Trials@uspto.gov
571-272-7822

Paper 35
Entered: September 8, 2021

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,
Petitioner,

v.

PACT XPP SCHWEIZ AG,
Patent Owner.

IPR2020-00539
Patent 9,552,047 B2

Before SALLY C. MEDLEY, SCOTT C. MOORE, and CHRISTOPHER L. OGDEN, *Administrative Patent Judges.*

OGDEN, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision
Determining Some Challenged Claims Unpatentable
35 U.S.C. § 318(a)

I. INTRODUCTION

In response to a Petition (Paper 2, "Pet.") by Petitioner Intel Corporation ("Intel"), the Board instituted an *inter partes* review of claims 1–24, 27, 28, 30, and 33 of U.S. Patent No. 9,552,047 B2 (Ex. 1003, "the '047 patent"). Paper 11. Patent Owner PACT XPP Schweiz AG ("PACT") filed a Patent Owner Response (Paper 20, "PO Resp."), Intel filed a Reply to the Patent Owner Response (Paper 26, "Pet. Reply"), and PACT filed a Sur-reply (Paper 30, "PO Sur-reply").

We held an oral hearing on June 15, 2021, and the transcript is entered on the record. Paper 34 ("Tr.").

This is a Final Written Decision under 35 U.S.C. § 318(a) as to whether the claims challenged in the *inter partes* review are patentable. By a preponderance of the evidence on the record before us, Intel has shown that claims 5, 6, 8, 12, 14–24, 28, 30, and 33 of the '047 patent are unpatentable, but has not shown that claims 2–4, 7, 9–11, 13, or 27 are unpatentable on the asserted grounds.

II. BACKGROUND

A. RELATED PROCEEDING

As related matters, the parties identify *PACT XPP Schweiz AG v. Intel Corp.*, No. 1:19-cv-00267-RGA (D. Del. terminated Apr. 24, 2019); *PACT XPP Schweiz AG v. Intel Corp.*, No. 6:19-cv-00273-ADA (W.D. Tex. terminated May 30, 2019); *Intel Corp. v. PACT XPP Schweiz AG*, No. 3:19-cv-02241 (N.D. Cal. terminated May 30, 2019); and *PACT XPP Schweiz AG v. Intel Corp.*, No. 1:19-cv-1006-JDW (D. Del. filed May 30, 2019) ("the co-pending litigation"). Paper 28, 2; Paper 3, 1–2.

B. THE '047 PATENT (EX. 1003)

The '047 patent describes "[a] multiprocessor that . . . provides for adjusting the clock frequency for at least some data processing units at runtime." Ex. 1003, code (57). The system includes a data processing unit ("VPU") comprising an arbitrary number of processing array elements ("PAEs," or "cells") arranged in a unidimensional or multidimensional matrix, and "connected to one another either directly or via a bus system." *Id.* at 1:38–39, 2:4–12.

The architecture of this system is reconfigurable. *See, e.g.*, Ex. 1003, 1:31–32, 5:10–14, 9:13–17. The system uses "[a] configuration unit (CT) which affects the interconnection and function of the [matrix of PAEs] through configuration . . . assigned to the entire matrix or parts thereof." *Id.* at 2:14–16. The VPU's configuration is stored in the form of "configuration words" stored in configuration registers, of which each PAE can have more than one to determine, for example, its interconnection, its "clock determination," and its internal function. *See id.* at 2:16–23.

As part of the VPU's configuration states, different PAEs can use different clocking speeds. *See* Ex. 1003, 2:53–54. For example, if the VPU is used to execute a complex operation in which individual cells perform sub-operations, the '047 patent states that "it is preferred if these cells are operated at different clock rates, namely in such a way that the cells for less complex operations, thus operations which run in fewer clock cycles, are clocked slower than the other cells." *Id.* at 4:20–29.

According to the '047 patent, "operat[ing a PAE] at a higher clock frequency requires more power," which substantially increases the cooling requirements. Ex. 1003, 2:24–29. "To determine the clock frequency for a

microprocessor based on the state is known," but "problems arise in the overall speed with which certain applications are carried out." *Id.* at 2:30–34. For the VPU to achieve optimal data processing results for speed and power consumption, the '047 patent suggests "that clocking takes place depending on the state, which means that no clock is preselected jointly for all cells based on a certain state, but rather an appropriate clock is assigned to each cell based on the state." *Id.* at 2:60–63.

As particular examples of factors that may be used to determine the operating frequency of each PAE, the '047 patent lists the condition of a power source, a maximum operating frequency, a temperature measurement, buffer filling levels, battery charge state, supply voltage, prioritization of tasks or computations, a user indication, or a hysteresis characteristic. Ex. 1003, 4:20–67, 5:29–57, 7:21–8:7, 13:37–14:16.

The '047 patent claims a priority date of March 5, 2001, which Intel does not contest in the Petition. Pet. 9; Ex. 1003, code (30), 1:16–20. PACT does not challenge March 5, 2001 as the date of invention for the challenged claims. *See, e.g.*, PO Resp. 3–9 (discussing the '047 patent and asserted prior art, without asserting an earlier conception and reduction to practice).

C. CHALLENGED CLAIM AND ASSERTED GROUNDS OF UNPATENTABILITY

Independent claim 1 is representative of the challenged claims:

> 1. A multi-processor system comprising:
>
> a plurality of data processing units, each of the data processing units:
>
> (a) being programmable; and

> (b) being adaptable for sequentially processing data in a clocked manner; and including
>
> (c1) at least one Arithmetic Logic Unit (ALU), and
>
> (c2) at least one data register set for storing intermediate results in sequential data processing;
>
> at least one bus system for at least interconnecting at least some of the data processing units, each being provided a supply voltage by a voltage supply;
>
> wherein for at least some of the data processing units, the clock frequency is adjustable at runtime without affecting the clock frequency of at least one of: one other of the data processing units, and the bus system; and
>
> wherein the voltage supply is adapted to supply higher supply voltages for data processing at higher clock frequencies.

Ex. 1003, 15:38–57. Claim 19 is also independent, and is essentially identical to claim 1 except that it is directed to a mobile device containing the multiprocessor system. *See id.* at 16:60–17:13. The remaining claims 2–18 and 20–24, 27, 28, 30, and 33 depend directly or indirectly from either claims 1 or 19. *See id.* at 15:58–16:59, 17:14–26, 17:50–56, 17:60–63, 18:11–12.

Intel argues four grounds for *inter partes* review, as summarized in the following table:

Ground	Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
1	1–7, 17–20, 22, 23, 27, 33	103(a)[1]	Nicol[2]
2	8–16, 21, 24, 28, 30	103(a)	Nicol, Bhatia[3]
3	1–7, 17–20, 22, 23, 27, 33	103(a)	Nicol, DeHon[4]
4	8–16, 21, 24, 28, 30	103(a)	Nicol, Bhatia, DeHon

Pet. 7. For Ground 1, Intel alleges that claims 1–7, 17–20, 22, 23, 27, 33 are unpatentable under § 103 as obvious over Nicol. *Id.* For Ground 2, Intel alleges that dependent claims 8–16, 21, 24, 28, 30 are unpatentable under § 103 as obvious over Nicol in view of Bhatia. *Id.*

Grounds 3 and 4 mirror Grounds 1 and 2, respectively, except that Intel has added DeHon as an additional reference. Intel provided these two sets of grounds alternatively, based on each of the parties' differing proposed claim constructions in the co-pending litigation. *See* Pet. 13. But in its Reply, Intel agreed to withdraw Grounds 3 and 4, and rely solely on its arguments in Grounds 1 and 2. *See* Pet. Reply 24–25.

D. EXPERT TESTIMONY

For its Petition and Reply, Intel relies on two declarations of Dr. Carl Sechen. Exs. 1001, 1055. For its Patent Owner Response and Sur-reply,

---

[1] 35 U.S.C. § 103(a) (2006), *amended by* Leahy–Smith America Invents Act, Pub. L. No. 112-29 § 103, sec. (n)(1), 125 Stat. 284, 287, 293 (2011) (effective Mar. 16, 2013). This version of § 103 applies because the effective priority date of the '049 patent is before the effective date of the AIA amendments. *See supra* part II.B.

[2] Nicol et al., US 6,141,762 (issued Oct. 31, 2000) (Ex. 1008).

[3] Bhatia et al., US 6,535,798 B1 (issued Mar. 18, 2003) (Ex. 1009).

[4] DeHon et al., US 6,052,773 (issued Apr. 18, 2000) (Ex. 1010).

PACT relies on a declaration by Dr. William Henry Mangione-Smith. Ex. 2027. We find both of them qualified, based on their education and experience, to offer expert testimony on the subject matter of their respective declarations. *See* Fed. R. Evid. 702(a); *see* Ex. 1001 ¶¶ 11–20; Ex. 1002 (Dr. Sechen's curriculum vitae); Ex. 2027 ¶¶ 4–9; Ex. 2028 (Dr. Mangione-Smith's curriculum vitae).

III. GROUNDS OF THE PETITION

For the reasons below, we determine that Intel has shown, by a preponderance of the evidence, that claims 1, 5, 6, 8, 12, 14–24, 28, 30, and 33 of the '047 patent are unpatentable under § 103, but has not shown that claims 2–4, 7, 9–11, 13, or 27 are unpatentable on the grounds asserted in the Petition. Before analyzing the Petition's asserted grounds in detail, we address two matters that will underlie our analysis: the level of ordinary skill in the art, and the construction we will apply to the claim terms.

A. LEVEL OF ORDINARY SKILL IN THE ART

The level of ordinary skill in the pertinent art at the time of the invention is one of the factual considerations relevant to obviousness. *See Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). It is also relevant to how we construe the patent claims. *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1312–13 (Fed. Cir. 2005) (en banc). To assess the level of ordinary skill, we construct a hypothetical "person of ordinary skill in the art," from whose vantage point we assess obviousness and claim interpretation. *See In re Rouffet*, 149 F.3d 1350, 1357 (Fed. Cir. 1998). This legal construct "presumes that all prior art references in the field of the invention are

available to this hypothetical skilled artisan." *Id.* (citing *In re Carlson*, 983 F.2d 1032, 1038 (Fed. Cir. 1993)).

Dr. Sechen opines that "a person of ordinary skill in the art would have been a person with at least a Master's degree in Electrical Engineering (or equivalent experience), and at least three years of experience with processor design." Ex. 1001 ¶ 90. In its Petition, Intel cites Dr. Sechen's testimony, but adds "memory architecture" to the required experience. Pet. 12. PACT and Dr. Mangione-Smith explicitly adopt Dr. Sechen's definition, but misquote it as including a degree in "electrical or computer engineering." *See* PO Resp. 12; Ex. 2027 ¶¶ 16–17.

