

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 IPR2020-00541

PETITIONER'S REPLY TO PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. GROUND I: NICOL IN VIEW OF KLING RENDERS CLAIMS 1-6 OBVIOUS.....	1
A. A POSITA Would be Motivated to Combine Nicol and Kling	1
B. Nicol’s Embodiments Share Common Functionality	6
C. Nicol in View of Kling Renders Claim 1 Obvious	8
1. “being adapted for sequentially processing data”	8
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”	10
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”	15
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”	20
5. “subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with . . . (b) a hysteresis characteristic”	21
D. Nicol in View of Kling Renders Obvious Dependent Claim 2.....	22
E. Nicol in View of Kling renders Obvious Dependent Claim 3	24
F. Nicol in View of Kling renders Obvious Dependent Claim 4	25
G. Nicol in View of Kling renders Obvious Dependent Claim 5	26
H. Nicol in View of Kling renders Obvious Dependent Claim 6	26
III. GROUNDS II AND III.....	27

IV. CONCLUSION.....	27
---------------------	----

TABLE OF AUTHORITIES

	Page(s)
 Cases	
<i>Bos. Sci. Scimed, Inc. v. Cordis Corp.</i> , 554 F.3d 982 (Fed. Cir. 2009)	7
<i>Elbrus Int’l Ltd. v. Samsung Elecs. Co.</i> , 738 F. App’x 694 (Fed. Cir. 2018)	6
<i>Facebook, Inc. v. Windy City Innovations, LLC</i> , 973 F.3d 1321 (Fed. Cir. 2020)	6
<i>KSR Int’l Co. v. Teleflex Inc.</i> , 550 U.S. 398 (2007).....	2
<i>Livanova, Inc. v. Neuro and Cardiac Techs., LLC</i> , IPR2019-00264, Paper 32 (PTAB May 18, 2020)	20
<i>Polaris Indus., Inc. v. Arctic Cat, Inc.</i> , 882 F.3d 1056 (Fed. Cir 2018)	2
<i>XPP Schweiz AG v. Intel Corp.</i> , Case No. 1:19-cv-1006-UNA (D. Del.).....	vi

PETITIONER’S UPDATED EXHIBIT LIST

Exhibit No.	Description
1001	Declaration of Carl Sechen, Ph.D. (“Ex. 1001”)
1002	Carl Sechen, Ph.D. <i>Curriculum Vitae</i>
1003	U.S. Patent No. 9,075,605 (the “’605 patent”)
1004	Certified File History of U.S. Patent Application No. 13/653,639 (“’605 file history”) (without copies of the prior art references)
1005	U.S. Patent No. 8,099,618 (the “’618 patent”)
1006	U.S. Patent No. 7,444,531 (the “’531 patent”)
1007	U.S. Patent No. 9,552,047 (the “’047 patent”)
1008	Certified File History of U.S. Patent Application No. 14/219,945 (“’047 file history”) (without copies of the prior art references)
1009	U.S. Patent No. 6,141,762 (“Nicol”)
1010	U.S. Patent No. 6,367,023 (“Kling”)
1011	U.S. Patent No. 6,535,798 (“Bhatia”)
1012	U.S. Patent No. 6,052,773 (“DeHon”)
1013	Complaint, <i>XPP Schweiz AG v. Intel Corp.</i> , Case No. 1:19-cv-1006-UNA (D. Del.) (No. 1) (“Complaint”)
1014	Kung, H.T., <i>Why Systolic Architectures?</i> , Computer Vol. 15 (Pub. January 1982)
1015	Reserved
1016	Joint Claim Construction Chart, <i>XPP Schweiz AG v. Intel Corp.</i> , Case No. 1:19-cv-1006-UNA (D. Del.) (No. 67) (“Joint Claim Construction Chart”)
1017	I. Hong, D. Kirovski, Gang Qu, M. Potkonjak and M. B. Srivastava, <i>Power Optimization of Variable-Voltage Core-Based Systems</i> , in 18 IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, NO. 12, 1702-1714 (Dec. 1999) (“Hong”)
1018	U.S. Patent No. 6,704,877 (“Cline”)
1019	RUDOLF F. GRAF, MODERN DICTIONARY OF ELECTRONICS (7th ed. 1999) (“Graf”)
1020	<i>Advanced Configuration and Power Interface (ACPI)</i> , Revision 1.0 (December 22, 1996)

Exhibit No.	Description
1021	Michael Peter Kennedy & Leon O. Chua, <i>Hysteresis in Electronic Circuits: A Circuit Theorist's Perspective</i> , in 19 INT. JOURNAL OF CIRCUIT THEORY AND APPLICATIONS, 471–515 (1991)
1022	John P. Hayes, <i>Computer Architecture and Organization</i> (3rd ed. 1998)
1023	Reserved
1024	Reserved
1025	Declaration of Rachel J. Watters in Support of Public Availability of <i>Hysteresis in Electronic Circuits: A Circuit Theorist's Perspective</i> , in 19 INT. JOURNAL OF CIRCUIT THEORY AND APPLICATIONS, 471–515 (1991) (Kennedy)
1026	U.S. Patent No. 6,397,343 (“Williams”)
1027-1032	Intentionally Left Blank
1033	Power of Attorney, U.S. Patent No. 8,471,593, filed in 2014
1034	Power of Attorney, U.S. Patent No. 7,928,763, filed in 2014
1035	Power of Attorney, U.S. Patent No. 9,552,047, filed in 2014
1036	Power of Attorney, U.S. Patent No. 9,037,807, filed in 2014
1037	Patent Assignment Agreement between PACT XPP Technologies AG & Scientia Sol Mentis AG, executed March 15, 2018
1038	Email chain between K. Bendix and A. Grunberger re acceptance of service, dated February 5-6, 2020
1039	Email chain between L. Tarpley and A. Grunberger re acknowledging receipt of electronic service, dated February 7 and 10, 2020
1040	Email chain between A. Grunberger and K. Bendix re service, dated February 10, 2020
1041	Power of Attorney, U.S. Patent No. 7,928,763 filed in 2020
1042	Affidavit of Raymundo Avila with Nationwide Legal re hand-service of Petitioner Intel Corporation's Petitions and Supporting Documents on PACT counsel on February 10, 2020 and Delivery Receipt
1043	Intel Corporation's Complaint for Declaratory Judgment, Dkt. No. 1, filed April 25, 2019

