UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,
Petitioner

v.

PACT XPP SCHWEIZ AG,
Patent Owner

U.S. PATENT NO. 9,075,605

Case No. IPR2020-00541

DECLARATION OF DR. CARL SECHEN IN SUPPORT OF PETITIONER INTEL CORPORATION'S REPLY TO PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

Page

I, Carl Sechen, Ph.D., do hereby declare as follows:

I. INTRODUCTION

1. I have been retained as an expert witness on behalf of Intel Corporation ("Intel") for the above-captioned Petition for *Inter Partes* Review ("IPR") of U.S. Patent No. 9,075,605 ("'605 patent") (Ex. 1003). I am being compensated for my time in connection with this IPR at my standard consulting rate of $375 per hour. My compensation is not affected by the outcome of this matter.

2. I previously submitted a Declaration in support of this IPR on February 7, 2020 (Ex. 1001). I apply the same viewpoint of a POSITA as expressed therein.

3. Since preparing my first Declaration, I have reviewed the materials cited in my first Declaration, in addition to the Declaration of William Henry Mangione-Smith ("Mangione-Smith Declaration"), which was submitted by Patent Owner in response. (Ex. 2024), the deposition transcript of Dr. Mangione-Smith (Ex. 1051), the deposition transcript of Dr. Mangione-Smith in the related proceeding for the '301 patent (Ex. 1052), the Wiley Electrical and Electronics Engineering Dictionary (Ex. 1053), the Claim Construction Order in the parallel District Court case (Ex. 1054), Patent Owner's Preliminary Response, the Institution Decision (Paper 11), and Patent Owner's Response (Paper 20). I have also considered my own knowledge

and experience in the field. For the reasons discussed further below, I disagree with Patent Owner's Response and Dr. Mangione-Smith.

II. BACKGROUND AND QUALIFICATIONS

4. My background and qualifications are set forth in my February 7, 2020 Declaration, which is incorporated herein by reference.

III. UNDERSTANDING OF PATENT LAW

5. In my first Declaration, I set forth the applicable principles of patent law that were provided to me by counsel. As appropriate, I have continued to apply those principles in providing my opinions in this Reply Declaration.

IV. GROUND I: NICOL IN VIEW OF KLING

A. A POSITA would have been motivated to combine Nicol and Kling

6. In my opinion, as I explained in my first Declaration, a POSITA would have been motivated to combine Nicol and Kling. Ex. 1001, ¶¶ 122-129.

7. There are numerous reasons a POSITA would have been motivated to combine Nicol and Kling, such as, for example, to use temperature measurements to improve voltage scaling as Nicol teaches and for over-temperature protection as Kling teaches. Similarly, in my opinion, a POSITA would have been motivated to use power measurements for both over-temperature protection as taught by Kling, and to detect excessive power consumption consistent with Nicol's purpose of minimizing power consumption.

8. I disagree with Patent Owner that there would be "prioritization" issues that would arise in combining Nicol and Kling. For example, Kling's over-temperature protection is a critical safety mode of operation such that the controller throttles (or puts the brakes) on the processors in order to avoid catastrophic device meltdown and failure due to over-heating. In my opinion, Kling's emergency braking does not interfere with the normal operation of Nicol's system so long as there is no risk of device failure due to over-heating, such as through excessive power consumption, in which case Kling's features take over Nicol's system with such a risk because "the processor may destroy itself by its own heat." Ex. 1010, 1:44-46.

9. It is also my opinion that a POSITA would not have had difficulty implementing normal power management modes of operation taught by Nicol alongside over-temperature protection modes taught by Kling. In my opinion, both Nicol and Kling teach power management features that have routinely been implemented together in computer systems. As Kling explains that: "Processor and computer system designers have developed numerous methods to deal with these [power management] issues. For example, processor designers implement specialized circuit design techniques *that [like Nicol] reduce power consumption*. In addition, modern computer systems are designed to shut down portions of the system that are

not needed during a particular period of time [*Nicol refers to this selective shutdown of PEs that are not needed as clock-gating*] … Some processor packages include a thermal sensor to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. *If these precautions are not taken, the processor may destroy itself by its own heat*." Ex. 1010, 1:30-46.