For this decision, we adopt Dr. Sechen's uncontested declaratory testimony, which we find consistent with the level of technical expertise assumed by the '047 patent disclosure and the asserted prior art references. Intel's addition of "memory architecture" and Dr. Mangione-Smith's addition of "computer engineering" do not materially affect the level of ordinary skill, and would not affect our analysis or the outcome of our decision.

B. Claim Construction

In an *inter partes* review, we construe a patent claim "using the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b)." 37 C.F.R. § 42.100(b) (2020). This includes "construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent." *Id.* The ordinary and customary meaning of a claim term "is its meaning to the ordinary

artisan after reading the entire patent," and "as of the effective filing date of the patent application." *Phillips*, 415 F.3d at 1313, 1321. In construing a claim, we consider "[a]ny prior claim construction determination concerning a term of the claim in a civil action . . . that is timely made of record" in this proceeding. 37 C.F.R. § 42.100(b) (2020).

1. "data processing unit," "data processing unit[s] . . . adaptable for sequentially processing data," "sequentially processing data"

In its Petition, Intel informs us that in the co-pending litigation, the parties had proposed competing constructions for the alternative terms "data processing unit," "data processing unit . . . adaptable for sequentially processing data," or "sequentially processing data." Pet. 12–13 (citing Ex. 1013, 11, 17). A key difference between each party's proposed construction is that Intel's required that the data processing unit be "reconfigurable." *See id.*; Pet. 72–73. Rather than relying solely on its own proposed construction, Intel covered both possibilities in its Petition by presenting two alternative sets of grounds: Grounds 1 and 2 are based on PACT's proposed construction, and Grounds 3 and 4 are based on Intel's proposed construction, which required data processing units to be reconfigurable. Pet. 13.

After Intel filed the Petition, the district court in the co-pending litigation construed the terms "data processing unit" and "data processing units . . . adapted for sequentially processing data" to mean "[s]equential data processors that pass results to one or more additional data processing units to perform additional computations." Ex. 1056 (first alteration in original), 2; Ex. 2029, 17. The district court declined to require, in its

construction, that the data processor be "reconfigurable." Ex. 2029, 19; *see also* Pet. Reply 24–25 (citing Ex. 1056, 2) (acknowledging that the district court's construction "does not require the unit to be reconfigurable").

PACT contends that the district court's construction resolves the claim construction dispute concerning this term. PO Resp. 3. And, to narrow the issues in this *inter partes* review, Intel agreed to withdraw Grounds 3 and 4, and rely on its arguments in Grounds 1 and 2, which did not assume that a data processing unit must be reconfigurable. *See* Pet. Reply 24–25.

Because the meaning of "data processing unit," "data processing unit[s] . . . adapted for sequentially processing data," and "sequentially processing data" are no longer in dispute, we do not need to construe those terms for our decision, so we adopt the uncontested construction from the co-pending litigation. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) ("[W]e need only construe terms 'that are in controversy, and only to the extent necessary to resolve the controversy.'" (quoting *Vivid Techs., Inc. v. Am. Sci & Eng'g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999))).

2. "the clock frequency is adjustable at runtime according to a state"

Claims 2–4 recite the term "the clock frequency is adjustable at runtime according to a state." *See* Ex. 1003, 15:58–67. Although the parties have not proposed explicit constructions for this term, the term is contested, and both sides have made substantial arguments as to its meaning in the briefs and in the oral hearing. *See* Pet. 36–40; PO Resp. 22, 24–25; Pet. Reply 6; PO Sur-reply 4–6; Tr. 12:1–17:2, 35:1–38:8, 40:1–41:1. We must

construe the term because its meaning is dispositive on the issue of whether claims 2–4 would have been obvious in view of the asserted prior art.

Intel identifies a "state" as a processor load, and contends that a clock frequency is "adjustable at runtime according to a state" if, for example, the clock frequency is set to ensure that computing tasks are completed within a required completion time, based on a priority assignment scheme. *See* Pet. 42; Pet. Reply 6. According to Intel, PACT's expert Dr. Mangione-Smith admitted during cross-examination that "[a] person of ordinary skill in the art could look at the load [of a processor] and say, 'That represents a state for some useful reason.'" Pet. Reply 6 (emphasis omitted) (quoting Ex. 1052, 159:13–16); *see also id.* ("[I]n general the clock frequency is at least in part based on the load." (alteration in original) (emphasis omitted) (quoting Ex. 1052, 159:22–24). Intel also argued, at the oral hearing, that the '047 patent explicitly associated "data processing time" with the term "state." Tr. 40:17–41:1 (citing Ex. 1003, 2:56–63 ("In order to achieve optimum data processing results, *in particular with regard to the required data processing time*, as well as the power consumption of the entire data processing unit, it is suggested that clocking takes place depending on the state . . . .") (emphasis added)).

PACT does not dispute that a processor load could be "a state" as recited in claims 2–4. *See* PO Sur-reply 4 ("[O]ne [concept of a load] could be understood as a state of the system."). But PACT distinguishes between an "actual load on the system" and an "anticipated load" or "workload," in which "a clock frequency determination [is] based on future tasks." *Id.* at 4–5; Tr. 35:7–8 ("[W]e think that workload is not reflective of the state of the system.").

More importantly, PACT disagrees with Intel's interpretation of the claim language as a whole.[5] Rather than focusing on whether a load is a "state," PACT focuses on causation: in PACT's view, the clock frequency is not "adjustable at runtime according to a state" if the processor load "is a result, not a cause" of the determined clock frequency. *See* PO Sur-reply 5–6.

We agree with PACT that Intel's analysis fails to consider the language "according to," which by its ordinary meaning implies causation. The plain meaning of claims 2–4, in view of the '047 patent specification as a whole, indicates that the clock frequency is adjusted *after* the multiprocessing system or the data processing unit enters a particular state. As the '047 patent specification expresses, "clocking takes place *depending on* the state," rather than the state depending on the clock frequency. Ex. 1003, 2:60 (emphasis added).

In many embodiments of the '047 patent, the relevant "state" that determines the clocking frequency is called a "configuration state," which is data stored in "configuration registers." *See* Ex. 1003, 2:16–23 ("The configuration of a VPU is determined by writing configuration words into configuration registers . . . . [including] one/or more words for the clock determination . . . ."); 2:64–67 ("[I]t is suggested that the clocking be designed . . . so that one calibration (configuration) mutually influences the

[5] According to PACT, Dr. Mangione-Smith's statement during cross-examination conveyed "the idea that load may be a state for some purposes, but he stated this is not under the claim language." PO Sur-reply 5 (citing Ex. 1052, 156:20–159:3); Tr. 49:17–50:3 (arguing that Dr. Mangione-Smith's statements were answers to hypothetical questions, not posed in the context of the entire claims at issue).

clocking of the total number of cells."); 3:1–6 ("[The clock preselecting means can] select the clocking of the cell based on its configuration as soon as this configuration is determined."); 4:5–16 (describing a method of "operating a field of clocked logic cells which may be set into different configuration states," in which a first clock is "determined dependent on the first state" and a second clock is separately "determined dependent on the second state"); 4:17–19 ("[C]locking may be preselected together with the configuration. The state is then the configuration state and/or is at least determined by it."). In all the above embodiments, the '047 patent contemplates that the "configuration state" is something that the VPU determines, and sets within the configuration registers, before using that state to set the clock frequency.

Besides a configuration state, the '047 patent refers to other factors that the VPU may use to set the clock frequency, and these are likewise states or measured characteristics of the system that exist prior to setting the clock frequency. These include "the condition of a power source," Ex. 1003, 4:59–60, 8:1–7; "a drop in supply voltage," *id.* at 4:61–63, 13:37–50, 14:1–2; "responding to a temperature sensor signal," *id.* at 4:63–67, 7:25–33; responding to "a user indication" through a hardware or software switch, *id.* at 14:4–13; 4:65–66; or responding to the fill level of a software buffer, *id.* at 7:34–67. In these examples, the system adjusts the clock frequency in response to an existing hardware or software condition, and sometimes a sensor is involved in measuring the condition. *See id.* at 13:51–60.

Similarly, the '047 patent describes examples of setting the clock frequency in response to a preexisting property of the PAEs or their configuration. This may include "the maximum operating frequency of the

slowest configuration which results as a function of a PAE configuration or a configuration of a group of PAEs." Ex. 1003, 5:29–57. It may also include a "hysteresis characteristic," which is beneficial "when a temperature-sensitive change of the clock frequencies is to be performed." *Id.* at 14:13–16. During prosecution of the '047 patent, the applicant argued that "the various processors can be separately clocked according to their functional and timing requirements" when, for example, "configurations of processors are limited by their interconnections (space, capacitance and other factors) such that propagation time is affected." Ex. 1004, 551–52.

The '047 patent also discusses the "prioritization of tasks or computations" as a factor that determines the clock frequency. Ex. 1003, 14:2–3; *see also id.* at 4:20–57 (describing the use of different clock rates to perform complex computations). We agree with Intel's observation during the oral hearing that the '047 patent ties together the idea of using different clocking rates with the goal of minimizing data processing time. *See* Tr. 40:17–41:1 (citing Ex. 1003, 2:56–63). But as part of its scheme for prioritizing tasks and computations, the '047 patent still makes use of specific configuration states within the reconfigurable VPU.

For example, if a set of PAEs are working together to perform a complex calculation, a particular cell may be "appropriately clocked down by . . . signals which indicate the state of [a previous] cell participating in a processing sequence." Ex. 1003, 3:45–47. The VPU may, in one embodiment, have a central clock generator that is "set or influenced by a configuration," and a local clock generating unit for individual PAEs that multiplies or divides the central clock frequency by a factor using a "frequency divider ratio" preselected "according to the clock determination

based on the state." *Id.* at 3:48–60. The frequency divider ratio is set based on an input from the PAE's configuration memory. *See id.* at 12:17–23, Fig. 2 (the frequency divider ratio, and thus the local clock frequency, is set based on input 0210a from configuration memory 0208). Thus, the clock frequency of each PAE is set *in response* to a configuration state stored in the system's configuration memory. *See id.* at 12:48–53 ("[A] clock dividing/multiplying signal . . . [is] transmitted to frequency divider/multiplier 0210, so that, *as a response*, 0210 may clock data processing unit 0206 and possibly also other units." (emphasis added)).