Exhibit No.	Description
1044	Letter from S. Li to S. Lotfollahi re PACT's claim swapping and narrowing, dated December 13, 2019
1045	Intentionally Left Blank
1046	Screenshot from Public PAIR Correspondence Record, U.S. Patent No. 9,552,047, taken February 10, 2020
1047	Request for Certificate of Correction, U.S. Patent No. 9,075,605, filed December 5, 2018
1048	Declaration of Brandon Brown In Support of Motion to Appear <i>Pro Hac Vice</i> On Behalf of Petitioner Intel Corporation
1049	Declaration of Sarah Mikosz In Support of Motion to Appear <i>Pro Hac Vice</i> On Behalf of Petitioner Intel Corporation
1050	Declaration of Kyle Calhoun In Support of Motion to Appear <i>Pro Hac Vice</i> On Behalf of Petitioner Intel Corporation
1051	Deposition Transcript of William Mangione-Smith, dated February 18, 2021
1052	Deposition Transcript of William Mangione-Smith, dated January 15, 2021
1053	Excerpt from Willey Electrical and Electronics Engineering Dictionary, Steven M. Kaplan (2004)
1054	<i>PACT XPP Schweiz AG v. Intel Corp.</i> , Case No. 1:19-cv-1006-JDW (D. Del.), Claim Construction Order, Dkt. No. 192
1055	Declaration of Carl Sechen in Support of Petitioner's Reply to Patent Owner's Response

I. INTRODUCTION

Faced with a strong unpatentability challenge, Patent Owner (“PO”) resorts to a scattershot rebuttal that relies on mischaracterizations of the asserted art, misapplications of the law, and illusory distinctions between what is claimed in the ’605 patent and what a POSITA would have known in view of the prior art. Ultimately, as described below, these positions are not supported by the ’605 patent, the prior art, or even PO’s own expert.

II. GROUND I: NICOL IN VIEW OF KLING RENDERS CLAIMS 1-6 OBVIOUS

A. A POSITA Would be Motivated to Combine Nicol and Kling

In its Patent Owner Response (“POR”), PO argues that (1) the Petition merely describes benefits without showing actual motivation to combine; (2) the Petition ignores “prioritization” issues that would arise in combining Nicol and Kling; (3) Kling’s disclosure would “dissuade a POSITA from using a thermal sensor”; and (4) the Petition fails to explain how the combined Nicol and Kling system would function. *See* POR, 54-57. For the reasons discussed below, PO’s arguments fail.

As an initial matter, it is facially incorrect that the Petition merely describes the benefits of combining Nicol and Kling without showing actual motivation to combine. Moreover, PO’s argument that “[m]ere benefit without any showing of actual motivation is plainly insufficient” is a *non sequitur* that rests on misstatements

of the law and the facts.¹ POR, 55; *see KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 401 (2007) (“If a person of ordinary skill in the art can implement a predictable variation, and would see the benefit of doing so, § 103 likely bars its patentability.”). The Petition specifically sets forth multiple reasons to combine. Pet., 23-35, 42-45. For example, the Petition explains that Nicol does not discuss over-temperature protection and thus “a POSITA seeking to improve Nicol would have been motivated to consult Kling, which discusses well-known techniques of providing over-temperature protection, as well as additional improvements.” Pet., 24-25. In particular, a POSITA would have been motivated to use temperature measurements both to improve voltage scaling (taught in Nicol) and for over-temperature protection (discussed in Kling). Pet., 42-43; Ex. 1055, ¶7. Likewise, a POSITA would have been motivated to use power measurements both for over-temperature protection (detailed in Kling) and to detect excessive power consumption consistent

¹ PO misquotes *Polaris* as purportedly holding that reliance on benefit evidences impermissible hindsight. POR, 55 (*citing Polaris Indus., Inc. v. Arctic Cat, Inc.*, 882 F.3d 1056, 1068 (Fed. Cir 2018)). *Polaris* in fact held the opposite—that disregarding evidence of benefit or stated preferences invites hindsight. *Polaris*, 882 F.3d at 1068 (internal citations omitted) (“[B]y completely disregarding certain teachings [in the prior art] as ill-defined ‘subjective preferences,’ the Board’s approach invited the ‘distortion caused by hindsight bias’ into the fold.”); *see also id.* at 1069 (“[S]tatements [in the prior art reference] regarding preferences are relevant to a finding regarding whether a skilled artisan would be motivated to combine that reference with another reference.”).

with Nicol’s objective of minimizing power consumption. Pet., 25; Ex. 1055, ¶7. PO and its expert failed to address these expressly stated reasons to combine Nicol and Kling. POR, 54-57.