10. I also disagree with Patent Owner that Kling's disclosure would dissuade a POSITA from using thermal sensors. POR, 56. Given Nicol expressly discusses "tracking temperature", which requires the use of temperature sensors, (Ex. 1009, 4:11-13), it is my opinion that a POSITA would have understood the advantage of utilizing Nicol's temperature sensor(s) to detect over-temperature conditions as well. Such a use was commonly done in prior art computer systems, as Kling explains: "Some processor packages include *a thermal sensor* to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. *If these precautions are not taken, the processor may destroy itself by its own heat*." Ex. 1010, 1:38-46. Despite Kling's teachings using an additional over-temperature protection mechanism with direct power measurement as a trip point (through power

thresholds), in my opinion a POSITA would have understood that Kling's teaching does not suggest or teach eliminating the conventional trip points based on temperature measurements, despite imperfect temperature sensor placement, using temperature sensors.

11. In my opinion, a POSITA would have understood that temperature sensors are necessary to quickly identify local "hot spots" that power measurement alone could miss. In fact, this is because the amount of power consumed by the system (*e.g.,* a chip) may not be high enough to trigger a power measurement trip point, but still could cause a local temperature on the chip to rise to an unsafe level if most of the power is consumed by one small circuit of the chip, which could create a dangerous "hot spot."

12. Further, as I previously explained, it is my opinion that a POSITA would have found it obvious to program Nicol's controller to handle the necessary over-temperature protection features taught by Kling, only taking over in periods of criticality and not interfering with Nicol's normal operation, which would yield predictable results as the implementation is primarily add-on software.

B. Nicol's embodiments share common functionality

13. In my opinion, the core functionality of Nicol's real-time operating system ("RTOS") software cuts across all embodiments. The core function of the RTOS is to allocate tasks among PEs and control their frequency and

voltage settings accordingly. This is true irrespective of the number and types of PEs, and regardless how PEs are connected to DC-DC converters (which supply their voltages under RTOS control) and PLLs (which supply their clock frequencies under RTOS control). *See* Ex. 1009, 2:18-27, 6:29-33.

14. For example, in Nicol's Figure 4 embodiment, where each PE has its own DC-DC converter and its own PLL, the RTOS has complete flexibility and can independently set the clock frequency and supply voltage for each PE. However, the RTOS can also set the same clock frequency and supply voltage for all PEs, which would be the case when RTOS can evenly divide the tasks such that each PE receives an even computational load, such that the PEs operate in unison exactly as in the Figure 2 embodiment. Ex. 1009, 5:53-65. Alternatively, the RTOS can group them in clusters such that each cluster of PEs operate under the same voltage and frequency as would be the case under the "middle ground" discussed in Nicol. Ex. 1009, 6:66-7:3.

C. Nicol in view of Kling renders obvious Claim 1

1. "being adapted for sequentially processing data"

15. As I explained in my first Declaration, there are numerous reasons for why a POSITA would understand Nicol to disclose and/or render obvious "being

adapted for sequentially processing data in a clocked manner." *See* Ex. 1001, ¶¶ 147-160.

16. In my opinion, as I previously explained, Nicol discloses a network allowing for communication between PE-to-PE. In fact, Nicol explicitly describes a synchronization schema in order to allow data to be communicated between PEs. Ex. 1009, 6:26-28, 6:14-18, Figure 4. In my opinion, therefore, Nicol discloses that through the synchronization schema, Nicol's PEs are adapted for sequentially processing data from one PE to another, satisfying this claim limitation. *See* Ex. 1009, 6:18-21 ("a synchronization schema must be implemented when there is a need to communicate data between PEs (or with other system elements) that operate at different frequencies."), 6:25-33 (discussing the need and the asynchronous communication network solution for "synchronizing the exchange of data among the PEs of a multiprocessor chip").

17. I understand that Patent Owner and its expert disagree with this clear disclosure in Nicol, and contend that Nicol's asynchronous communication network ("ACN") does not meet this claim element, arguing that it does not allow for PE-to-PE communication (*i.e.*, such as between the green PEs below), but rather only communication between the controller and each individual PE. *See* POR, 22-23; Ex. 2024, ¶41.

