-4:5. In my opinion, a POSITA would have understood that Nicol's clock frequency is jointly determined by the processing load and the required

completion time. Specifically, Nicol's clock frequency is minimized to the value necessary to finish the tasks within the required completion time. A POSITA would have understood that the clock frequency in Nicol is not determined by voltage but rather that the voltage is determined by the frequency.

68. I understand Petitioner relies on Kling for this element. *See* Pet., 46. Kling relates to sending throttle signals "if the power consumed by the portion of the computer system reaches a threshold." Ex. 1010, Abstract. The power is measured by a meter, "tak[ing] a measurement of a parameter that is approximately proportional to the power consumed from the power supply by at least a portion of a computer system. This parameter may include, for example, the current or **voltage between the power supply and the portion of the computer system**, or the duty cycle of a switching signal in the power supply." *Id*. (emphasis added). I understand that Petitioner alleges that the voltage measured in Kling discloses the claimed "processor voltage." See Pet. at 46 ("'a parameter that is approximately proportional to the consumed power, ***such as voltage*** . . .' is tracked. Ex. 1010, 2:31–37. . . . "[C]urrent consumed by a processor is tracked by measuring 'a voltage across a resistor' of the processor as would have been well-known in the art." Ex. 1010, 3:57–59; Ex. 1001 ¶195.).

69. I understand Petitioner further alleges that "current consumed by a

processor is tracked by measuring '**a voltage across a resistor' of the processor** as would have been well-known in the art." Pet., 46 (emphasis added). In my opinion, the underlined statement is incorrect, because the referenced "resistor" is not part of the processor. In my opinion, a POSITA would have understood that the voltage measured in Kling is not the processor voltage, but is the voltage drop between the power supply and the processor. For example, the Abstract quoted above explicitly identified the measured voltage being the "voltage between the power supply and the portion of the computer system." Ex. 1010, Abstract. Claim 23 of Kling says: "the parameter includes a voltage across a **resistor between the power supply and the portion of the computer system**." And claim 16 explains that the "portion of the computer system" is the component that consumes power. In my opinion, a POSITA would have understood that the voltage measured in Kling is the voltage drop between the power source and the load (*e.g.,* a processor). Such voltage drop is not processor voltage because processor voltage is the voltage between the load (*e.g.,* a processor) and the ground. Therefore, a POSITA would have understood that Kling fails to disclose a measurement of the "processor voltage" and Kling's throttling is not based on the "processor voltage."

70. Even if I assume for argument sake that Kling's throttling is based on the "processor voltage," it is my opinion that a POSITA would not have a motivation

to combine Kling with Nicol. Nicol already minimizes the power consumption based on processing load and required completion time. Ex. 1009, Abstract ("by dynamically controlling processing load of chips and controlling . . . the operating voltages of those chips, [Nicol] minimize[s] overall power consumption."); 3:66-67 ("the operating system of PE 100 needs to ascertain the required completion time"). In my opinion, a POSITA would have understood that a further reduction of clock frequency based on supply voltage would cause the tasks to not finish on time impacting performance, and would not have a reasonable expectation of success in combining Kling with Nicol.

71. Thus, it is my opinion that Petitioner fails to establish that Nicol in view of Kling renders this claim obvious.

C. Petitioner has not Carried its Burden to Show That Nicol in view of Kling Renders Obvious Dependent Claim 3 ([DH])

72. In my opinion, Petitioner has not shown that Nicol in view of Kling discloses "said clock frequency is determined in part by available power." Petitioner relies on Kling for this element. I understand Petitioner quotes Kling as disclosing "the upper limit is a constant value … associated with an approximately maximum power that can be reliably consumed by the computer system." Pet., 48. I further understand that Petitioner argues that "[a] POSITA would have understood that this

functionality can be readily implemented in Nicol's system with reasonable expectation of success" because "[t]he functionality involves relatively simple add-on software to Nicol's more complex task scheduler software residing in the controller, and relies on well-known measurement techniques, such as measuring a voltage across an appropriately placed resistor, and would have yielded predictable results." I disagree with Petitioner's argument for three reasons: (1) Petitioner misrepresents what is disclosed by omitting not just the context of this quote, but the remaining portion of the sentence, which undercuts Petitioner's claims; (2) Petitioner has not shown that Kling's frequency is dependent on an available power within the context of a changing clock frequency in claim 1, from which claim 3 is dependent; (3) Petitioner's claim that Kling's maximum power limitation is "simple add-on software" that could have been readily implemented with "an appropriately placed resistor" is false.