In other words, although the '047 patent describes setting the clock frequency based on task prioritization needs or time requirements, it does so by responding to the existing configuration state of the VPU, as stored in configuration registers. *See, e.g.*, Ex. 1003, 5:58–61 (describing "an optimized embodiment" in which "the frequency is adapted only to the configurations which are *currently carried out on the VPU*, in other words, the global frequency may be reset/reconfigured with each configuration." (emphasis added)).

Based on the above evidence in the '047 patent and its prosecution history, we determine that a person of ordinary skill in the art would have understood the term "the clock frequency is adjustable at runtime according to a state" to mean that the clock frequency is adjusted as a result of, or in response to, an existing state of the multiprocessing system or its data processing units.

C. GROUND BASED ON NICOL

Intel alleges that claims 1–7, 17–20, 22, 23, 27, and 33 are unpatentable under 35 U.S.C. § 103 as obvious over Nicol. Pet. 7.

A claim is unpatentable under § 103 for obviousness if the differences between the claimed subject matter 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 said subject matter pertains." *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). For a combination of known elements that are not explicitly found together in the prior art, we consider "whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue." *Id.* at 418 (citing *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)). A successful petition must "articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness." *In re Magnum Oil Tools Int'l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016) (citing *KSR*, 550 U.S. at 418); *see also* 35 U.S.C. § 322(a)(3); 37 C.F.R. §§ 42.22(a)(2), 42.104(b)(4) (2020).

We base our obviousness inquiry on factual considerations including (1) the scope and content of the prior art, (2) any differences between the claimed subject matter and the prior art, (3) the level of skill in the art, and (4) any objective indicia of obviousness or non-obviousness that may be in evidence. *See Graham*, 383 U.S. at 17–18.

Based on these factors,[6] Intel has shown by a preponderance of the evidence that claims 1, 5, 6, 8, 12, 14–24, 28, 30, and 33 are obvious over

[6] The parties do not point to any objective indicia in their briefs, and neither side's expert testifies that there are any such indicia. *See* Ex. 1001 ¶¶ 334–

Nicol, but has not shown that claims 2–4, 7, 9–11, 13, or 27 are obvious over Nicol. We present our findings and conclusions below.

1. Overview of Nicol

Nicol describes technology for reducing power consumption of multiprocessor chips by dynamically controlling their operating voltages. *See* Ex. 1008, 2:18–31, code (57). The system includes a controller for "allocat[ing] tasks to the individual processors to equalize processing load among the chips," "lower[ing] the clock frequency on the chip to as low a level as possible while assuring proper operation," and "reduc[ing] the supply voltage." *Id.* at 2:22–31.

Nicol's Figure 2, reproduced below, is a block diagram illustrating a multiprocessor chip with supply voltage control for the whole chip. Ex. 1008, 2:37–40.

340 (Dr. Sechen stating that he has not seen any evidence supporting any objective indicia of non-obviousness); Ex. 2027 ¶ 19 (Dr. Mangione-Smith omitting objective indicia in his discussion of the *Graham* factors).

FIG. 2

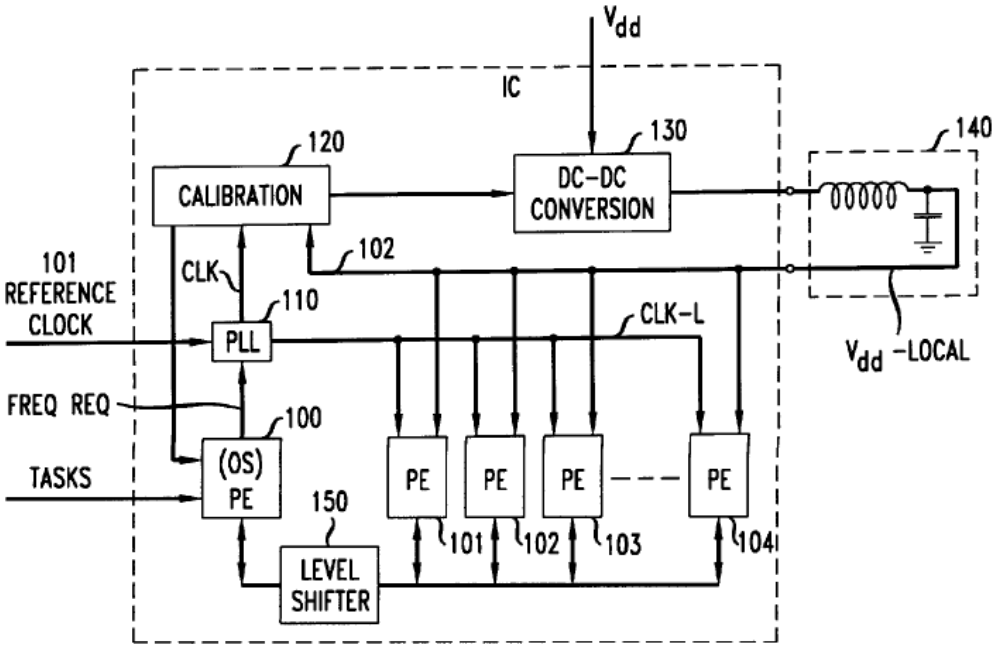
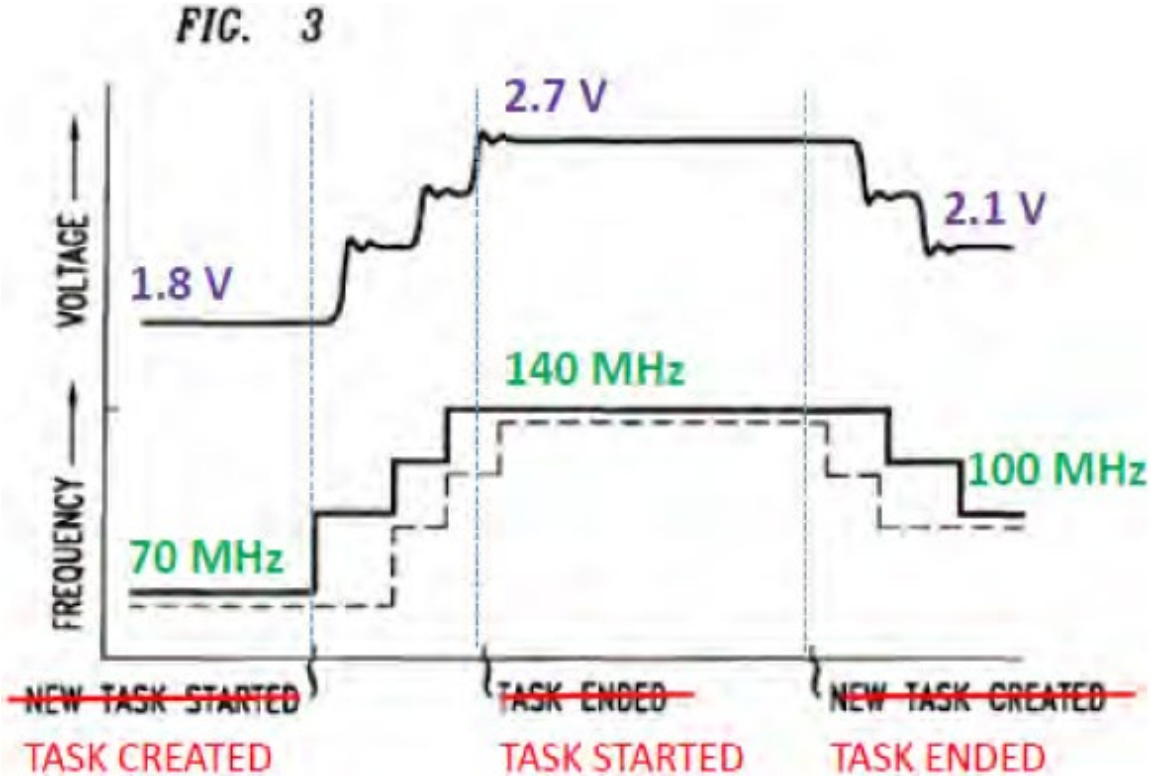


Figure 2 depicts processing elements (PEs) 100–104, each containing a central processing unit (CPU) and a local cache memory. Ex. 1008, 2:50–53. PE 100 runs a real-time operating system ("RTOS"), which allocates tasks to the other PEs. *Id.* at 2:54–55, 3:9–13. The operating system maintains a "to-do" list, and "[a]s more tasks are entered into the 'to-do' list, the operating system of PE 100 computes the correct way to balance the additional computational requirements and allocates the tasks to the processors. It then computes the required operating frequency." *Id.* at 5:24–28.

Nicol teaches that when an application is subdivided among the PEs into concurrent task streams, this allows the clock frequency to be reduced, which also allows the system to reduce the supply voltage, which reduces the power consumption. *Id.* at 3:22–42. As part of regulating frequency and voltage, programmable phase lock loop (PLL) frequency synthesizer circuit 110 generates two clock signals: Clk (the main clock signal) and Clk-L (a clock signal that lags behind Clk to ensure stability). *See id.* at 4:17–25,

4:38–53. PLL 110 sends the Clk signal to calibration circuit 120, which generates a supply voltage command, and then DC–DC converter 130 and L–C circuit 140 create a supply voltage V_{dd}-local which feeds back to calibration circuit 120 over line 102. *Id.* at 4:29–35. This embodiment applies the same supply voltage to all of the PEs. *Id.* at 4:35–37.

Nicol's Figure 3, reproduced below with Intel's annotations, illustrates the relationship between voltage V_{dd}-local and clock signals Clk and Clk-L in Figure 2:



Pet. 16. Figure 3 is a plot, over time, of supply voltage V_{dd}-local (top), clock signal Clk (bottom, solid line), and clock signal Clk-L (bottom, broken line) during the lifespan of a computing task. Ex. 1008, 4:54–58; *see also id.* at 2:40–43. As Intel asserts, and PACT does not dispute, the original Figure 3 contains obvious errors in the time-axis labels, which Intel corrected as

shown above in red. *See* Pet. 16–17 n.5 (citing Ex. 1001 ¶ 99; Ex. 1008, 4:58–66); PO Resp. 21 (reproducing Intel's annotated version of Figure 3, without disputing its accuracy). Intel has also added uncontested annotations, in green, showing the relevant voltages and clock frequencies as described in the associated text of the '047 patent. *See* Pet. 17; PO Resp. 21.