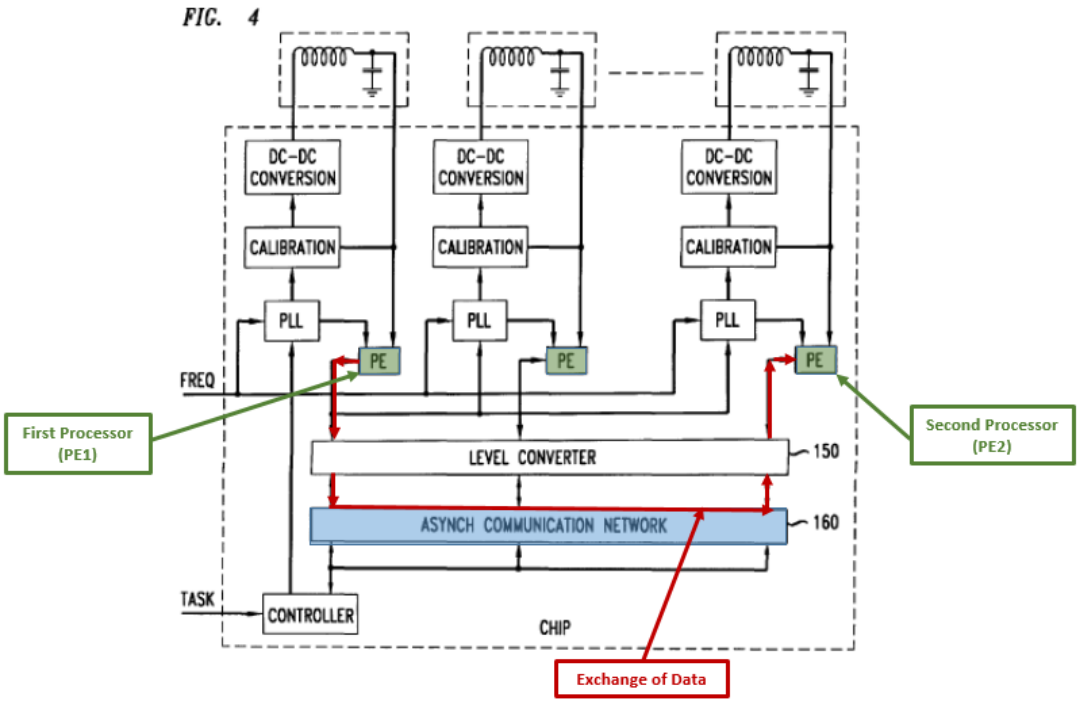

Ex. 1009, Figure 4 (annotated).

18. I disagree that Nicol's ACN would only be used to transfer data between the green PEs and the controller, and instead as I set forth in my first Declaration, I believe a POSITA would have understood that the ACN facilitates data transfer from PE-to-PE. In my opinion, the inter-PE communication and sharing of data satisfies being adapted for sequentially processing data, thereby allowing for sequential data processing to occur as data is passed from the output of one PE to the input of another. Therefore, Nicol discloses that through the synchronization schema, Nicol's PEs are

adapted for sequentially processing data from one PE to another, satisfying this claim limitation. *See* Ex. 1009, 6:18-21.

2. "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"

19. I understand that Patent Owner disagrees that Nicol in view of Kling renders this claim element obvious. POR, 28-34. I disagree. In my opinion, as I explained in my first declaration, Nicol in view of Kling renders this claim element obvious. Ex. 1001, ¶¶ 161-172.

20. Nicol's Figure 3 illustrates the controller, which sets the frequency of certain PEs (*e.g.*, the PEs of Figure 4) initially to 70MHz. Ex. 1009, Figure 3. In my opinion, this would be the case when, for example, the controller evenly distributes tasks among available PEs. In such a case, the PEs of Figure 4 operate in unison just as the PEs in the Figure 2 embodiment.

21. Nicol explains that the Figure 4 embodiment allows the controller to flexibly adjust the voltage and frequency of each PE when tasks cannot be divided evenly among available PEs, however, Nicol emphasizes that "chip-wise voltage scaling" in which the same voltage and frequency is applied to all PEs should be used "when the load of the computation can be evenly distributed across all of the PEs." Ex. 1009, 5:53-6:3; *see also id.*, 3:66-4:5 (explaining that the controller seeks to "divide the collection of tasks as

evenly as possible" among PEs). And, as annotated in Figure 3 below, the initial 70MHz frequency setting is the ***minimum*** frequency needed to handle the tasks at hand by the PEs. This takes into account the number of pending operations of the PE identified as the "first processor" (PE1). Ex. 1009, 2:23-27 ("the controller ***lowers the clock frequency*** on the chip to ***as low a level as possible*** while assuring proper operation."), 3:13-18 (controller "determine[s] the best way to allocate the tasks to the available processors" by taking into account "***the number of instructions*** that need to be executed for each task").