PO admits, as it must, that Nicol and Kling both teach a power management system in which a controller dynamically adjusts clock frequencies of PEs. POR, 56 (agreeing with Petitioner’s expert that “Kling teaches a controller ... to adjust clock frequenc[ies] of the processors ... based on sensor measurements” and that “Nicol teaches that the controller allocates tasks to the individual processors ...[and] lowers the clock frequency on the chip to as low a level as possible.”); Ex. 1051, 157:13-18 (admitting that Kling “definitely does relate to power management, yes.”). PO nevertheless contends there are “prioritization issues that would arise from these two different methods for changing frequency” and speculates without support that it would be “difficult to implement them both simultaneously.” POR, 56. PO is incorrect. *First*, there are no prioritization issues because Kling’s over-temperature protection is a critical safety mode of operation whereby the controller puts the brakes on (“throttles”) the processors to avoid catastrophic device meltdown due to over-heating. Pet., 16-18 (*citing* Ex. 1010, 2:45-51, 4:50-5:8); Ex. 1055, ¶8. Kling’s emergency braking does not interfere with Nicol’s normal mode of operation as long as there is no risk of device failure due to over-heating (*e.g.*, through excessive power consumption), and Kling’s features rightfully take over when there

is such a risk because “the processor may destroy itself by its own heat.” Ex. 1010, 1:44-46; Ex. 1055, ¶8.

Second, there is no difficulty implementing normal power management modes of operation (as taught in Nicol) alongside over-temperature protection modes (as taught in Kling). Pet., 23-25. PO fails to recognize that Nicol and Kling teach power management features routinely implemented together in computer systems. Pet., 23-25; Ex. 1055, ¶9; Ex. 1010, 1:30-46 (“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 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.*”).²

PO’s other argument that “the disclosure of Kling would dissuade a POSITA from using a thermal sensor as proposed by Petitioner” (POR, 56) fares no better.

² All emphasis added unless noted otherwise.

As discussed in the Petition, Nicol expressly discusses “tracking temperature”, requiring the use of temperature sensors. Pet., 42-43 (citing Ex. 1009, 4:11-13). A POSITA would have understood the obvious advantage of using those temperature sensors to also detect over-temperature conditions, as was commonly done in prior art computer systems. Pet., 43; Ex. 1055, ¶¶10-11; Ex. 1010, 1:38-46. In fact, PO’s expert testified in a related proceeding that Nicol explicitly contemplates the need for (and ability to address) thermal management, noting that Nicol is “saying that thermal management may be a concern that you need to address and that perhaps you can address it” Ex. 1052, 134:16-135:9. While Kling teaches an additional over-temperature protection mechanism using direct power measurement as a trip point, a POSITA would have understood that it does not teach eliminating the conventional trip points based on temperature measurements (however imperfect) using temperature sensors. Pet., 43; Ex. 1055, ¶10. A POSITA would have understood that temperature sensors are necessary to quickly identify local “hot spots” that power measurements alone could miss. Ex. 1055, ¶10. This is because the same 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 to rise to an unsafe level if most of the power is consumed by one small circuit of the chip, creating a dangerous “hot spot.” *Id.*, ¶11.

PO's response further hinges on its misguided contention that the Petition must "explain how a system combining Nicol and Kling would function." POR, 55. This is irrelevant, and PO fails to cite any authority for the proposition that Petitioner must show *how* a combined system *would function* for a motivation to combine analysis. In fact, it is settled law that no such showing is required for the obviousness analysis. *See, e.g., Facebook, Inc. v. Windy City Innovations, LLC*, 973 F.3d 1321, 1343 (Fed. Cir. 2020) (rejecting patentee's argument "that the Board's analysis was infected by hindsight and that it did not adequately explain *how* someone skilled in the art would build the combined Roseman and Rissanen system.") (emphasis in original); *Elbrus Int'l Ltd. v. Samsung Elecs. Co.*, 738 F. App'x 694, 698–99 (Fed. Cir. 2018) ("Therefore, Elbrus's argument that combining Sukegawa and Lu could lead to an unworkable circuit is basically irrelevant.") (citations and quotations omitted). In any event, as discussed above and explained in the Petition, a POSITA would have found it obvious to program Nicol's controller to handle the necessary over-temperature protection features taught by Kling, yielding predictable results as the implementation is primarily add-on software. Pet., 24-25, 42-45; Ex. 1001, ¶129; Ex. 1055, ¶12.

B. Nicol's Embodiments Share Common Functionality

PO repeatedly contends that Petitioner failed to show reasons to combine embodiments of Nicol in support of its POR arguments. *See* POR, 17-18, 28-30, 34.