As shown above in annotated Figure 3, a new task is created ("TASK CREATED") at an initial state in which both Clk and Clk-L are set to 70 MHz while supply voltage V_{dd}-local is 1.8 V. Ex. 1008, 4:58–59. At this point, load Clk "is increased in three steps to 140 MHz," while V_{dd}-local increases to 2.7 V. *Id.* at 4:60–61. During this ramp-up, Clk-L lags behind Clk and only increases after each step when V_{dd}-local has stabilized (e.g., there is no longer any ringing). *See id.* at 4:38–53. When V_{dd}-local stabilizes after the third and final step, "a new task is enabled for execution" ("TASK STARTED"). *Id.* at 4:60–62. Sometime later, the task completes ("TASK ENDED"), "which reduces the load on the multiprocessor," and this permits lowering Clk to 100 MHz and lowering V_{dd}-local to 2.1 V. *Id.* at 4:62–66. "This, too, is accommodated in steps (two steps, this time), with Clk-L preceding Clk to insure, again, that the PEs continue to operate properly while the supply voltage is decreased." *Id.* at 4:66–5:2.

Nicol teaches that "[t]he 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. 1008, 5:52–56; *see also id.* at 6:34–37 ("The principles disclosed above for a multiprocessor [are] extendible to . . . systems with a plurality of separate processor elements that operate at different frequencies and operating voltages.").

Accordingly, Figure 4, reproduced below, illustrates a multiprocessor chip that can separately control each PE's supply voltage. Ex. 1008, 2:44–46.

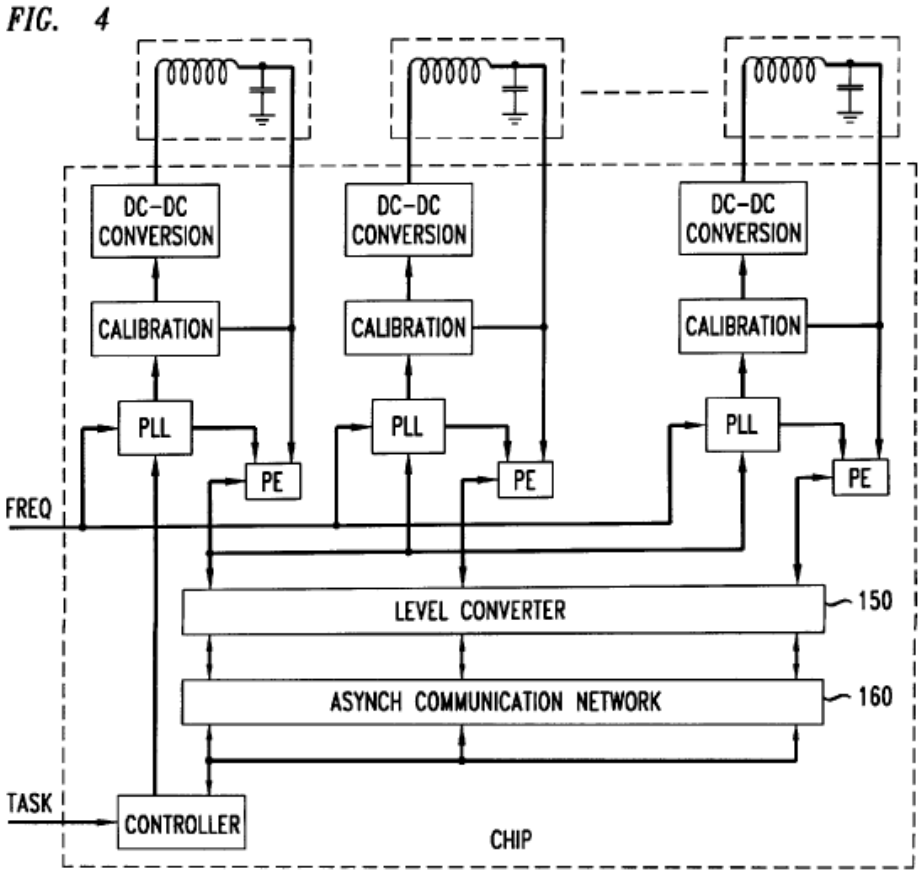


Figure 4 depicts three PEs, each with its own PLL, calibration circuit, DC–DC converter, and L–C circuit. *Id.* at 6:8–10. Thus, the system can regulate the frequency and supply voltage for each PE independently. *Id.* at 5:55–6:2. Because the PEs operate at different frequencies, the PEs are each connected to level converter 150, which "converts the variable voltage swings of the PEs to a fixed level swing," and asynchronous communication network 160, which "resolves the issue of different clock domains." *Id.* at 6:18–21, 6:29–33. In the bottom left corner, there is a single "controller . . . which may be just another processing element in the integrated circuit that contains the

other processing elements," and which "assigns the tasks given to the multi-processor chip of FIG. 4 among the PEs." *Id.* at 6:6–9, 11–13.

In addition to the embodiments in Figures 2 and 4, Nicol teaches that there is a "middle ground" embodiment in which "the PEs of a multi-processor chip can be divided into groups, and each group of PEs can be arranged to operate from its own controlled supply voltage." Ex. 1008, 6:67–7:3. According to Nicol, "a skilled artisan may easily make" this and other modifications within the scope of Nicol's claims. *Id.* at 6:58–61.

2. Independent claims 1 and 19

Independent claims 1 and 19 are materially the same, apart from the preambles. *Compare* Ex. 1003, 15:38–57, *with id.* at 16:60–17:13; *see also* Pet. 50 ("Claim 19 claims a mobile device (preamble) that includes the multiprocessor chip claimed in claim 1, and is otherwise identical to claim 1."). Thus, though we focus on claim 1 when explaining our analysis below, we reach the same findings and conclusions with respect to claim 19.

(a) Claim preamble

The preamble of claim 1 recites "[a] multiprocessor system." Ex. 1003, 15:38. Intel argues that Nicol discloses the preamble by disclosing a multiprocessor system. Pet. 21–22 (citing Ex. 1008, 2:44, code (54), Fig. 4; Ex. 1001 ¶ 119).

The preamble of claim 19 recites "[a] mobile device." Ex. 1003, 16:60. Although Intel contends the preamble is non-limiting, Intel argues that Nicol renders the preamble obvious for essentially the same reasons discussed below in relation to claim 18. Pet. 50; *see infra* part III.C.7.

PACT does not contest Intel's argument, and we agree that Nicol clearly discloses a multiprocessor system. PACT disputes Intel's arguments regarding claim 18 which are applicable to claim 19, but as we discuss below, we find Intel's arguments persuasive. *See infra* part III.C.7. Thus, we do not need to decide whether the preambles of claims 1 and 19 are limiting. *See Nidec v. Zhongshan*, 868 F.3d at 1017.

(b) "a plurality of data processing units, each of the data processing units: (a) being programmable"

Claim 1 requires a plurality of programmable data processing units. *See* Ex. 1003, 15:39–41. Citing Dr. Sechen's testimony, Intel argues that such programmable data processing units correspond to the processing elements (PEs) depicted in Nicol's Figures 2 and 4, which are programmable. *See* Pet. 22–24 (citing Ex. 1008, 2:50–53, 3:13–18, 6:6–13, 6:40–41, Figs. 2, 4; Ex. 1001 ¶¶ 121–122, 125–126, 129; Ex. 1011 ("Hayes"[7]), 48).

PACT does not contest Intel's arguments. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition.

(c) "each of the data processing units . . . (b) being adaptable for sequentially processing data in a clocked manner"

Intel argues that Nicol's multiprocessor chip is adaptable for sequentially processing data in a clocked manner, as recited in this limitation, because Nicol's asynchronous communication network 160

[7] John P. Hayes, *Computer Architecture and Organization* (3d ed. 1998).

(shown in Figure 4, *see supra* part III.C.1) passes results from a calculation in one PE to a second PE, which subsequently processes the data. Pet. 24–25 (citing Ex. 1008, 6:14–18, 6:26–28, Fig. 4; Ex. 1001 ¶¶ 132–133).

According to Intel,[8] "[a]synchronous communication networks were well-known in the art." Pet. 25 (citing Ex. 1001 ¶ 135; Ex. 1011, 491–93 (describing buses with asynchronous timing, where each device can generate bus-control signals at its own rate)). Citing Dr. Sechen's testimony, Intel argues that a person of ordinary skill in the art "would have understood that, in Nicol's system, 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.* (citing Ex. 1008, 6:19–21; Ex. 1001 ¶ 135). Intel also contends that this data exchange over the asynchronous network "is performed in a clocked manner, as required by the claim." *Id.* (citing Ex. 1001 ¶ 136; Ex. 1011, 492 ("[D]ata transmission can be asynchronous in the sense that the number of clock periods between bus events (signal changes) can be indeterminate, while the events themselves are synchronized by the clock.")).[9]

Intel also relies on Dr. Sechen's testimony that Nicol's multiprocessor system is designed to handle a variety of digital signal processing

[8] PACT interprets Intel's arguments as a "three step analysis." PO Resp. 12; Sur-reply 3 n.1. In its Reply, Intel clarifies that its analysis includes "three independent reasons," and we agree. Pet. Reply 3.

[9] Intel asserts that PACT "does not dispute that Nicol's PEs operate 'in a clocked manner.'" Pet. Reply 3 n.2 (citing PO Resp. 12–20). Thus, Intel contends, and we agree, that the only disputed issue regarding this limitation is "whether Nicol discloses PEs adaptable for sequentially processing data." *Id.* PACT did not dispute this characterization in its Sur-reply. *See* PO Sur-reply 3–4.

applications, which "involve complex operations commonly performed in a pipeline fashion, whereby partial results from one PE are passed to another PE of the pipeline for future processing." Pet. 27 (citing Ex. 1001 ¶¶ 139–140; Ex. 1018 ("Williams"[10]), 10:56–67, Fig. 4 (describing a system in which graphical data is processed successively through stages of a graphics pipeline)).[11]

In response, PACT first argues that Intel has not shown that Nicol actually discloses that its chip passes data or instructions between processors. PO Resp. 13–18. According to PACT, in Nicol's Figure 4, ACN 160 "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," but "[t]he use of the ACN does not necessitate communication across the PEs, much less imply that sequential data processing is performed." *Id.* at 15 (citing Ex. 1008, 6:14–21; Ex. 2027 ¶ 61). PACT also disputes that Hayes is evidence that Nicol's "ACN necessarily implies communication between PEs," because it is a non-analogous system and the communication between its processors is by a different type of internal bus. *Id.* at 15–18 (citing Ex. 1008, 2:51–53; Ex.

[10] Williams et al., US 6,397,343 B1 (issued May 28, 2002).