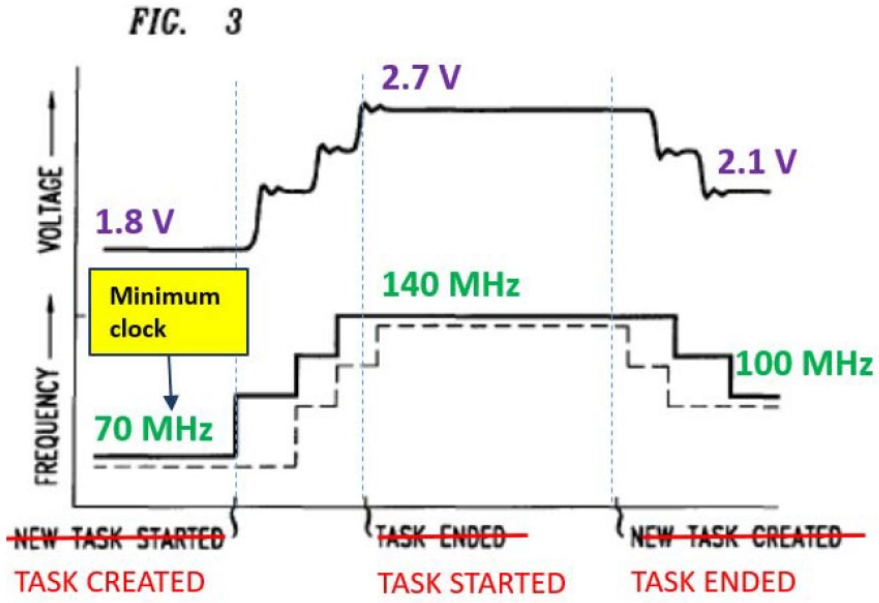

Ex. 1009, Figure 3 (annotated).

22. Moreover, as I previously explained, Nicol's clock gating feature provides another way of satisfying the claim element. Ex. 1001, ¶¶ 168-169. Under

this mechanism, Nicol's controller sets the clock frequency of some or all PEs to zero ("clock gating") when those PEs are not needed, *e.g.,* because the number of pending operations for those PEs is zero (*i.e.*, no load). Ex. 1009, 5:42-44 ("the operating system employs gated clocking techniques to 'shut down' PEs that are not needed."). In my opinion, a POSITA would have known that clock gating was (and still is) commonly used in computer systems to temporarily shut down circuits that are not needed to save power, as discussed in Kling. Ex. 1010, 1:30-46 ("[M]odern computer systems are designed to shut down portions of the system that are not needed during a particular period of time.").

23. Contrary to Patent Owner's argument, I disagree that Nicol's clock gating is not done in accordance with the number of pending operations of the PEs. POR, 31-32. Patent Owner quotes the following passage for support: "Of course, Vdd-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, 5:41-44. However, a POSITA would know that Nicol is simply explaining that PEs are designed to operate over a certain V_{dd} range and that they cannot operate below a minimum design voltage, for example in Figure 1, that minimum voltage is approximately 1.1V. Ex. 1009, 5:42-45, Figure 1. Therefore, in my opinion

when a PE is idling (no pending operations), the controller would not reduce its V_{dd} supply voltage below its minimum voltage to save power because the circuit would fail. Instead, as taught by both Nicol and Kling and as was well-known in the art, the PE should be clock-gated and its voltage maintained at the minimum voltage or higher when the PE is not needed (no pending operations).

24. Therefore, in my opinion, Nicol in view of Kling renders this claim element obvious.

3. "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"

25. As I previously explained, in my opinion, Nicol in view of Kling renders this claim element obvious. Ex. 1001, ¶¶ 173-179.

26. As I previously explained, Nicol's Figure 3 illustrates the controller raising the clock frequency of PEs from 70MHz (minimum clock frequency) to a higher frequency, illustrated as 140MHz (maximum clock frequency) in Figure 3, to handle additional workload. Ex. 1001, ¶177. The higher frequency can be the maximum frequency, annotated below in Figure 3 that the PEs are designed to operate at. Ex. 1009, 2:65-3:4 ("For a maximally utilized system, all of the available processors ought to be operating at full speed"), 3:63-64 (discussing "the maximum [frequency] that the chip can support").