(a) **Petitioner Misrepresents Kling**

73. I understand that the sentence from Kling, on which Petitioner relies, recites "the upper limit is a constant value … associated with an approximately maximum power that can be reliably consumed by the computer system ***from a power outlet into which the system is plugged before the outlet's fuse or circuit breaker trips***." Ex. 1010, 5:22-25 (emphasis added). With this full context, it is my

opinion that Kling is properly understood to describe a system that sets its own voltage limit that is reliably below a maximum voltage that can be supplied before an outlet fuse is tripped. Therefore, Kling is not limited by available power; it is limited by a self-imposed safety measure.

(b) **Petitioner has not shown that the clock frequency is determined in part on available power**

74. I understand that independent claim 1, from which claim 3 depends, describes a dynamically changing clock frequency that requires an initial setting of a clock frequency, subsequently increasing the clock frequency to a maximum frequency, and reducing the clock frequency in accordance with an operating temperature threshold and hysteresis characteristic. Claim 3 specifies that this frequency must be determined, in part, by the available power. I understand Petitioner points to a static upper voltage limit as the available power upon which the clock frequency depends. In my opinion, this argument is inconsistent with the claim language. First power and voltage are very different concepts and not at all equivalent. Furthermore, only in the case where the frequency is increased to a maximum would the static upper limit potentially impact that clock frequency. In the setting of the initial clock frequency and the reduction of clock frequency based on temperature and hysteresis, the static upper voltage limit would have no effect.

Indeed, it is my opinion that because the upper voltage limit is static, there would be no need to base the clock frequency on the limit.

(c) Incorporating Kling's upper power limit would not be a simple matter as Petitioner asserts

75. In my opinion, the idea that a POSITA would easily incorporate the invention of Kling into Nicol via simple add-on software is untrue. Balancing upper limits with Nicol's power management system would have been a complex task in which a POSITA would have to balance the power needs of individual processors to accomplish tasks within a designated time along with temperature consideration which is also an element of the claim. A POSITA may spend their entire career developing such a system that responds appropriately and efficiently to these variables.

D. Petitioner has not Carried its Burden to Show That Nicol in view of Kling Renders Obvious Dependent Claim 4 ([DH])

76. Dependent claim 4 recites "said clock frequency is determined in part by a user setting." For a user setting, I understand that Petitioner refers to Kling lower threshold which Kling describes as being modifiable by a user. Pet., 49. However, in my opinion, the frequency does not depend on a lower threshold. The frequency depends on either the upper threshold, which causes throttling, or normal operating conditions which sets a high core frequency. Ex. 1010, 10:10-13 ("If the

processor does not receive the throttle signal (or receives a deasserted throttle signal), normal operation of the processor at the high core frequency continues"). The lower threshold determines only whether the frequency ceases to depend on the upper threshold. *Id*. 8:43-45 ("The pipeline continues to be stalled at step 610 until the power has reached a lower threshold . . . ."). Accordingly, assuming for argument's sake that the lower threshold can fairly be characterized as a user setting, it is my opinion that Petitioner has failed to show that the frequency is dependent upon it.

E. Petitioner has not Carried its Burden to Show That Nicol in view of Kling Renders Obvious Dependent Claim 5 ([DH])

77. In my opinion, Petitioner has failed to show that Nicol and Kling disclose the elements of claim 5 for at least the reasons discuss in Sections VII.C. and VII.D. In addition, Claim 5 requires that a clock frequency depend on both available power and a user setting. In my opinion, Petitioner's arguments with respect to each of these elements fails to detail a situation in which a clock frequency is dependent on both these elements – according to Petitioner's analysis, the clock frequency depends on available power OR a user setting.