[11] Intel makes another argument that a person of ordinary skill in the art "would have known that the PEs in Nicol's multiprocessor system can be configured to sequentially process data in a manner similar to the well-known systolic architecture." Pet. 28 (citing Ex. 1001 ¶¶ 142–143; Ex. 1014, 37, 45, Fig. 1 (describing a systolic array)). When PACT challenged this argument in its Response (PO Resp. 19–20), Intel chose not to respond, given the strength of its other arguments. Pet. Reply 4–5 n.3; *see also* Sur-reply 3 (interpreting Intel's statement as abandoning its argument based on systolic architectures). We do not rely on this argument for our decision.

1011, 159, Fig. 3.14; Ex. 2027 ¶¶ 63–64). In its Sur-reply, PACT argues that "[t]he statements from Nicol only show an ability to communicate between PEs, and fall short of Petitioner's claim that Nicol shows that data output from one PE is passed to another PE." PO Sur-reply 4. PACT cites testimony of Dr. Mangione-Smith that "[i]n the case of communications among the PEs, there is a whole world of scheduling aspects, scheduling issues, that would have to be addressed because one task is now dependent upon another[,] none of which are discussed in Nicol." *Id.* (first alteration in original) (quoting Ex. 1052, 47:15–20) (citing Ex. 1052, 95:3–6).

Next, PACT argues that just "because Nicol performs digital signal processing and digital signal processing is typically performed using a chain of processors," this does not mean that "Nicol must be performed in the same manner as, for example, Exhibit 1018 ('Williams')." PO Resp. 18. According to PACT, "there are many architectures for digital signal processing," and passing data in a sequence "is not even the most common method off performing digital signal processing." *Id.*; *see also* Ex. 2027 ¶ 65 (opining that "[b]y 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."). PACT also argues that Williams describes a graphics processor rather than a digital signal processor, and thus its operation requires additional elements that Nicol does not disclose. PO Resp. 18 (citing Ex. 1018, 10:56–11:19, Ex. 2027 ¶ 65).

Considering the arguments and evidence on both sides, we find that Intel has persuasively shown that the PEs in Nicol's Figure 4 are "adaptable for sequentially processing data in a clocked manner." As we discuss above, the parties do not contest in this proceeding that "data processing units . . .

adapted for sequentially processing data" means "[s]equential data processors that pass results to one or more additional data processing units to perform additional computations." Ex. 1056, 2 (first alteration in original); Ex. 2029, 17; *see supra* part III.B.1. Nicol clearly discloses that data is passed from one PE to another PE. *See* Ex. 1008, 6:18–21, 6:25–33. And we are persuaded, given the context of Nicol's disclosure and Intel's other evidence—particularly the persuasive testimony of Dr. Sechen (Ex. 1001 ¶¶ 137–141)—that a person of ordinary skill in the art would have understood that one beneficial purpose of this PE-to-PE data transfer is so that, in a sequence of computations, one PE can send the data resulting from a computation to other PEs to perform additional computations.

PACT's arguments primarily address whether Nicol "*necessarily* implies that data is being communicated between PEs," but as PACT acknowledges, Nicol teaches a situation "when there is a need to communicate data between PEs." PO Resp. 15 (emphasis added) (quoting Ex. 1008, 6:19–21) (citing Ex. 2027 ¶ 61). And in its Sur-reply, PACT likewise acknowledges that "[t]he statements from Nicol . . . show an ability to communicate between PEs." PO Sur-reply 3–4. As Intel points out in its Reply, Dr. Mangione-Smith admitted during cross-examination that not only was PE-to-PE communication in Nicol's Figure 4 possible, but that a person of ordinary skill in the art would have envisioned this possibility. Pet. Reply 4–5 (citing Ex. 1052, 60:18–61:9, 95:11–15, 107:14–108:1, 110:21–25).

Finally, we address PACT's arguments that the ACN in Hayes or the graphics pipeline in Williams are not sufficiently analogous or compatible with the structures depicted in Nicol. *See* PO Resp. 15–18. Intel is not citing Hayes or Williams as prior art to be combined with Nicol. But we find that

these references support Intel's argument that a person of ordinary skill in the art would have understood that Nicol's chip has the ability to sequentially process data. *See* Pet. 25–29.

For the above reasons, we find that Nicol teaches data processing units that are "adaptable for sequentially processing data in a clocked manner" as recited in claim 1. Ex. 1003, 15:42–43.

(d) "each of the data processing units . . . including (c1) at least one Arithmetic Logic Unit (ALU), and (c2) at least one data register set for storing intermediate results in sequential data processing"

Intel argues that Nicol renders this limitation obvious. Pet. 29 (citing Ex. 1001 ¶ 148). According to Intel, each PE in Nicol's system "contains a central processing unit (CPU) and a local cache memory." *Id.* (quoting Ex. 1008, 2:50–53). Citing Dr. Sechen's testimony, Intel argues that a person of ordinary skill in the art "would have understood that a PE having a CPU and cache includes arithmetic logic unit(s) (ALU) and registers." *Id.* (citing Ex. 1001 ¶ 149; Ex. 1011, 33 (Hayes), Fig. 1.17). For example, Intel points to Hayes, a textbook showing a PE with multiple ALUs and associated registers. Pet. 29–30 (citing Ex. 1011, Fig. 1.17). According to Intel, "[s]ome of the registers are dedicated to storing results of intermediate computations performed by an ALU; those intermediate results can then be further processed." *Id.*

PACT does not contest Intel's arguments. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition.

(e) "at least one bus system for at least interconnecting at least some of the data processing units, each being provided a supply voltage by a voltage supply

Intel argues that asynchronous communications network 160 in Nicol's Figure 4 is a bus system as recited in this limitation. Pet. 30–31 (citing Ex. 1008, 6:26–33 ("[To effect] synchroniz[ation of] the exchange of data among the PEs of a multiprocessor chip . . . , each PE within the FIG. 4 arrangement is connect[ed] to an arrangement comprising 150 and 160."), Fig. 4; Ex. 1001 ¶ 153). Intel also argues that, as shown in Figure 4, each PE has its own voltage supply. Pet. 31–32.

PACT does not contest Intel's arguments in its Response. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition.

(f) "wherein for at least some of the data processing units, the clock frequency is adjustable at runtime without affecting the clock frequency of at least one of: one other of the data processing units, and the bus system"

Citing testimony of Dr. Sechen, Intel argues that Nicol discloses this limitation. Pet. 32 (citing Ex. 1001 ¶ 158). According to Intel, each of the PEs in Nicol's Figure 4 "is subjected to its own, individually controlled, supply voltage." *Id.* (quoting Ex. 1008, 6:63–66); *see also* Ex. 1001 ¶ 159. Citing testimony of Dr. Sechen, Intel argues that this "allow[s] the operating system to *independently program* the lowest operating frequency and corresponding lowest supply voltage *for each PE*." Pet. 32–33 (alteration in original) (quoting Ex. 1008, 5:66–6:2) (citing Ex. 1008, 6:14–18 (describing an embodiment where "the frequencies at which the individual PEs operate

differ from one another and from other elements within the system"); Ex. 1001 ¶ 160).

Intel also argues that in Nicol, "the clock frequency of each PE is adjustable at runtime, as illustrated in Nicol's annotated Figure 3" (reproduced above, *see* part III.C.1). Pet. 33 (citing Ex. 1008, 4:54–5:2, Fig. 3).

Other than arguing that Intel has failed to show a motivation to combine the different embodiments in Nicol (which we address below, *see infra* part III.C.2(h)), PACT does not contest Intel's arguments as to this limitation. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition.

(g) "wherein the voltage supply is adapted to supply higher supply voltages for data processing at higher clock frequencies"

Intel argues that Nicol discloses this claim limitation. Pet. 34 (citing Ex. 1001 ¶ 162). Citing testimony of Dr. Sechen, Intel argues that in Nicol's system, higher voltage is associated with higher clock frequencies. Pet. 34–35 (citing Ex. 1008, 4:39–41, 4:54–5:5, Fig. 1, Fig. 3; Ex. 1001 ¶¶ 162–163, 166).

Other than arguing that Intel has failed to show a motivation to combine the different embodiments in Nicol (which we address below, *see infra* part III.C.2(h)), PACT does not contest Intel's specific arguments as to this limitation. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition.

(h) Motivation to combine Nicol's embodiments

In its Petition, Intel relies on combining the teachings specifically relating to Nicol's Figures 2 and 3 with other disclosed embodiments including that of Figure 4. *See* Pet. 14–17, 32–49. In particular, Intel contends that the plot of voltage and frequency over time shown in Figure 3 is applicable to adjustments of "one or more PEs (or . . . an entire chip)." *See* Pet. 15 (citing Ex. 1001 ¶ 99), 17 (citing Ex. 1001 ¶ 101); *see also* Pet. 37 ("Nicol teaches a range of control strategies in which adjustments can be done selectively for each PE, can be common to a group of PEs, or can be common to all PEs of a chip.").

PACT argues that "the Figure 4 embodiment cannot simply be combined with [the] Figure 2 embodiment without modification." PO Resp. 10. According to PACT, for the Figure 2 embodiment, Nicol teaches balancing the load "'as evenly as possible' because the purpose is to find the allocation of tasks that would allow operation of the entire chip at the lowest clock frequency." *Id.* at 9 (citing Ex. 2027 ¶ 57). By contrast, in the Figure 4 embodiment, each PE can operate at an independent frequency, so according to PACT, "[i]n this scenario, there is no longer a need to allocate tasks as evenly as possible." *Id.* at 10 (citing Ex. 1008, 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).")). Thus, PACT argues that "[t]he features of Figures 2 and 4 are not interchangeable, as Figure 2 and Figure 4 deal with different scenarios." *Id.* And according to PACT, Intel "has not described at all how they should be combined," other than simply treating the embodiments as interchangeable. *Id.* at 11.

In response, Intel argues that "[c]ombining two embodiments disclosed adjacent to each other in a prior art patent does not require a leap of inventiveness." Pet. Reply 1 (quoting *Boston Scientific Scimed, Inc. v. Cordis Corp.*, 554 F.3d 982, 991 (Fed. Cir. 2009)). But in any event, Intel argues that its obviousness argument "relies on the core function of Nicol's [RTOS] software, which cuts across all embodiments." *Id.* at 1–2 (citing Ex. 1008, 2:18–27, 6:29–33). According to Intel, Nicol teaches that the RTOS can set the clock frequency and supply voltage independently for each PE (as in the Fig. 4 embodiment), the same for each PE (as in the Figs. 2 and 3 embodiment) or according to the disclosed "middle ground" between the two embodiments. *Id.* at 2 (citing Ex. 1008, 5:53–65, 6:61–7:3; Ex. 1051, 60:4–7, 65:17–25; Ex. 1055 ¶¶ 6–7).