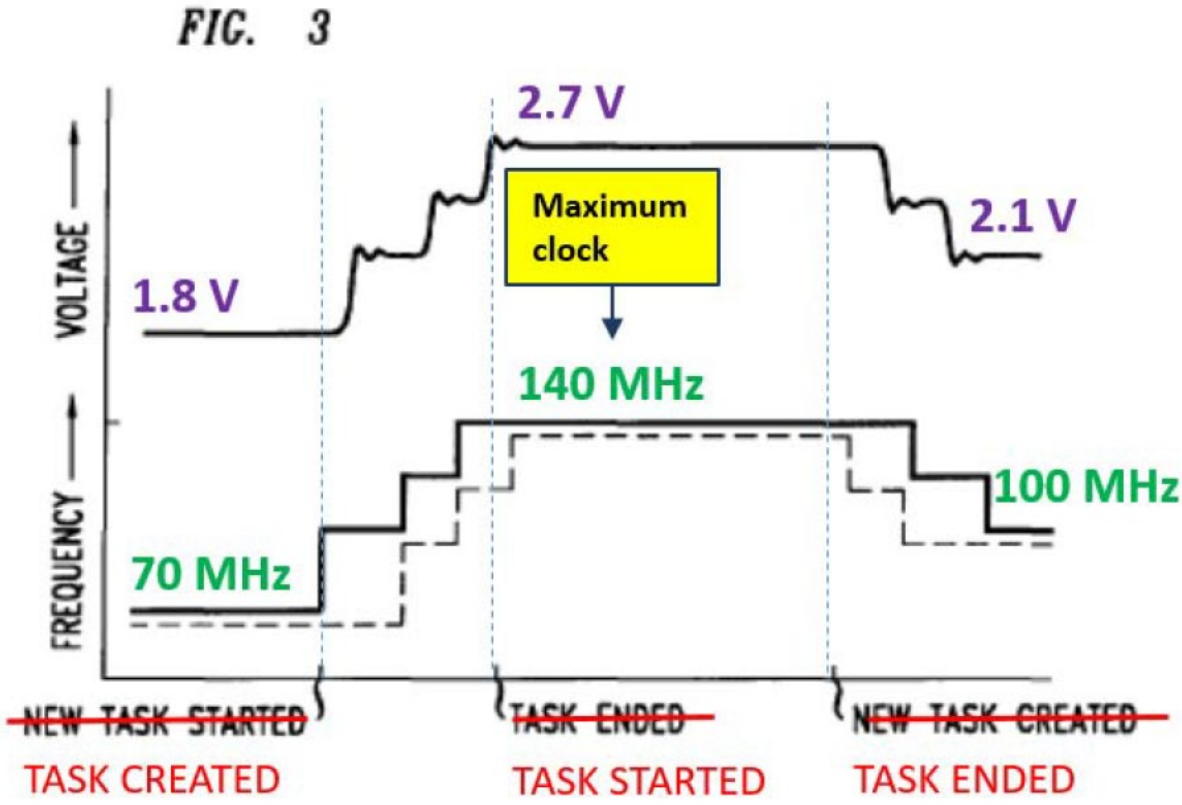

Ex. 1009, Figure 3 (annotated).

27. I disagree with Patent Owner's contention that Nicol's dynamic voltage and frequency scaling ("DVFS") as illustrated in Figure 3 is somehow inapplicable to the hardware configuration illustrated in Nicol's Figure 4. POR, 34. Instead, in my opinion, a POSITA would have understood that Nicol's DVFS is the core teaching of Nicol, and it applies to all hardware configurations, including the Figure 2 configuration (in which all PEs share the same DC-DC converter and PLL), the Figure 4 configuration (in which each PE has its own DC-DC converter and PLL), in and the "middle

ground" configuration (in which each group of PEs shares a DC-DC converter and PLL). As such, it is also my opinion that a POSITA would have understood that Nicol's Figure 3 illustrates DVFS applied to one PE, a group of PEs, or all PEs in unison. *See* Ex. 1009, 2:40-43, 4:54-64.

28. I also disagree with Patent Owner's argument that the controller in Figure 4 can *only* set the clock frequency of each PE according to its own workload. POR, 35. In my opinion, Nicol makes clear that Figure 4 illustrates the most *flexible* hardware configuration whereby the controller *can* independently adjust the clock frequency and voltage of each PE when tasks cannot be evenly divided among PEs, however, Figure 4 is not only limited to such an arrangement. Ex. 1009, 5:53-56 ("The notion of adjusting operating frequency to load and adjusting supply voltage to track the operating frequency can be ***extended*** to allow each PE to have its own supply voltage."). For example, when tasks can be evenly divided among available PEs, in my opinion, Nicol's controller adjusts the PEs' clock frequency and supply voltage in unison ("chip-wise" scaling), taking into account all of the PE workloads. Ex. 1009, 5:56-59. For example, as illustrated below, the controller sets the maximum clock frequency of PE1 (first processor), taking into account the workload (*e.g.,* number of pending operations) of all PEs,

including the PE mapped to the claimed second processor (PE2) when evenly distributed.