78. I understand Petitioner argues, with reference to annotated Figure 2A (reproduced below), that the upper limit is available power, Pet., 47-48, and the

lower limit is a user setting, *id*. at 49-50. I understand Professor Sechen explains that "if the measured power reaches the upper limit, the controller sends a 'throttle signal' to the processor to instruct it to reduce its clock frequency to reduce power consumption." Ex. 1001, ¶211. Professor Sechen further quotes Kling as reciting "[o]nce [the] lower threshold is reached, the throttle signal may be deasserted, and, in response, the IC resumes normal operation." *Id*. at 212; Ex. 1010, 5:48-50. Under Professor Sechen's analysis of Kling, the clock frequency is either (1) throttled according to the upper limit or (2) not throttled and determined according to normal operational techniques if the upper limit has not been reached, or if the lower limit has been reached. At no point does the clock frequency depend on both the upper and lower limits. Furthermore the available power is set based upon the current carrying capability of a power supply external to the computer system, for example a circuit breaker elsewhere in a building. In contrast the upper limit discussed here is set by either a user or system software with no disclosed knowledge of what the actual available power is. In fact, a POSITA would know that the only way to determine the available power would be to increase the power consumption of the system until an external failsafe, such as a circuit breaker or fuse, was triggered. Accordingly, in my opinion, at no point does Kling's clock frequency depend on both an available power and a user input under Petitioner's own analysis.

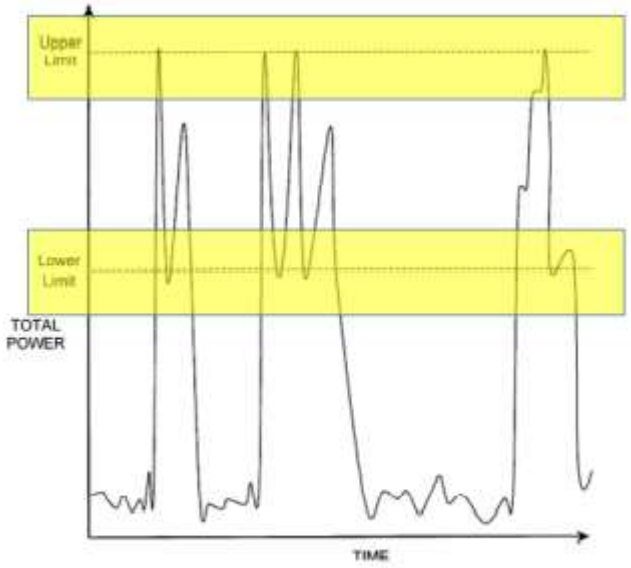

Figure 2A

F. Petitioner has not Carried its Burden to Show That Nicol in view of Kling Renders Obvious Dependent Claim 6

79. I understand Petitioner relies on Nicol alone for this element. It points to Nicol's disclosure of "the number of instructions" and the existence of an "instruction queue." Pet. at 51. However, in my opinion, a POSITA would not have concluded that this disclosure means the number of instructions or pending operations is "determined by" the fill level. There are many ways to determine the number of pending operations, such as "by taking an average of recent workload and assuming the average represents the near-term workload." Pet. at 51-52. In my opinion, Nicol simply does not disclose any particular way to determine the number of instructions.

80. I understand Petitioner points to the file history of the '605 patent, alleging that Patent Owner admitted that "the fill level of a data buffer [is] representative of the number of waiting operations or processor load." Pet. at 51. However, it is my opinion that this statement does not render the claim obvious. It was made in a context of written support rejection. Specifically, Patent Owner explained to the Examiner that the patent specification discloses "a clock frequency set in accordance with a processor load" because it discloses "the clock frequency may . . be determined from the filling level of the particular data buffers" and a POSITA would have known that the filling level of the data buffers indicates the processor load. Ex. 1004 at 377. However, in my opinion, this does not mean the reverse is true. A POSITA would not have understood that processor loads or number of pending operations are necessarily determined through the buffer fill level.