PO is wrong in law and in the facts. “Combining two embodiments disclosed adjacent to each other in a prior art patent does not require a leap of inventiveness.” *Bos. Sci. Scimed, Inc. v. Cordis Corp.*, 554 F.3d 982, 991 (Fed. Cir. 2009). Here, Nicol’s embodiments at issue, illustrated for example in Figures 2 and 4, are adjacent embodiments. In any event, Petitioner relies on the core function of Nicol’s real-time operating system (“RTOS”) software, which cuts across all embodiments. The core function of the RTOS is to allocate tasks among PEs and control their frequency and voltage settings accordingly, regardless 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). Pet., 12-16 (*citing* Ex. 1009, 2:18-27, 6:27-33); Ex. 1055, ¶13. For example, in Nicol’s Figure 4 embodiment, in which each PE has its own DC-DC converter and its own PLL, the RTOS has the utmost flexibility and can independently set the clock frequency and supply voltage of each PE. Ex. 1055, ¶14. But the RTOS can also set the same clock frequency and supply voltage for PEs when tasks are evenly distributed among them, in which case the PEs operate in unison exactly as in the Figure 2 embodiment. Ex. 1009, 5:53-65. The RTOS can also group PEs into PE groups such that each group operates under the same voltage and frequency as would be the case under the “middle ground” discussed in Nicol. *Id.*, 6:61-7:3. Notably, PO’s expert concedes that Nicol’s “middle ground” is a

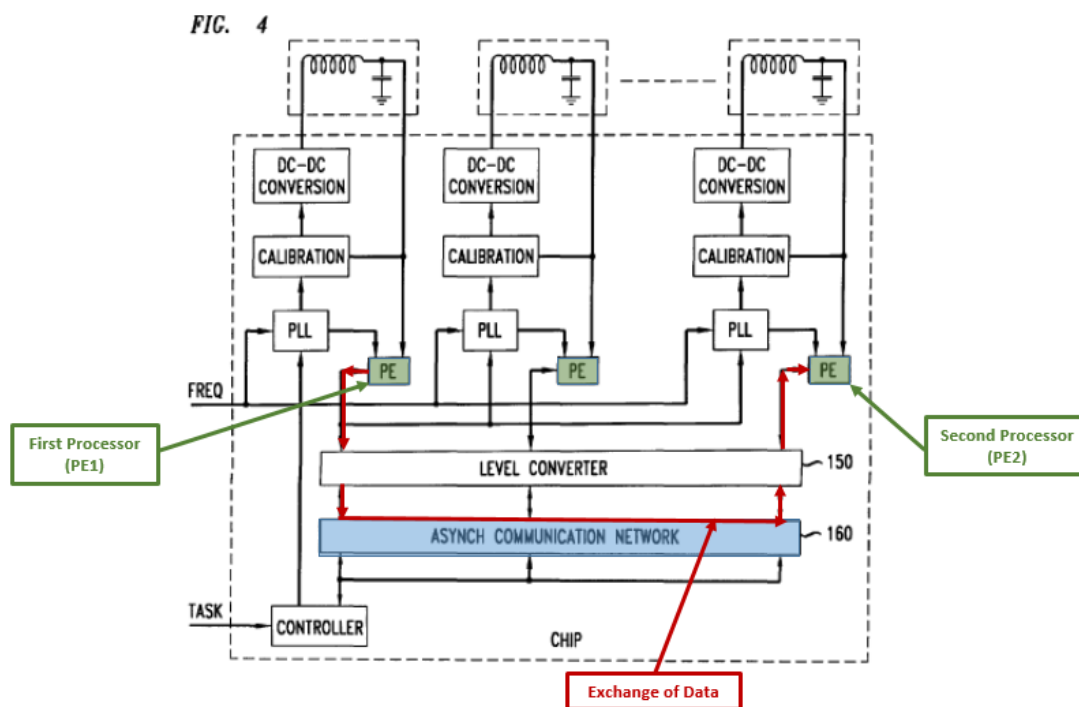
combination of Nicol's Figure 2 and Figure 4 embodiments that, moreover, operates under RTOS control just like other embodiments. Ex. 1051, 60:4-7, 65:12-25 ("whether it's a Figure 2 system in isolation or sort of this middle ground that *combines* elements of Figure 2 with Figure 4, the Figure 2 components, PE 1 and 2, ... will be configured to operate at a clock speed that satisfies the requirements of whichever one is most heavily loaded.").

C. Nicol in View of Kling Renders Claim 1 Obvious

1. "being adapted for sequentially processing data"

The Petition provides three independent reasons for why a POSITA would understand Nicol to disclose, or at least render obvious, "being adapted for sequentially processing data." See Pet., 30-35. While the POR mischaracterizes these independent grounds as being part of a "three step analysis" (POR, 20-21), no such effort is necessary: as explained in the Petition, Nicol discloses that data can be processed by one PE and then sequentially passed to another PE for further processing. Pet., 30-32. Nicol discloses this functionality via its PE-to-PE asynchronous communication network ("ACN") which passes data from one PE to another PE for sequential processing. See *id.*; 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”);
Ex. 1055, ¶16.



Pet., 30-32; Ex. 1009, Figure 4 (annotated). Despite this clear disclosure in Nicol, PO contends that the ACN only provides point-to-point communication between the controller and each PE, and does not necessitate “communication across the PEs” among PEs (*e.g.*, between the green PEs in the figure above). POR, 22-23. PO’s position is eviscerated by the admissions of its expert who admits that Nicol’s ACN “would be used for the processing *PEs to talk to each other*” (Ex. 1052, 107:14-108:1) because “it certainly is possible, and was possible, to put two tasks *with a dependency between them* on two separate processing elements in a multiple PE

system in general.”³ *Id.*, 95:11-15, 60:18-61-9 (“[Q]: You’re saying that in the instances that Nicol described, so you have to take one task and try to break it up into multiple subtasks to divide across multiple PEs, there’s almost always going to have to be communication *between* those individual PEs about what [tasks] they’re working on; right? A. *Yeah....*”); *id.*, 107:14-108:1; *see also* Ex. 1055, ¶18. PO’s expert further concedes that “a person of ordinary skill in the art would see that *Nicol anticipates that figure 4 could be used to build a system where the PEs communicate and envisions that as a possibility.*” *Id.*, 110:21-25.