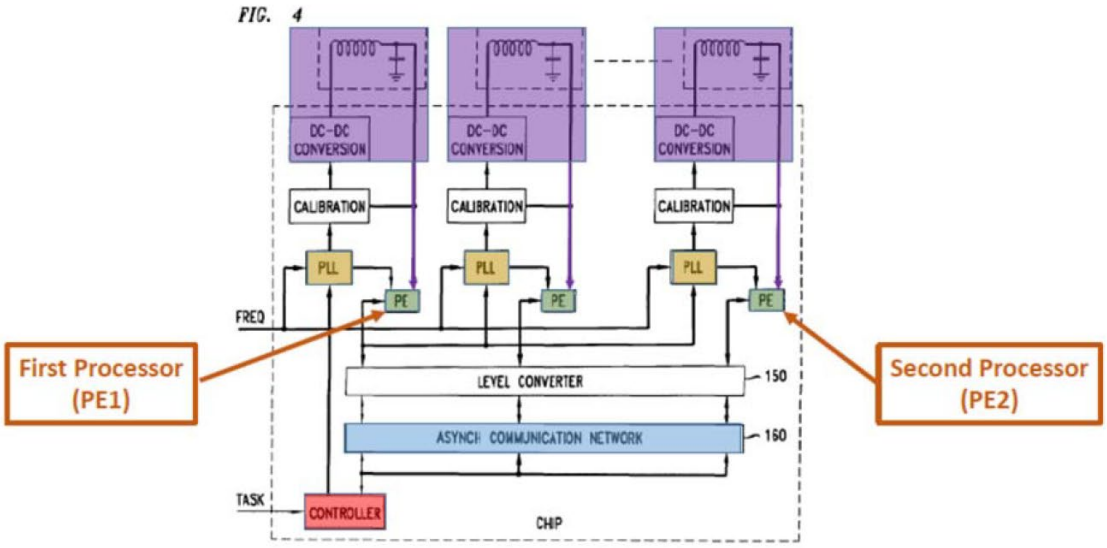

Ex. 1009, Figure 4 (annotated).

29. It is also my opinion that the same PE does not have to be the "most heavily loaded" at all times. Indeed, both PE1 and PE2 can be the "most heavily loaded" when the tasks are evenly divided among available PEs. Those available PEs (which includes PE1 and PE2) would be equally loaded, both the "most heavily loaded" compared, for example, to clock-gated PEs with no load.

30. While the above discussion utilizes Nicol's Figure 4 hardware configuration (wherein each PE has its own DC-DC converter and PLL) as an example, in my opinion, the same applies to both Nicol's Figure 2 and the "middle ground" embodiments where all PEs of the chip (*i.e.* Figure 2) or the PEs

within a group (*i.e.* middle ground) share the same DC-DC converter and PLL. Ex. 1009, 6:66-7:3. In such a case, all PEs on the chip (Figure 2), or PEs within the same group (middle ground) operate in unison, with the same clock frequency and supply voltage applied to the PEs. Ex. 1009, 3:9-42. There, Figure 3 illustrates an exemplary minimum frequency of 70MHz common to the PEs, determined based on the workload of the most heavily loaded PE ("PE1") at some point in time prior to the "new task" creation. Ex. 1009, 4:54-60, Figure 3. Then, subsequently, Nicol's Figure 3 also illustrates a maximum frequency of 140MHz (common to all PEs or the group of PEs) that is determined based on the workload of the most heavily loaded PE ("PE2"), taking into account the additional workload due to the "new task." *Id.*, 3:66-4:6.

31. I also disagree with Patent Owner's argument that Nicol "teaches away" from running the PEs at "maximum" frequency. POR, 40. In my opinion, a POSITA would not read Nicol to teach away from running the PEs at a maximum frequency when needed. In fact, Nicol expressly states that in "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." Ex. 1009, 2:65-67.

32. Therefore, it is my opinion that Nicol in view of Kling renders this claim element obvious.

4. "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"

33. In my opinion, as I previously explained, Nicol in view of Kling renders this claim element obvious. Ex. 1001, ¶¶ 180-188.