81. I understand Petitioner also points to Cline and Hong to show that inspecting the fill level of buffers was well known. Petitioner appears to provides no analysis as to a POSITA's motivation to apply such knowledge to Nicol and any expectation of success of such application. In my opinion, Cline and Hong shows no more than inspecting the buffer fill level was independently known. Cline and Hong's independent demonstration of elements would not suggest to a POSITA that

the invention was obvious. In my opinion, a POSITA would need to be motivated to combine the teachings of those references to achieve the claimed invention with the expectation of a reasonable expectation of success. In my opinion, Petitioner fails to show such motivation, and thus, fails to establish that claim 6 is obvious in view of Nicol.

G. Petitioner has not Carried its Burden to Show That POSITA Would Combine Nicol and Kling

82. I understand Petitioner proposes three potential motivations to combine Nicol and Kling: (1) Nicol and Kling are directed to thermal and power management aspects that are allegedly critical and ubiquitous in processor-based systems, Pet., 23; (2) Nicol and Kling involve multiprocessor systems, software-managed power supplies for each processor, and immaterial differences between software such that a POSITA would find it obvious to introduce Kling's functionality into Nicol's controller OS, *id*. at 23-24; and (3) Nicol is concerned with power consumption, and over-temperature protection is allegedly critical for a multiprocessor system, and a POSITA seeking to improve Nicol allegedly would have incorporated the benefits of Kling to help with over-temperature protection, *id.* at 24-25. In my opinion, these alleged motivations, taken together, amount to a suggestion that Nicol and Kling manage aspects of a computer system and it would be beneficial to combine the

aspects that they manage. In my opinion, mere benefit in retrospect would have been insufficient to motivate a POSITA to combine these reference.

83. I understand that Petitioner nonetheless alleges, supported only by parroting language of Professor Sechen, that a "POSITA would have found it obvious to program Nicol's controller to handle the additional software functionalities discussed in Kling," and "[i]mplementing this added functionality would not require undue experimentation, has a reasonable expectation of success, and would yield predictable results because the added functionality is primarily add-on software." Pet., 24-25. I disagree. I don't understand Petitioner to make any attempt to explain how a system combining Nicol and Bhatia would function, likely because, in my opinion, managing a plurality of metrics to create an efficient processing system is extremely difficult. A POSITA might spend their entire career developing intricate systems for managing processor temperature and clock frequency among a plurality of processors, and Petitioner seems to dismiss it as a mere triviality.

84. Lastly, it is my opinion that Petitioner ignores any issues that would arise. *Id*. For example, Nicol and Kling change frequencies for substantially different reasons. I understand Professor Sechen recognizes that "Kling teaches a controller … to adjust clock frequency of the processors . . . based on sensor measurements,"

Ex. 1001, ¶123, whereas "Nicol teaches that the controller allocates tasks to the individual processors . . . so that it can equalize the processing load among the PEs, . . . and then the controller lowers the clock frequency on the chip to as low a level as possible" *id*. at ¶164. In my opinion, Petitioner ignores any prioritization issues that would arise from these two different methods for changing frequency. It is difficult to implement them both simultaneously, and there is no discussion in Nicol or Kling to suggest that they could be implemented as seamlessly and simply as Petitioner argues.

85. In fact, the disclosure of Kling would dissuade a POSITA from using a thermal sensor as proposed by Petitioner. As analyzed above, Kling abandoned the indirect way of measuring power consumption through a thermal sensor, so that "temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided." Ex. 1010, 2:56-58. By not using a thermal sensor, "the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition." *Id.*, 2:58-62. Nevertheless, Petitioner argues that "Kling discusses conventional techniques to address those issues. In particular, to address critical overtemperature, a thermal sensor is typically used to monitor the temperature of a processor." Pet. at 16.

However, given Kling's disclosure, a POSITA would not have a motivation to combine the traditional thermal sensor with Nicol because of the stated inaccuracies and slow response time caused by the thermal sensor.

IX. GROUND II FAILS

86. In my opinion, Petitioner fails to show that Nicol in view of Kling and DeHon renders claims 1-6 obvious at least because Petitioner has not shown that a POSITA would be motivated to combine DeHon with Nicol.