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”


PO’s 5-page argument directed to this element of Claim 1 is deeply flawed, relying on misrepresentations and taking positions contradicted by its own expert. *First*, PO misrepresents that Petitioner improperly points to features from different embodiments without explaining how they are interchangeable or fit together. POR,

³ The Petition explained that Nicol’s PEs can also sequentially process data in a manner similar to the known systolic architecture. Pet., 34-35. Given the admissions of PO’s expert, further discussion of systolic architecture is not necessary. *See* Ex. 1052, 60:11-61:9, 95:11-15, 107:14-108:1, 110:21-25.

28-30. Not so. The Petition in fact makes clear that both Figures 2 and 4 embodiments of Nicol satisfy claim 1. Pet., 28 (discussing “Nicol’s multiprocessor chip, illustrated in both Figure 4 [] and Figure 2”). And, contrary to PO’s assertion, the Petition also relies on the Figure 4 embodiment to satisfy this claim element. Pet., 35-36 (“Nicol explains that the controller (red) in its multiprocessor system”) (citing Ex. 1009, Figure 4). More fundamentally, the Petition relies on Nicol’s core teaching, illustrated in Figure 3, that the clock frequency of one or more PEs is reduced when a “reduced load permits lowering the clock frequency.” Ex. 1009, 4:54-66; Pet. 37-38. This core teaching cuts across all embodiments and is not limited to the Figure 2 embodiment, as PO incorrectly asserts. POR, 29-30; *see also supra* Section II.B.

As detailed in the Petition, Nicol’s Figure 3 illustrates setting the frequency of certain PEs (*e.g.*, PEs of Figure 4) initially to 70MHz. Pet., 37-38. This would be the case when the controller evenly distributes tasks among available PEs. Pet. 36-37. In this case, the PEs of Figure 4 operate in unison exactly the same way as in the Figure 2 embodiment. Ex. 1055, ¶20. Indeed, 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, but emphasizes that “chip-wise voltage scaling”—whereby the same voltage and frequency is applied to all PEs—is 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, the initial 70MHz frequency setting is the *minimum* frequency needed to handle the tasks at hand by the PEs, which takes into account the number of pending operations of the PE identified in the Petition as the “first processor” (PE1). Pet., 36-37; 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. 1055, ¶21. PO does not dispute this, and PO’s only argument is the misguided argument that the voltage/frequency scaling illustrated in Figure 3 does not apply to Nicol’s Figure 4 embodiment. POR, 29.



Ex. 1009, Figure 3 (annotated).

The Petition also explains that Nicol’s clock gating mechanism provides another independent way of satisfying this claim element.⁴ Pet., 37-38. Under this mechanism, the controller in Nicol 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. Pet., 37-38; Ex. 1009, 5:42-44 (“the

⁴ PO argues that the Petition fails to provide motivation to combine these different ways of satisfying the claimed “minimum” frequency. POR, 33-34. No such motivation is needed because Nicol teaches using them all. Pet., 37-39.


operating system employs gated clocking techniques to ‘shut down’ PEs that are not needed.”); *see also* Ex. 1051, 52:5-17.

PO and its expert incorrectly contend that clock-gating is instead used in Nicol to avoid circuit failure. POR, 31-32. In doing so, PO and its expert misrepresent Nicol’s teaching. The language PO block quotes from Nicol (“ V_{dd} -local can only be reduced so-far before the circuits start to fail”) simply states that PEs are designed to operate over a certain V_{dd} range and cannot operate below a minimum design voltage (in Figure 1, that minimum voltage is approximately 1.1V). Ex. 1009, 5:42-45, Figure 1; Ex. 1055, ¶23. Thus, when a PE is idling, reducing its V_{dd} supply voltage below its minimum voltage to save power is not appropriate as the circuit would fail and, instead, the PE should be clock-gated and its voltage maintained at minimum voltage or higher. Ex. 1055, ¶23. Indeed, clock gating was (and still is) commonly used in computer systems to temporarily shut down circuits that are not needed to save power. Pet., 38; Ex. 1055, ¶22; 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.”).⁵

⁵ The Petition also discusses the use of the ACPI standards’ C3 state as a third independent way of satisfying the claimed “minimum.” Pet., 38-39.

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”

As detailed in the Petition, 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 in Figure 3, to handle additional workload. Pet., 39-41. The higher frequency can be the maximum frequency (annotated below) that the PEs are designed to operate at. Pet., 41; 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. 1055, ¶26.