34. In my opinion, a POSITA would have known that the use of temperature sensors to detect over-temperature was commonly done in computer systems, such as described in Kling. Ex. 1010, 1:38-46 ("Some processor packages include ***a thermal sensor*** to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. ***If these precautions are not taken, the processor may destroy itself by its own heat***."). Indeed, Kling teaches an additional and improved over-temperature protection system that relies on power measurements. In my opinion, Kling teaches that if the measured power exceeds a threshold indicative of an over-temperature condition, an over-heating processor would be throttled to "reduce its core frequency." Ex. 1010, 2:57-51. And, as I previously explained, Nicol teaches "tracking temperature" of a PE (Ex. 1009, 4:11-13) and as such, a POSITA would have understood that the tracked temperature

is advantageously used to determine that the temperature exceeds a critical temperature threshold (over-temperature threshold) and thus immediately place an over-heating PE in a low power mode by reducing its frequency and voltage to avoid catastrophic meltdown, as taught in Kling. *See* Ex. 1001, ¶¶185-188.

35. Therefore, in my opinion, Nicol in view of Kling renders this claim element obvious.

5. "subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with . . . (b) a hysteresis characteristic"

36. In my opinion, as I previously explained, Nicol in view of Kling renders this claim element obvious. Ex. 1001, ¶¶ 189-192. I also note that while Patent Owner notes that Kling only has a "single mention of the term hysteresis" (POR, 41-42), the '605 patent specification fails to describe any implementation of hysteresis or provide any figure illustrating the use of hysteresis, and like Kling, only mentions "hysteresis" exactly once. Ex. 1003, 13:50. In my opinion, this lack of disclosure confirms that this claim element is devoid of any novelty

37. As I previously explained, a POSITA would have known that hysteresis was (and still is) commonly used to avoid excessive switching. Ex. 1001, ¶¶189-190. And, hysteresis is commonly implemented using a simple two-level

threshold as illustrated in Kling Figure 2A. Ex. 1001, ¶190; Ex. 1010, 5:51-54, Figure 2A. Indeed, the Wiley dictionary provides the well-known thermostat example of hysteresis using a two-level threshold. *See* Ex. 1053 (Wiley), 355 ("[A] thermostat utilizes the right amount of hysteresis to maintain the desired temperature. Too little hysteresis would result in constant switching between cycles, and too much would allow too large a temperature interval between cycles.")). In my opinion, just like a thermostat control system and as taught in Kling, an over-temperature protection system requires the use of hysteresis to avoid constantly switching between a normal operation mode (*e.g.,* high frequency and high voltage setting to achieve maximum PE performance) and an over-temperature protection mode (*e.g.,* reduced voltage and frequency to "cool off" an over-heating PE). Ex. 1010, 5:45-54.

38. Therefore, Nicol in view of Kling renders this claim element obvious.

D. Nicol in view of Kling renders obvious Claim 2

39. In my opinion, as I previously explained, Nicol in view of Kling renders obvious Claim 2's requirement that the "clock frequency is determined in part by processor voltage." Ex. 1001, ¶¶ 193-197.

40. For example, Kling teaches reducing clock frequency when a voltage proportional to the power consumed by a processor reaches a critical

threshold. Ex. 1010, 2:31-51 ("the power consumed by at least a portion of a computer system is monitored by measuring a parameter that is approximately proportional to the consumed power …Once the consumed power … reaches a threshold, a throttle signal is sent to one or more ICs of the computer system by the controller. In response to receiving this throttle signal, one or more of the ICs reduces its power consumption. For example, … a processor of the computer system ***reduces its core frequency*** … "), 4:55-58 ("the term '***power***' includes either actual power or a value, such as current, ***voltage***, duty cycle, or other measurement, that is proportional to power.").

41. Patent Owner disagrees, and argues that the consumed power of Kling is reflected in a voltage drop between the power source and the processor (which is allegedly thus not the claimed processor voltage). POR, 46. I disagree. In my opinion, Kling undisputedly teaches reducing clock frequency based on the power consumed by a processor. And, a POSITA would have known that Kling's clock frequency reduction is based in part on processor voltage because the power consumed by a processor is the product of its voltage and current (P= V x I), as was well-known in the art.

42. I also note that Patent Owner admits that the maximum power limit that triggers frequency reduction in Kling is a voltage limit: "Kling is properly

understood to describe a system that sets its own ***voltage limit*** that is reliably below a maximum voltage" as a "self-imposed safety measure." POR, 48; *see also id.* at 49 ("Petitioner points to a static upper ***voltage limit*** as the available power upon which the clock frequency depends …").