87. While I am a POSITA as proposed by the petitioners in this specific area of reconfigurable hardware I have specific expertise. I have published quite widely in the exact area addressed by DeHon, know the man quite well, and in fact have worked with him on reconfigurable hardware. See, for example, William H. Mangione-Smith, Brad Hutchins, David L. Andrews, André DeHon, Carl Ebeling, Reiner W. Hartenstein, Oskar Mencer, John Morris, Krishna V. Palem, Viktor K. Prasanna, Henk A. E. Spaanenburg: Seeking Solutions in Configurable Computing. Computer 30(12): 38-43 (1997).

88. The DeHon processor includes two components, "fixed processor core cell 10" and "a configurable array 100 on the same silicon die or chip" as shown below:

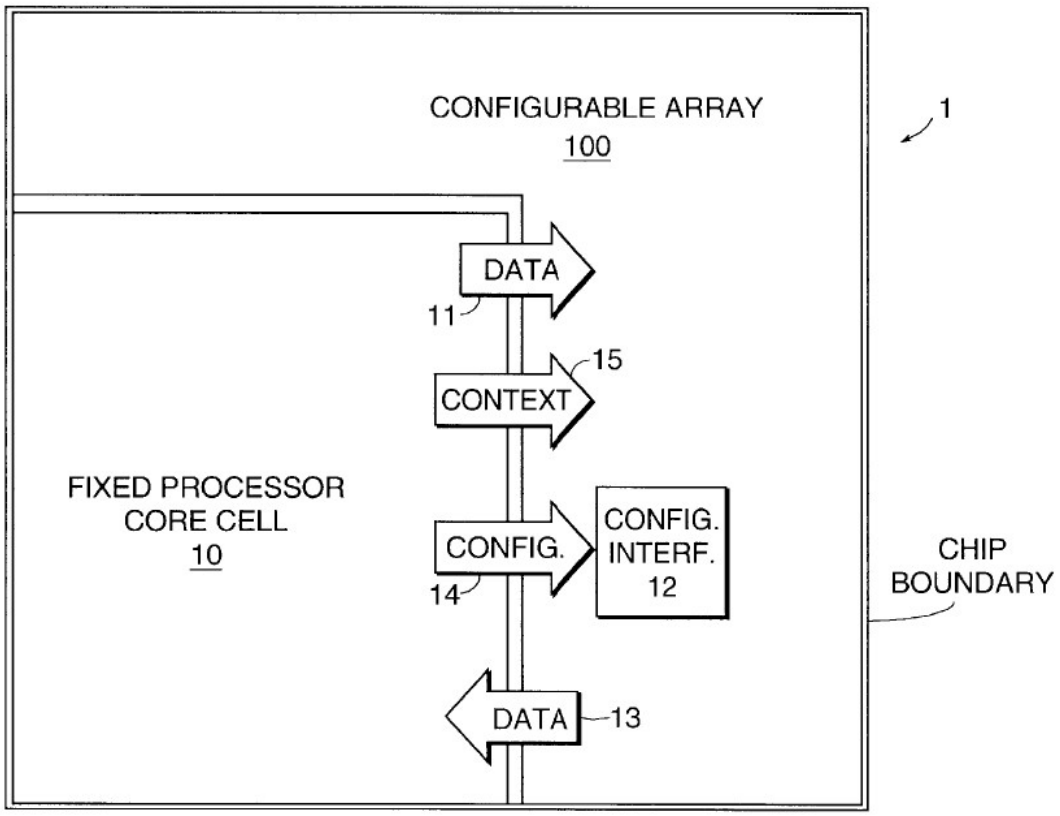

FIG. 1

Ex. 1010, Figure 1, 5:37-39. In order to configure the "configurable array 100," DeHon needs a compiler that "selects the hardware configurations and library routines that best accelerate each application." Ex. 1012, 6:47-65. In the proposed combination, I understand that Petitioner replaces each PE in Nicol with the integrated chip in DeHon so that the combination becomes a multi-chip system. *See* Petition at 55-56 ("The figure [] combines Figure 4 of Nicol with Figure 1 of DeHon to illustrate Nicol's multiprocessor system in which the PEs are DeHon's

reconfigurable processor **chips**.") (emphasis added). In my opinion, a POSITA would have understood that in this multi-chip system, the controller PE runs the operating system. In other words, it is my opinion that in the proposed combination, it is the controller PE that would configure the DeHon "configurable array 100" on separate chips.