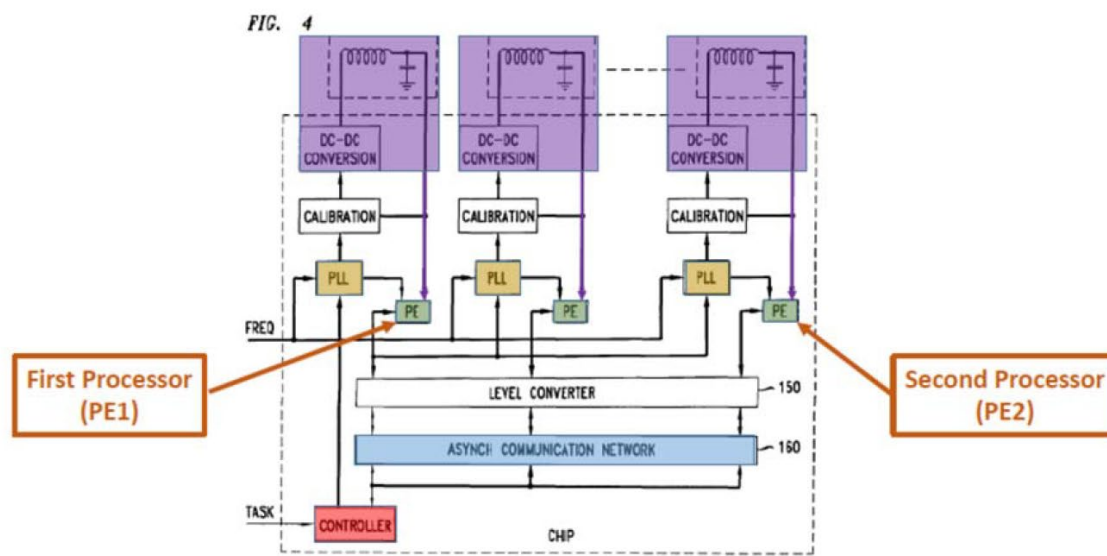


Ex. 1009, Figure 3 (annotated).

PO's 7-page discussion of this claim element is replete with fundamental errors. *First*, as discussed above, PO erroneously contends that Nicol's dynamic voltage and frequency scaling (DVFS) illustrated in Figure 3 is somehow inapplicable to the hardware configuration in Figure 4. POR, 34. To the contrary, a POSITA would have understood that DVFS is the core teaching of Nicol that applies to all hardware configurations, including the Figure 2 configuration (all PEs share the same DC-DC converter and PLL), Figure 4 configuration (each PE has its own DC-DC converter and PLL) and the "middle ground" configuration (each group

of PEs shares a DC-DC converter and PLL). *See supra* Section II.B; Ex. 1055, ¶27. Relatedly, PO erroneously contends that Figure 3 illustrates DVFS applied to a single PE. POR, 38-39. Not so. A POSITA would have understood that Figure 3 illustrates DVFS applied to one PE, a group of PEs, or all PEs in unison. Ex. 1055, ¶27; *see also* Ex. 1009, 2:40-43, 4:54-64.

Second, PO erroneously contends that the controller in Figure 4 can only rigidly set the clock frequency of each PE according to its own workload. POR, 35. Nicol makes clear that Figure 4 illustrates the most flexible hardware configuration whereby the controller *can* independently adjust the frequency and voltage of each PE when tasks cannot be evenly divided among PEs, but is not so limited. 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.”); Ex. 1055, ¶28. When tasks can be evenly divided, the controller adjusts their frequency and supply voltage in unison (“chip-wise” scaling), taking into account all of the PE workloads. Ex. 1009, 5:56-59. Indeed, PO’s expert admits that “the load on the other PEs might be taken into account in allocating the [] newly introduced tasks.” Ex. 1051, 61:21-23. In particular, as explained in the Petition and illustrated below, the controller sets the maximum clock frequency of PE1, taking into account the workload (*e.g.*, number of pending operations) of all PEs, including the PE mapped to the claimed second processor (PE2). Pet., 39-40.



Ex. 1009, Figure 4 (annotated).

Third, PO appears to contend that the “most heavily loaded” processor cannot be both PE1 and PE2. POR, 37; Ex. 1055, ¶29. But that would in fact be the case when tasks are evenly divided among available PEs. Those available PEs (including PE1 and PE2) would be equally loaded, and would be the “most heavily loaded” compared, for example, to clock-gated PEs with no load. *Id.*

The above discussion uses the Figure 4 hardware configuration (each PE has its own DC-DC converter and PLL) as an example. The same applies to both the Figure 2 and “middle ground” embodiments where all PEs of the chip (Figure 2) or PEs within a group (middle ground) share the same DC-DC converter and PLL. Ex. 1009, 6:66-7:3. Here, all PEs on the chip (Figure 2), or PEs within the same group (middle ground) operate in unison at the same clock frequency and supply voltage.

Ex. 1009, 3:9-42; Ex. 1055, ¶30. In those cases, 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 a point in time prior to the “new task” creation. Ex. 1009, 4:54-60, Figure 3; Ex. 1055, ¶30. 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”) at a later point in time, taking into account the additional workload due to the “new task.” Ex. 1009, 3:66-4:6. Indeed, PO’s expert admits that in the middle ground embodiment, “if there’s a PE that is not the most heavily loaded one, its operating state will be set based on whatever the load is of the most heavily loaded PE.” Ex. 1051, 69:13-16; *see also id.*, 69:18-25 (admitting group of PEs in middle ground operates the same as in Figure 2).

PO also argues that Nicol “teaches away” from running the PEs at “maximum” frequency. POR, 40. This is incorrect: Nicol expressly states that PEs may run at “full speed”: 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; Ex. 1055, ¶31. PO’s expert also admits that “...a particular load situation may cause [] all [PEs] to be running at their highest possible clock frequency and voltage.” Ex. 1051, 54:2-14. In any event, PO fails to support its allegation of “teach[ing] away” because PO has not provided any

evidence that Nicol discredits or criticizes running PEs at their maximum frequency. *Livanova, Inc. v. Neuro and Cardiac Techs., LLC*, IPR2019-00264, Paper 32, at 38-39 (PTAB May 18, 2020).