In its Sur-reply, PACT responds that even if Intel has shown a commonality between the embodiments, this is insufficient because Nicol discloses little detail about how its RTOS works in any of Nicol's embodiments. PO Sur-reply 1–2. Also, according to PACT, Intel's argument is "based on constructed hypotheticals" as to the functions of the RTOS, "without any guidance as to how those functions would or should be performed." *Id.* at 2 (citing *InTouch Techs., Inc. v. VGO Commc'ns, Inc.*, 751 F.3d 1327, 1352–54; *TQ Delta, LLC v. Cisco Sys.*, 942 F.3d 1352, 1360-61 (Fed. Cir. 2019)). PACT also contends that "Nicol is 'underspecified," and thus non-enabling." *Id.* (quoting Ex. 1052, 54:18–24, 56:18–57:5) (citing Ex. 1052, 55:17–19; *Raytheon Techs. Corp. v. General Electric Co.*, 993 F.3d 1374, 1380 (Fed. Cir. 2021)).

Considering the arguments and evidence on both sides, we find that a person of ordinary skill in the art would have had reason, and an expectation

of success, for combining Nicol's various embodiments. First, we credit Dr. Sechen's testimony (Ex. 1055 ¶¶ 6–7) that combining the embodiments in Figures 2 and 4 would have been within the ordinary skill in the art at the time of the claimed invention. This is consistent with Nicol's teaching that a person of ordinary skill in the art could "easily" make a "middle ground" embodiment between Figures 2 and 4. *Id.* at 6:6:57–7:3. Although Dr. Mangione-Smith testifies on cross-examination that Nicol is "underspecified" (Ex. 1052, 54:18–24, 56:18–57:5; *see also id.* at 55:17–19), we do not find that testimony credible, based on the other evidence of record. An obviousness analysis "can take account of the inferences and creative steps that a person of ordinary skill in the art would employ." *KSR*, 550 U.S. at 418. Given the relatively high level of ordinary skill in the art (*see supra* part III.A), Intel has persuasively shown that making this combination would have been within that level of ordinary skill, and that an ordinarily skilled artisan would have reasonably expected that this combination would be successful.

Moreover, Intel's argument does not rely on bodily combining the Figure 2 and 4 embodiments; it is based primarily on applying the teachings of Figure 3 and other general descriptions of Nicol's RTOS, to the embodiment in Figure 4. Nicol explicitly teaches this combination: As part of its transition from discussing Figure 3 to discussing Figure 4, Nicol teaches, "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. 1008, 5:52–56; *see also id.* at 6:34–37 ("The principles disclosed above for a multiprocessor [are] extendible to . . . systems with a plurality of separate processor elements that

operate at different frequencies and operating voltages."). Thus, we find that a person of ordinary skill in the art would have had motivation, and a reasonable expectation of success, in adapting the RTOS control system illustrated in Nicol's Figure 3 to each PE, individually, as in the Figure 4 embodiment, or to a group of PEs as in the "middle ground" embodiment.

(i) Conclusion as to Claims 1 and 19

For the above reasons, Intel has shown persuasively that Nicol discloses or teaches each limitation of claims 1 and 19. We also find that an ordinarily skilled artisan would have had reason and the ability to combine the teachings of Figure 3 and the associated description with other embodiments in Nicol, including that of Figure 4 and the "middle ground" embodiment, with a reasonable expectation of success. Thus, weighing all of the evidence of record, Intel has shown by a preponderance of the evidence that claims 1 and 19 are unpatentable as obvious over Nicol.

3. Claims 2–4

Claims 2–4 each depend from claim 1 and further recite "wherein for at least some of the data processing units, the clock frequency is adjustable at runtime according to a state" of either "the multiprocessor system" (claim 2), "the data processing entity" (claim 3) or "others of the data processing units" (claim 4). Ex. 1003, 15:58–67. Intel contends that Nicol renders obvious these claim elements, but we do not find Intel's arguments persuasive. Pet. 35, 37, 39 (citing Ex. 1001 ¶¶ 168–169, 174–175, 181–182).

Intel argues that Nicol's clock frequency adjustments are made at runtime, as shown in Nicol's Figure 3. Pet. 35–36 (citing Ex. 1008, Fig. 3). According to Intel, these "adjustments can be done selectively for each PE,

can be common to a group of PEs, or can be common to all PEs of a chip." Pet. 37 (citing Ex. 1008, 6:57–7:3, Fig. 2); *accord* Pet. 39 (also citing Ex. 1001 ¶ 183).

For claim 2, in which the state is "of the multiprocessor system," Intel contends that Nicol's controller PE (PE 100 shown in Figure 2, *supra* part III.C.1) "'needs to ascertain the *required completion time*'—*i.e.*, priority—of the competing tasks, 'divide the collection of tasks' among available PEs 'and adjusts the clock frequency [of the PEs] to insure that *the most heavily loaded PE* carries out its assigned tasks with the *required completion time*.'" Pet. 36 (alteration in original) (quoting Ex. 1008, 3:66–4:5). For example, referencing the annotated version of Nicol's Figure 3 (*supra* part III.C.1), Intel argues that "prior to tasks being created, the multiprocessor system is in a low-activity state or a non-operational state and the controller maintains the clock frequency at a low value (70 MHz illustrated)." *Id.* (citing Ex. 1008, 4:54–62). According to Intel, "[a]fter task creation, the multiprocessor system transitions to a high-activity operational state and the controller adjusts the clock frequency to a high value (140 MHz illustrated)." Pet. 36–37 (citing Ex. 1008, 4:54–62).

Intel makes a similar argument for claims 3 and 4, in which the state is "of the data processing entity" or "of others of the data processing units." Pet. 37–41 (citing Ex. 1008, 3:66–4:5, 4:54–62, 6:61–7:3, Figs. 2, 3; Ex. 1001 ¶¶ 177, 179–180, 184–185). In these cases, Intel identifies the recited "state of the data processing entity" or "state of the others of the data processing units" as the state of "the most heavily loaded PE" in Nicol's system. Pet. 39, 41 (quoting Ex. 1008, 4:3–4) (citing Ex. 1001 ¶¶ 180, 187).

In response, PACT argues that, "[t]o the extent that the[] different voltage[s] and frequencies [depicted in Intel's annotated Figure 3] might be considered states within the meaning of claims 2–4, Nicol does not teach that 'the clock frequency is adjustable at runtime according to a state.'" PO Resp. 21 (citing Ex. 2027 ¶ 70). Rather, according to PACT, "the frequency is established according to time and load limitations and it is the frequency that drives the state." *Id.* at 22 (citing Ex. 1008, 2:56–58 ("The load of the FIG. 2 system is time varying and is dependent on the applications that are being executed at any given time."), 3:15–18 ("[A] scheduler within the operating system can use [the number of instructions that need to be executed for each task] to determine the best way to allocate the tasks to the available processors in order to balance the computation."), 4:3–5 ("[Operating system PE 100] adjust[s] the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time."), 5:29–30 ("[T]he frequency is gradually programmed into the system (as shown by the stepped changes in FIG. 3)."). Thus, according to PACT, "Nicol discloses the opposite of 'the clock frequency is adjustable at runtime according to a state': the alleged 'states' in Nicol are adjustable according to the clock frequency." *Id.* (citing Ex. 2027 ¶ 71).

In its Reply, Intel argues that the '047 patent admits that "[t]o determine the clock frequency for a microprocessor based on the state *is known*," thus indicating that "this claim element wholly lacks novelty." Pet. Reply 5 (quoting Ex. 1003, 2:30–31). Intel also contends that during cross-examination, Dr. Mangione-Smith admitted "that Nicol teaches adjusting the clock frequencies of PEs at runtime based at least on load, and load 'represents a state.'" *Id.* at 6 (citing Ex. 1052, 159:4–16, 159:21–24).

Further, Intel points to passages in Nicol indicating a relationship between the clock frequency and the "load on the multiprocessor." *Id.* (citing Ex. 1008, 4:54–66; Ex. 1055 ¶ 13).

In its Sur-reply, PACT counters that "[t]he load that is described in Nicol is anticipated load, as Nicol makes a clock frequency determination based on future tasks. A target frequency is first determined, prior to any frequency change, based on tasks that have not yet been executed." PO Sur-reply 4–5 (citing Ex. 2032, 93:1–16 (Dr. Sechen testifying, on cross-examination, that Nicol's operating system controller determines the voltage and frequency when it receives a task, "before [the controller] sends it to a PE or group of PEs")). According to PACT, "[t]he frequency is then increased in steps, and only after the frequency has stabilized does execution of the task begin." *Id.* at 5 (citing Ex. 2032, 95:2–20 (Dr. Sechen testifying that "once the ultimate voltage and frequency are achieved, then that's when execution starts")). Thus, PACT argues that "in Nicol the actual load on the system is never part of what determines the clock frequency." *Id.*

Addressing the statement in the '047 patent that "[t]o determine the clock frequency for a microprocessor based on the state is known" (Ex. 1003, 2:30–31), PACT argues that "this isolated statement does not motivate a [person of ordinary skill in the art] to apply it to a system like Nicol, particularly since Nicol does not set the clock frequency based on a state." PO Sur-reply 5.

Considering the arguments and evidence on both sides, including the supporting declaratory and cross-examination testimony, we do not find persuasive Intel's argument that adjusting the frequency based on the processor "load" meets the requirement that "the clock frequency is

adjustable at runtime according to a state" as recited in claims 2–4. As we discuss above, we construe this term to mean that the clock frequency is adjusted as a result of, or in response to, an existing state of the multiprocessing system or its data processing units. *See supra* part III.B.2. We agree with PACT that the "load" mentioned in Nicol is not a load reflecting the actual state of the multiprocessing system or any individual PE, but an anticipated or calculated load, based on the tasks listed in the operating system's "to-do" list. *See* Ex. 1008, 5:23–30.

Intel has not identified anything in Nicol that corresponds to an actual state of either the multiprocessor system or any of its individual PEs. Rather, the evidence shows that Nicol sets the frequency *before* sending tasks to the PEs to increase their actual load. *See* Ex. 1008, 4:58–62 (disclosing that "the new task is enabled for execution" only after the clock frequency is increased and the voltage supply is stable); Ex. 2032, 93:12–16; Pet. 16 (annotated Fig. 3, showing that the task is started, and thus the actual processor load is increased, only after the operating system increases the clock frequency to its target of 140 MHz). Thus, Intel has not shown by a preponderance of the evidence that in Nicol's system, "the clock frequency is adjustable at runtime according to a state."