43. I also disagree with Patent Owner that Nicol does not teach determining frequency based in part on processor voltage. POR, 44. In my opinion, a POSITA would have understood that Nicol's PE(s) are designed to operate over a range of supply voltages, and therefore the maximum clock frequency of the PE is determined in part by the maximum supply voltage of the chip. *See* Ex. 1009, 2:34-26 ("FIG. 1 illustrates the ***maximum operating frequency … as a function of supply voltage***.").

44. Therefore, in my opinion, Nicol in view of Kling renders this claim obvious.

E. Nicol in view of Kling renders obvious Claim 3

45. As I previously explained, in my opinion Nicol in view of Kling renders obvious Claim 3's requirement that "said clock frequency is determined in part by available power." Ex. 1001, ¶¶ 198-206.

46. Indeed, Kling's upper limit reflects the available power that can be safely consumed by the system, reflecting the maximum frequency (power is proportional to frequency), and in response to hitting such a threshold, Kling throttles the clock frequency. *See* Ex. 1010, 5:21-25, 5:38-40, 2:48-51.

47. Therefore, in my opinion, Nicol in view of Kling renders this claim obvious.

F. Nicol in view of Kling renders obvious Claim 4

48. As I previously explained, in my opinion Nicol in view of Kling renders obvious Claim 4's requirement that "said clock frequency is determined in part by a user setting." Ex. 1001, ¶¶ 207-213.

49. Indeed, Kling discloses that a user can set both the upper and lower thresholds in total power consumed "for all or a portion of the computer system". Ex. 1010, 5:9-12; *see also id.*, 5:16-19 ("This [upper] threshold … may be modifiable by the system designer or system user"), 5:57-59 ("the lower threshold may be modifiable by a user"). And, I understand that Patent Owner at least agrees that "[t]he frequency depends on either ***the upper threshold***…" POR, 49-50. And, as I previously explained, the frequency also depends on the lower threshold, because the lower threshold determines when the throttle signal is deasserted. And, Kling explains "[o]nce the lower threshold is reached, the throttle signal may be deasserted, and, in response, the IC resume normal operation." Ex. 1010, 5:48-51; Ex. 1001, ¶¶210-213; *see also* Ex. 1010, 5:38-40, 2:48-51. .

50. Therefore, in my opinion, Nicol in view of Kling renders this claim obvious.

G. Nicol in view of Kling renders obvious Claim 5

51. As I previously explained, in my opinion Nicol in view of Kling renders obvious Claim 5's requirement that "said clock frequency is determined in part by available power and a user setting." Ex. 1001, ¶¶ 214-216.

52. Indeed, as previously explained, a user can set both the upper and lower thresholds of Kling. Ex. 1010, 5:16-19 ("This [upper] threshold … may be modifiable by the system designer or system user"), 5:57-59 ("the lower threshold may be modifiable by a user"). And, the upper threshold can be set by a user in accordance with the maximum power that can be reliably consumed by the system, such as via a power delivery specification. Ex. 1010, 5:21-25 (explaining that "***the upper limit*** is a constant value … associated with an ***approximately maximum power that can be reliably consumed by the computer system*** …"), Claim 12.

53. Therefore, in my opinion, Nicol in view of Kling renders this claim element obvious.

H. Nicol in view of Kling renders obvious Claim 6

54. As I previously explained, in my opinion Nicol in view of Kling renders obvious Claim 6's requirement that "said number of pending operations is determined at by the fill level of one or more buffers." Ex. 1001, ¶¶ 217-223.

55. Indeed, a POSITA would understand that using the fill level of a buffer was commonly done to determine the number of pending instructions. And, as Nicol already teaches estimating processing loads based on the number of instructions for each of the "to do" tasks to be allocated to the PEs, a POSITA would have found it obvious to utilize the fill level of a buffer to do so. Ex. 1009, 3:13–18 ("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.").

56. Therefore, in my opinion, Nicol in view of Kling renders this claim obvious.

V. CONCLUSION

57. I hereby declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct, and that all statements made of my own knowledge are true and that all statements made on information and belief are believed to be true. I understand that willful false statements are punishable by fine or imprisonment or both.

See 18 U.S.C. § 1001.

Date: 8 March 2021

Respectfully submitted,

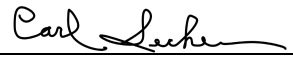

Dr. Carl Sechen