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”

PO contends that this claim element is not met because there is no reason to combine Nicol with Kling, and because a reduction in clock frequency to avoid overheating allegedly would prevent tasks from being completed on time. POR, 40-41. PO’s argument is without merit as detailed above in Section II.A. Indeed, the use of temperature sensors to detect over-temperature was commonly done in computer systems, as described in Kling. Pet., 42-43; 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.*”). Nicol teaches “tracking temperature” of a PE (Ex. 1009, 4:11-13) and a POSITA would have understood that the tracked temperature is advantageously used to determine that if the temperature exceeds a critical temperature threshold, the controller should place an over-heating PE in a low power mode, reducing its frequency and voltage to avoid catastrophic meltdown, as taught in Kling. Pet. 42-43; Ex. 1055, ¶34. Indeed, PO’s expert

admitted that a POSITA “concerned about overheating” Nicol’s system would “know of ways that he could do temperature management using a temperature sensor” as was done for example in Intel Pentium processors. Ex. 1052, 129:23-130:9.

Kling in fact teaches an improved over-temperature protection system that further relies on power measurements. Ex. 1055, ¶34. If the measured power exceeds a threshold indicative of an over-temperature condition, an over-heating processor is “throttled” to “reduce its core frequency.” Ex. 1010, 2:47-51; Pet., 17-18 (*citing* Ex. 1010, 2:47-51, 3:28-29, 4:50-5:8).

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

At the outset, PO’s repeated—and false—argument that Petitioner relies on nothing more than Kling’s “single mention of the term hysteresis” (POR, 41-42) is ironic given that the ’605 Patent specification mentions “hysteresis” exactly once. Ex. 1003, 13:50. This total lack of disclosure confirms this claim element is devoid of any novelty. Ex. 1055, ¶36.

Indeed, as explained in the Petition, hysteresis was (and still is) commonly used to avoid excessive switching; moreover, hysteresis is commonly implemented using a simple two-level threshold as illustrated in Kling’s Figure 2A. Pet., 44-45; Ex. 1010, 5:51-54, Figure 2A; Ex. 1055, ¶37. PO’s expert confirms this, admitting

the common use of “a two-level threshold band such as in a thermostat” is “a good example of hysteresis writ large.” Ex. 1051, 138:2-12; *see also* Ex. 1055, ¶37 (*citing* 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.”)). This admission is fatal to PO’s strained position. Just like thermostat control and as taught in Kling, an over-temperature protection system requires the use of hysteresis to avoid constantly switching between normal operation mode (*e.g.*, high frequency/voltage for maximum performance) and an over-temperature protection mode (*e.g.*, reduced frequency/voltage to “cool off” an over-heating PE). Pet., 44-45; Ex. 1055, ¶37; Ex. 1010, 5:45-54; Ex. 1021, 471.

D. Nicol in View of Kling Renders Obvious Dependent Claim 2

As detailed in the Petition, Nicol in view of Kling renders obvious claim 2’s requirement that the “clock frequency is determined in part by processor voltage.” Pet., 45-46. Kling teaches reducing clock frequency when a voltage representing the power consumed by a processor reaches a threshold. *Id.*; 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 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.”).

PO largely does not dispute this and only quibbles whether the consumed power is reflected in a voltage drop between the power source and the processor (in which case PO contends that the voltage drop is not the claimed processor voltage), or between the processor and ground (in which case PO admits it is the claimed processor voltage). POR, 46. This is merely an unsupported box drawing exercise and misses the point. Kling undisputedly teaches reducing clock frequency based on the power consumed by a processor; the frequency reduction is based in part on processor voltage because the power consumed by a processor is the product of its voltage and current ($P = V \times I$), as was well-known in the art. Ex. 1055, ¶¶41-42. Indeed, PO admits that the maximum power limit that triggers frequency reduction in Kling is a voltage limit. According to PO, “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; *id.*, 49 (“Petitioner points to a static upper *voltage limit* as the available power upon which the clock frequency depends ...”).

PO also wrongly contends that Nicol does not teach determining frequency based in part on processor voltage. POR, 44. A POSITA would have understood that a PE is designed to operate over a range of supply voltages, and that its maximum clock frequency is determined in part by its maximum supply voltage. Ex. 1055, ¶43; Ex. 1009, 2:34-36 (“FIG. 1 illustrates the *maximum operating frequency ... as a function of supply voltage.*”).

E. Nicol in View of Kling renders Obvious Dependent Claim 3

PO’s argument that Nicol in view of Kling does not render obvious “said clock frequency is determined in part by available power” is baseless. POR, 47. PO does not deny that Kling teaches an available power limit that, when exceeded, would trigger frequency reduction. POR, 49; *see also* Ex. 1055, ¶46. Nothing more is required to satisfy this claim element.⁶

PO’s argument that Kling’s available power limit is a “self-imposed safety measure” (*id.*, 48-49) is inapposite and undermined by the admission of its expert, who admits that Kling teaches reducing frequency when an available power limit is

⁶ PO argues that implementing Kling’s available power limit in Nicol’s power management system is a “complex task.” POR, 49. This argument is without support (PO’s expert parrots the same language) and is addressed in Section II.A.

reached. Ex. 1051, 149:12-150:2 (“**Q:** ... Figure 2A shows that Kling implements an upper limit ... of power consumption as part of its disclosure; right? **A:** Yes. **Q:** [W]hat it says is ‘*when we get close to the upper limit threshold, we should scale back frequency so that we don’t use more power than [] we are allowed to use*’; is that right? **A:** ... yes”), 147:13-23.