We also disagree with Intel that the statement in the '047 patent that "[t]o determine the clock frequency for a microprocessor based on the state is known" (Ex. 1003, 2:30–31) suggests that claims 2–4 fail to recite patentable subject matter. This passage only explicitly relates to determining (not necessarily adjusting) the clock frequency for an *individual microprocessor* based on a state, but does not explicitly admit a lack of

inventiveness for the entire solution recited in claims 2–4 for adjusting frequency in a multiprocessor system, based on a state.

Finally, in its Reply, Intel raises an alternative argument that "Nicol's RTOS takes into account other PEs' states besides their load in determining how to allocate tasks to PEs and adjusting frequencies at runtime." Pet. Reply 6 (citing Pet. 36). According to Intel, Nicol's disclosure of "employ[ing] gated clocking techniques to 'shut down' PEs that are not needed" would have indicated to a person of ordinary skill in the art "that Nicol's RTOS takes into account the fact that certain PEs may[ ]be in such a low-activity or non-operational state (and clock gated) when it allocates tasks that enter the 'to-do' list to the remaining active PEs and sets their clock frequencies accordingly." *Id.* at 6–7 (quoting Ex. 1008, 5:41–45) (citing Ex. 1008, 5:23–27, 10:34–36; Ex. 1055 ¶¶ 14–15); *see also* Ex. 1055 ¶ 15 ("Depending on the time requirements, [Nicol's] RTOS will likely include [a non-gated] PE in task allocation, but possibly not [a clock-gated] PE unless it is truly needed because it must take into account the time it takes for the [clock-gated] PE to come back online.").

First, we agree with PACT that this argument is untimely. *See* PO Sur-reply 22. Although Intel cites a portion of its petition that refers to a multiprocessor system being in "a non-operational state," Pet. Reply 7 (citing Pet. 36), this was in the context of the anticipated load on the processors, and did not mention clock gating. *See* Pet. 36. Intel did not raise the theory that a clock-gated processor was in a different "state" than a non-clock-gated processor (apart from their different anticipated loads) until its Reply, which deprived PACT of the ability to submit an expert declaration on the issue. PO Sur-reply 22.

Second, even if Intel's argument were timely, we would still find it unpersuasive. Dr. Sechen's testimony that a person of ordinary skill in the art would have understood that a PE's clock-gated status would be a factor that Nicol's operating system uses to allocate tasks is conclusory and inconsistent with his cross-examination testimony. Pet. Reply 6–7; Ex. 1055 ¶ 15. In his deposition, Dr. Sechen testified that clock gating involves setting the clock frequency to zero, but not turning off the processor's voltage supply. *See* Ex. 2030, 34:17–23; Ex. 2033, 138:18–21, 140:7–15. Although Dr. Sechen testified that shutting off the voltage to a processor would result in "thousands of [clock] cycles to get back to a stable clock," he said that gate clocking merely involves "just the delay of an AND gate, [and] you can be back having the desired clock frequency accurately." Ex. 2033, 144:18–23. This testimony undermines the idea that a clock-gated PE in Nicol's system would take such a long time to come back on line that Nicol's operating system would consider that delay in its criteria for allocating tasks to other PEs. Instead, we find PACT's arguments persuasive on this point. Sur-reply 6–7 (citing Ex. 2030, 35:2–9; Ex. 2033, 137:18–144:23; Ex. 1055 ¶ 56).

For the above reasons, we determine that Intel has not shown, by a preponderance of the evidence, that claims 2, 3, or 4 are unpatentable as obvious over Nicol.

4. Claims 5, 17, and 22

Claim 5 depends from claim 1 and further recites "wherein for at least some of the data processing units, the clock frequency is adjustable at runtime according to a priority of a sequential task." Ex. 1003, 16:1–3. Similarly, claim 17 depends from claim 1 and further requires that the

system is "adapted to selectively set the clock frequency of at least some of the processing units in response to a prioritization of a task assigned." *Id.* at 16:55–57. Claim 22 depends from independent claim 19 and further recites "wherein for at least some of the data processing units, the clock frequency is adjustable at runtime according to a priority of a sequential task." *Id.* at 17:18–21.

Intel argues that Nicol renders the limitation of claim 5 obvious. Pet 41. First, Intel argues that in Nicol, the clock frequency is adjustable at runtime, as shown in Figure 3. Pet. 41–42 (citing Ex. 1008, Fig. 3). Second, Intel argues that Nicol "teaches that the clock frequency adjustment of PEs is made according to a priority of a sequential task." Pet. 42 (citing Ex. 1001 ¶ 190). According to Intel, this is because Nicol teaches that there is a dynamic mix of applications with different load requirements and sequential tasks that involve executing multiple operations, and Nicol's controller PE "'needs to ascertain the required completion time' of the competing tasks, 'divide[] the collection of tasks' among available PEs, 'and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time.'" *Id.* (emphasis omitted) (alteration in original) (quoting Ex. 1008, 3:66–4:5) (citing Ex. 1008, 2:53–64, 3:13–18). Intel contends that a person of ordinary skill in the art "would have understood that Nicol demonstrates a priority assignment scheme based on PE workload," which allocates the highest clock frequency to the most heavily loaded PE, so that it can complete the assigned tasks within the required completion time. *Id.* (citing Ex. 1001 ¶ 193). Intel relies on the same argument for claim 22. See Pet. 50–51.

Intel relies on the same arguments for claim 17. *See* Pet. 48–49. But unlike claims 5 and 22, claim 17 does not explicitly require that the task be "sequential." *Compare* Ex. 1003, 16:1–3, and *id.* at 17:18–21, *with id.* at 16:55–57. Thus, Intel argues that Nicol's Figure 4 embodiment discloses the added limitation of claim 17 by allowing "the operating system to *independently* [*i.e.*, selectively] program the lowest *operating frequency* and corresponding lowest supply voltage *for each PE*." Pet. 47–48 (alteration in original) (quoting Ex. 1008, 5:66–6:2) (citing Ex. 1008, 6:14–18, Fig. 4).

PACT makes a number of arguments in response. PO Resp. 25–32. First, relying on the same arguments it made for claim 1, PACT contends that Intel has not shown that Nicol discloses a sequential task. *Id.* at 26 (citing Ex. 2027 ¶ 77). We address those arguments above in the context of claim 1, and we likewise disagree with them in the context of claims 5, 17, and 22. *See supra* part III.C.2(c). We also find persuasive Intel's argument that during cross-examination, Dr. Mangione-Smith admitted that "[t]he PEs in Nicol could . . . execute[] instructions *in sequence* to complete tasks." Pet. Reply 7 (alterations in original) (quoting Ex. 1052, 88:3–6).

PACT's next set of arguments relate to the combinability of Nicol's various embodiments. *See* PO Resp 26–32. First, for claims 5 and 22, PACT argues that Intel relies on the embodiment in Nicol associated with Figures 2 and 3, which only describes a single set of frequency and voltage changes for all PEs together, and does not tie the argument to running multiple tasks with different priorities. PO Resp. 26–27 (citing Ex. 1008, 3:31–42, 3:66–4:6, Figs. 2, 3; Ex. 2027 ¶ 78; Ex. 2030, 21:20–22:7; Pet. 41–42; Ex. 1001 ¶¶ 188–193). For claim 17, in which Intel also relies explicitly on the embodiment in Nicol's Figure 4, PACT argues that Intel has not explained

"how Figure 4 discloses 'prioritization of a task assigned,' instead only quoting language indicating the ability to independently program operating frequency and supply voltage for each PE." *Id.* at 26.

PACT also argues that in the embodiment of Nicol's Figures 2 and 3, "Nicol never identifies any task as prioritized[,] as that is not a goal or purpose of Nicol." PO Resp. 27 (citing Ex. 2027 ¶ 80). According to PACT, Intel's argument relies on frequency being set according to the processor load (i.e., the "*number of* tasks" assigned to each PE), which associates priority with a *processor*, rather than with any particular *task*. *Id.* at 25 & n.1, 27–28 (citing Ex. 2027 ¶ 80). Also, PACT argues that since the Figure 2 embodiment sets each PE to the same frequency, "the most heavily loaded PE would actually finish last, in the sense that it is not prioritized." *Id.* at 28 (citing Ex. 2027 ¶ 82).

In its Reply, Intel counters that Nicol specifically teaches real-time tasks that have required completion times, as well as non-real-time tasks for which "there might not be any particular requirement for when the tasks are to be finished." Pet. Reply 8 (quoting Ex. 1008, 3:49–50) (citing Ex. 1008, 3:47–49, 4:5; Ex. 1051, 86:10–15 (Dr. Mangione-Smith agreeing that Nicol describes that there are "some prioritized real-time tasks along with some tasks that don't have time requirements"); Ex. 1055 ¶ 18). According to Intel, Dr. Mangione-Smith admitted that Nicol's "RTOS prioritizes real-time tasks such as voice calls—which have stringent completion time requirements—over non-real-time tasks that might have more relaxed, or no completion times." *Id.* (citing Ex. 1051, 74:9–15; Ex. 1055 ¶ 19).

We find Intel's arguments persuasive, and we disagree with PACT. Intel has persuasively shown that Nicol's disclosure that some (but not

necessarily all) tasks in Nicol's system may have required completion times, in light of Nicol's teachings as a whole, would have led a person of ordinary skill in the art to assign higher priorities to tasks in order to meet those timing requirements. We credit Dr. Sechen's testimony to this effect (Ex. 1055 ¶¶ 18–19) and we find that Dr. Mangione-Smith's cross-examination testimony (Ex. 1051, 74:9–15, 86:10–15) undermines his declaratory testimony on this topic (Ex. 2027 ¶¶ 81–82).

Next, PACT relies on its arguments for claim 1 to assert that Intel has not provided any motivation to combine Nicol's Figure 4 embodiment with the embodiment in Figures 2 and 3. PO Resp. 28–29 (citing Ex. 2027 ¶¶ 83–84). We address these arguments above, and we disagree for the same reasons. *See supra* part III.C.2(h).