89. In my opinion, a POSITA would have understood that this proposed combination violates a key requirement of DeHon. DeHon states that integrating all components on one die is a key feature:

> The combination of the processor cell 10 and the configurable array 100 on the same silicon die is **a key feature** of the invention enabling high bandwidth communications with a low latency interface. **When the configurable array is merely attached to a host processor as an external peripheral, the high overhead implicit in communication with the configurable engine located on a different chip limits the avenues for acceleration and prevents close cooperation**. The low bandwidth similarly makes communication expensive and limits the throughput achievable with the reconfigurable array.

Ex. 1012, 6:19-29. It is my opinion that a POSITA would have understood that the combination proposed by Petition is equivalent to "the configurable array [being] merely attached to a host processor as an external peripheral," where the controller PE becomes the host processor that configures the configurable array on separate chips. This would substantially increase the overhead implicit in communication

between the controller PE and the DeHon processors on different chips. In addition, DeHon states that "the integration of the processor core cell 10 and the reconfigurable array 100 also decreases reconfiguration overhead. As a mere peripheral, the reconfiguration time for the array must be amortized over a large amount of processing to justify the array's reconfiguration." Ex. 1012, 6:31-35. In my opinion, a POSITA would have understood that using the controller PE to configure the array on different chips as proposed by Petitioner would significantly increase reconfiguration overhead, and Nicol does not disclose the amount of processing that could justify the array's configuration time. Thus, it is my opinion that a POSITA would not have a motivation to use separate DeHon chips in the Nicol system. Given the extra overhead introduced in the proposed combination, a POSITA would not have a reasonable expectation of success. For example, the reconfiguration overhead and latency created in the proposed combination would substantially increase the complexity for the "Asynchronous Communication Network 160" and would no longer meet the requirement that the "real-time operating system resides in PE 100."

90. In my opinion, Petitioner does not provide an explain as to why a POSITA would ignore the "**key feature**" of DeHon and apply the DeHon invention in a multi-chip environment. Instead, I understand Petitioner only to touch on the

motivation to combine at a very high level. I understand that Petitioner's motivation to combine analysis in the present case is almost a verbatim copy of its motivation to combine Nicol and DeHon in its petition in IPR2020-00539. *See* IPR2020-00539, Paper 2 at 73-78. There, I understand that the Board summarized the motivation to combine arguments in the Institution Decision. *See* IPR2020-00539, Paper 11 at 34-36. I understand first, Petitioner argues that "Nicol explains that its multiprocessor system is intended for use in a variety of digital signal processing applications," and DeHon teaches that its multiprocessor "can serve as the principle building block for a wide range of applications," so a person of ordinary skill in the art would have chosen DeHon's processor as each of the PEs in Nicol's system. IPR2020-00539, Paper 11 at 34. However, in my opinion, this shows no more than Nicol could potentially be combined with DeHon, which I do not believe would be sufficient to motivate a POSITA to combine the references.

91. Second, I understand Petitioner argues that "a person of ordinary skill in the art would have been motivated to use a reconfigurable processor, such as DeHon's, to meet the needs of specific applications." IPR2020-00539, Paper 11 at 34. I understand Petitioner specifically also points to some advantages in DeHon, such as "handl[ing] a variety of compute-intensive functions 'such as compression, encryption, signal processing, or protocol processing'" and "replac[ing] or be[ing]

deployed alongside traditional fixed processors." Pet. at 57. However, in my opinion, Petitioner does not provide any specific applications in Nicol that would benefit from DeHon's reconfigurable processor or any needs that Nicol alone cannot meet. Given Nicol is a self-contained complete system, it does not have any need to be modified. Accordingly, it is my opinion that Petitioner's motivation is based on hindsight.