PO also argues that this limitation is not met because Kling’s available power limit is allegedly “static ... [and] cannot determine *all of these* frequencies” POR, 48-49. Claim 3 does not require determining “all” frequencies based on available power.

F. Nicol in View of Kling renders Obvious Dependent Claim 4

PO’s arguments that Nicol in view of Kling does not render obvious “said clock frequency is determined in part by a user setting” is without merit and eviscerated by the admission of its expert. PO admits that the clock frequency in Kling “depends on [] the upper threshold, which causes throttling...”. POR, 49-50. Thresholds are determined by user settings, as detailed in the Petition and admitted by PO’s expert. Pet., 49-50; Ex. 1055, ¶49; Ex. 1051, 152:17-22 (“**Q:** ...in the Kling reference, *the upper and lower threshold [] may both be set by a user setting*; right? **A:** *That’s my understanding from what they disclosed. Yes.*”), 153:13-20; 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”).

G. Nicol in View of Kling renders Obvious Dependent Claim 5

PO's argument that Nicol in view of Kling does not render obvious "said clock frequency is determined in part by available power and a user setting" is without merit, as detailed above under Sections II.E and II.F. Kling teaches reducing clock frequency when an available power limit is reached, and the limit is determined by user settings. Pet., 49-50; Ex. 1010, 2:31-51, 5:16-19; Ex. 1055, ¶¶51-53. Indeed, PO's expert admitted "there is certainly a user-specified setting for the upper limit ... [which is] simply related to the maximum power that could be consumed by the power supply inside the device" in Kling. Ex. 1051, 155:12-24; *see also id.*, 156:2-8.

H. Nicol in View of Kling renders Obvious Dependent Claim 6

PO argues that Nicol in view of Kling does not render obvious "said number of pending operations is determined at by the fill level of one or more buffers." POR, 52-53. This argument is without merit and contrary to PO's representation to the Patent Office during prosecution of the '605 Patent.

PO does not dispute that Nicol discloses using "the number of instructions" waiting to be executed (number of pending operations) in its task allocation. POR, 52-53. PO also does not dispute that it was well-known in the art to determine the number of pending operations by monitoring buffer fill levels, as confirmed by Cline and Hong. *Id.*, 53. Nothing more is required to render obvious this claim element,

as a POSITA would have found it obvious to utilize the fill level of a buffer and PO's argument that that determination is not "*necessarily*" through buffer fill levels (POR, 53) is inapposite. *See also* Ex. 1055, ¶55. PO's argument is also contrary to PO's representation to the Patent Office during prosecution of the '605 Patent. Pet., 51-52 (*citing* Ex. 1004, 377 ("one of ordinary skill in the art would understand that the fill level of a data buffer [is] representative of the number of waiting operations or processor load.")).

III. GROUNDS II AND III

To narrow the issues before the Board, and in light of the District Court's finding that "data processing unit" need not be reconfigurable, and that no construction is necessary for "minimum",⁷ Petitioner hereby withdraws Grounds II and III. The claims associated with Grounds II and III are obvious for the reasons stated in Ground I.

IV. CONCLUSION

For the reasons explained above, the Board should find that claims 1–6 are obvious in view of Nicol in combination with Kling.

⁷ Ex. 1054, 2.

Date: March 8, 2021

Respectfully Submitted,

/s/ Robert A. Appleby
Kevin Bendix (Reg. No. 67,164)
kevin.bendix@kirkland.com
KIRKLAND & ELLIS LLP
555 South Flower Street, Suite 3700
Los Angeles, CA 90071
Telephone: (213) 680-8400
Facsimile: (213) 680-8500

Gregory S. Arovas, P.C. (Reg. No. 38,818)
garovas@kirkland.com
Robert A. Appleby, P.C. (Reg. No. 40,897)
robert.appleby@kirkland.com
KIRKLAND & ELLIS LLP
601 Lexington Avenue
New York, New York 10022
Telephone: (212) 446-4800
Facsimile: (212) 446-4900

CERTIFICATE OF COMPLIANCE

This Petition complies with the type-volume limitations as mandated in 37 C.F.R. § 42.24, totaling 5,544 words. Counsel has relied upon the word count feature provided by Microsoft Word.

/s/ Robert A. Appleby
Robert A. Appleby (Reg. No. 40,897)

CERTIFICATE OF SERVICE

The undersigned hereby certifies that a copy of the foregoing document with accompanying Exhibits 1051-1055, were served on March 8, 2021 through the Patent Review Processing System, as well as by e-mailing copies to:

Ziyong Li, Esquire
QUINN EMANUEL URQUHART & SULLIVAN, LLP
50 California Street, 22nd Floor
San Francisco, CA 94111
(415) 875-6373
seanli@quinnemanuel.com

Nima Hefazi, Esquire
Joseph M. Paunovich, Esquire
QUINN, EMANUEL, URQUHART & SULLIVAN, LLP
865 S. Figueroa Street, 10th Floor
Los Angeles, CA 90017
(213) 443-3000
nimahefazi@quinnemanuel.com
joepaunovich@quinnemanuel.com

Attorneys for Patent Owner

/s/ Robert A. Appleby
Robert A. Appleby