92. Third, I understand Petitioner argues that "both references 'relate to the same field of invention, optimizing processor systems by dynamically controlling processors.'" Id. However, in my opinion, a POSITA would understand that the field of processor is very broad. This characterization covers very small slow processors such as those used in pagers (which I worked on at Motorola) and extremely high-performance supercomputers (which I worked on both in graduate school and at Motorola). The characterization covers the controller for the display in a microwave oven as well as the controller for a reactor core on a nuclear submarine. All of these areas attempt to optimize processors systems by dynamically controlling processors. And yet they have little if anything in common beyond that shockingly broad characterization. I am hard-pressed to think of a computer system or computer related patent where that characterization would not apply.

93. I understand that the Board noted, "[Petitioner] did not devote much discussion in its Petition specifically relating to an expectation of success in combining DeHon with Nicol (as modified by Bhatia with respect to Ground 4)." IPR2020-00539, Paper 11 at 35. Given that integrating on one chip is a "key feature of the [DeHon] invention," it is my opinion that Petitioner's failure of explanation of a reasonable expectation of success is fatal to its argument

94. Thus, in my opinion, Petitioner fails to establish the motivation to combine Nicol and DeHon and a reasonable expectation of success of such combination by the preponderance of the evidence.

X. GROUNDS III FAILS

95. In my opinion, Petitioner fails to show that claims 1-6 are obvious over Nicol in view of Kilng, DeHon, and Bhatia at least because Petitioner fails to show that POSITA would be motivated to combine Bhatia with Nicol, Kling, and DeHon.

96. As discussed above, it is my opinion that a POSITA would not be motivated to combine Nicol with Kling or DeHon. A POSITA would also not be inclined to combine Nicol with Bhatia, much less to combine all four reference together.

97. I understand Petitioner proposes three potential motivations to combine Nicol and Bhatia: (1) Nicol and Bhatia are directed to thermal and power

management aspects that are allegedly critical and ubiquitous in processor-based systems, Pet., 65; (2) POSITA would allegedly have found it obvious to add additional functionality in Bhatia to Nicol because the increased functionality would be an improvement, *id*. at 65-66; (3) A POSITA would have had a reasonable expectation of success in doing so because Bhatia explains that those low activity states were commonly used in the industry, *id*. at 66. In my opinion, these alleged motivations to combine amount only to a suggestion that Nicol and Bhatia manage aspects of a computer system and it would be beneficial to combine the aspects that they manage. In my opinion, these allusions to benefit without any showing of actual motivation is insufficient to suggest a POSITA would combine these reference.

98. I understand that Petitioner alleges, supported only by parroting language of Professor Sechen, that a "POSITA would have found it obvious to program Nicol's controller to handle the additional software functionalities discussed in Bhatia," and "[i]mplementing this added functionality would not require undue experimentation, has a reasonable expectation of success, and would yield predictable results because the added functionality is primarily add-on software." Pet., 19-20. I disagree. I understand Petitioner to make no attempt to explain how a system combining Nicol and Bhatia would function, likely due to the fact that managing a plurality of metrics to create an efficient processing system is

extremely difficult. *Id*. A POSITA might spend their entire career developing intricate systems for managing processor temperature and clock frequency among a plurality of processors, and Petitioner dismisses it as a mere triviality.

99. Lastly, in my opinion, Petitioner ignores any issues that would arise from combining Nicol and Bhatia. Indeed, it is my opinion that combining Nicol and Bhatia would frustrate the purpose of each. Nicol is designed to minimize power consumption by **lowering the clock frequency** and therefore the voltage demanded by the processors. *See* Ex. 1006, Abstract; 2:22-27. Conversely, Bhatia is designed to **maintain** a system operating at a high performance, characterized by **a high frequency** which requires a high voltage, without exceeding over-temperature thresholds. 1:47-49; 3:14-23; 6:41-51. Thus Nicol and Bhatia pull against each other.

XI. CONCLUSION

100. I understand that this Declaration will be filed as evidence in a contested case before the PTAB. I acknowledge that I may be subject to cross-examination in the case and that cross-examination will take place in the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

101. I declare that all statements made in this Declaration of my own knowledge are true to the best of my knowledge and that all statements made on

information and belief are believed to be true. These statements were made with the knowledge that willful false statements and the like are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: December 4, 2020 By: _____
William Henry Mangione-Smith, Ph.D