Finally, PACT argues that Intel's proposed combination of Nicol's embodiments would run counter to Nicol's purpose "to reduce power, by reducing the clock frequency, based on processor load." PO Resp. 29 (citing Ex. 2027 ¶ 85). According to PACT, a person of ordinary skill in the art "would understand that Nicol attempts to reduce the clock frequency whenever possible, depending on processor load." *Id.* at 29–30 (citing Ex. 1008, 3:2–4, 3:24–25, 5:66–6:2; Pet. 14; Ex. 1001 ¶¶ 95, 101, 215; Ex. 2027 ¶¶ 85–86). However, PACT contends that Intel's argument relies on *increasing* the frequency of the most heavily loaded PEs. *Id.* at 31 (citing Ex. 1001 ¶¶ 193, 213 (Dr. Sechen testifying that "'the most heavily loaded PE' has the highest priority and would receive the highest clock frequency such that it can complete 'its assigned task in the required completion time.'" (quoting Ex. 1008, 3:66–4:5)); Ex. 2030, 14:16–22 (Dr. Sechen testifying on cross-examination that "[y]ou would typically set the frequency of the most

heavily loaded PE to the lowest frequency that would meet the required computation time, completion time, . . . such that you therefore minimize power")). According to PACT, this "is facially contradictory and logically contrary to the disclosure and teachings of Nicol, as well as [Intel's] and Sechen's other statements in the record." *Id.* at 31–32 (citing Ex. 2027 ¶¶ 88–89).

We disagree. Even if Nicol teaches reducing the clock frequency whenever possible, this does not preclude increasing the clock frequency whenever it is *not* possible to decrease the frequency while still ensuring that a task will complete within its allotted time. *See* Ex. 1008, 5:60–64 (teaching that for some PEs, a higher frequency is required if tasks cannot be partitioned into evenly loaded threads); *id.* at 7:51–55 (reciting a step of "ascertain[ing] the lowest frequency that, when employed at the processor carrying the largest load of sub-tasks processing, is sufficient to complete its assigned sub-tasks processing within the allotted time").

For the above reasons, Intel has shown persuasively that Nicol teaches each limitation of claims 5, 17, and 22, and that a person of ordinary skill in the art would have had reason to achieve the claimed invention, with a reasonable expectation of success, by combining the embodiment of Figure 4 with Nicol's teachings about its RTOS as particularly shown in Figure 3. Thus, weighing all of the evidence of record, Intel has shown by a preponderance of the evidence that claims 5, 17, and 22 are unpatentable as obvious over Nicol.

5. Claims 6, 23

Claims 6 and 23 depend from claims 1 and 22, respectively, and recite "wherein the processor is adapted to reduce the supply voltage when the clock frequency is reduced." Ex. 1003, 16:4–6, 17:22–24.

Citing testimony by Dr. Sechen, Intel argues that Nicol discloses this limitation, as shown in Figure 3. Pet. 43–44 & n.7 (citing Ex. 1008, code (57), 1:26–30, 4:28–31, 4:64–5:2, Fig. 3; Ex. 1001 ¶¶ 194–195, 198–199); *id.* at 51 (for claim 23, relying on the same arguments as for claim 6).

PACT does not contest Intel's arguments. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition, and because Nicol's Figure 3 clearly shows that when the clock frequency is reduced, the system responds by reducing the supply voltage. *See* Ex. 1008, Fig. 3. Thus, we conclude that claims 6 and 23 are unpatentable as obvious over Nicol.

6. Claims 7 and 27

Claims 7 and 27 depends from claims 6 and 20, respectively, and require that the system or mobile device is "adapted to subsequently reduce the clock frequency in accordance with a hysteresis characteristics [sic]." Ex. 1003, 16:7–9, 17:50–52.

Although neither party proposes an explicit construction for "a hysteresis characteristics [sic]," and the '047 patent itself contains no helpful description,[12] Intel describes hysteresis as follows:

[12] Intel contends that "[t]he '047 Patent specification is devoid of any description of a hysteresis characteristics [sic]." Pet. 45 n.8. We agree, although the '047 patent does include an example of when it may be appropriate to provide "a hysteresis characteristic . . . in determining the

> Hysteresis involves the nonlinear response of a circuit, in which the response to a particular set of input conditions depends both on the instantaneous values and the recent past of the input and output signal, with such behavior further characterized as the inability to retrace exactly on the reverse swing of a particular set of input and output conditions.

Pet. 45 (citing Ex. 1012[13], 357 (definition 5 of "hysteresis"); Ex. 1001 ¶ 203). PACT does not dispute this understanding of hysteresis (although it does not endorse it either). PO Resp. 34 n.2; Ex. 2027 ¶ 45 (Dr. Mangione-Smith stating, "I will respond to this definition for the purposes of my analysis but that is not an indication that I accept it as accurate or useful in this context."). Since it is not disputed, we do not need to explicitly construe the term and will accept that understanding of the term in our analysis. *See Nidec*, 868 F.3d at 1017.

Intel argues that Nicol renders obvious the added limitations of claims 7 and 27, and relies on the same arguments for both claims. Pet. 45, 52 (citing Ex. 1001 ¶ 200). Intel argues that "Nicol is particularly concerned with avoiding unwanted noise when adjusting frequencies of the PEs and uses hysteresis to avoid such unwanted noise," and "to avoid excessive switching" due to false triggers. *Id.* (citing Ex. 1008, 5:29–36; Ex. 1016,[14] 471; Ex. 1001 ¶¶ 201–202); Pet. Reply 8–9.

---

clock rates," which is "when a temperature-sensitive change of the clock frequencies is to be performed." Ex. 1003, 14:14–16.

[13] Ridolf F. Graf, *Modern Dictionary of Electronics* (9th ed. 1999)

[14] Michael Peter Kennedy & Leon O. Chua, *Hysteresis in Electronic Circuits: A Circuit Theorist's Perspective*, 19 Int'l J. Circuit Theory & Applications 471 (1991).

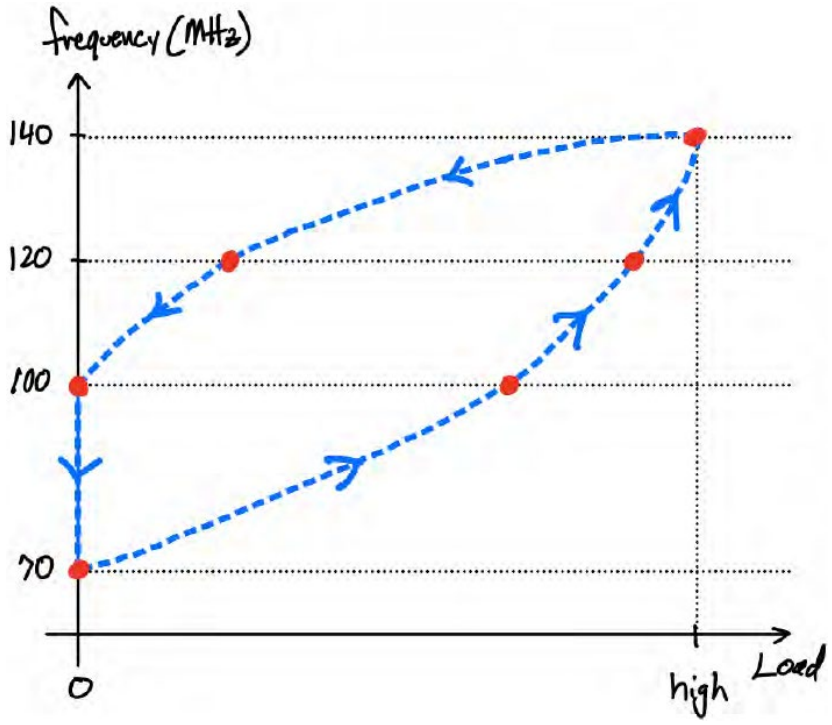
According to Intel, a person of ordinary skill in the art "would have understood that the clock frequency adjustment in Figure 3 (annotated) exhibits hysteresis, depicting a non-linear frequency response to the change in the load (addition of tasks) on the system." Pet. 45–46 (citing Ex. 1001 ¶¶ 204–206). In particular, Intel argues that the Clk-L signal (broken line) exhibits hysteresis because "[t]he response to an initial change in load (depicted as 'task started') shows a non-linear, delayed, gradual response with the increasing step frequency changes affected by past loads." Pet. 46 (citing Ex. 1008, Fig. 3; Ex. 1001 ¶¶ 204–207); *see supra* part III.C.1 (annotated Figure 3). According to Intel, "[a] different non-linear behavior is depicted when the task ends ('task ended') and the reduced load permits lowering the frequency, taking into account the historical behavior." *Id.* (citing Ex. 1001 ¶¶ 204–207).

In other words, Intel argues that "Nicol's RTOS deliberately introduces lag to slow the [voltage and frequency] adjustments" and this lag, itself, is a form of hysteresis. Pet. Reply 10 (citing Ex. 1008, 5:34–35; Ex. 1012, 357 (definitions 6 ("The lag in the response of an instrument or process when a force acting on it changes abruptly."), 8 ("This term literally means to lag behind.")); Ex. 1053,[15] 355 ("3. Any phenomenon in which there is a lag between the cause and the induced or observed effect."); Ex. 1055 ¶ 21). In this way, Intel argues that "Nicol uses hysteresis to ensure that small workload fluctuations do not unnecessarily trigger [voltage or frequency] step changes, as would have been well-known to a [person of

[15] Steven M. Kaplan, *Wiley Electrical & Electronics Engineering Dictionary* (2004).

ordinary skill in the art]." Pet. Reply 9 (citing Ex. 1055 ¶ 26); *id.* at 10–11 (citing Ex. 1051, 35:4–18, 37:3–21; Ex. 1055 ¶¶ 25–26).

Intel also argues that, "in accordance with hysteresis, the frequency (output) does not 'retrace exactly on the reverse swing' when the load is then reduced, *i.e.* starting frequency at 70 MHz, and ending at 100 MHz." Pet. 47 (citing Ex. 1012, 357; Ex. 1001 ¶¶ 204–207). In support, Intel relies on a figure provided by Dr. Sechen, "depicting the locus of points for how the frequency in Nicol's Fig. 3 varies with changes in load," which we reproduce below:



Pet. Reply 11–12; *accord* Ex. 1055 ¶¶ 23–24. Dr. Sechen's figure, above, is a plot of frequency (MHz) versus "Load," identifying the path of the system's load and clock frequency with blue arrows. According to Intel, "the clock frequency starts at 70 MHz and ramps up to 140 MHz when a new task is created." Pet. Reply 11 (citing Ex. 1055 ¶ 23; Ex. 1008, 4:54–62).

Then, "[w]hen the task is completed and the corresponding load goes to zero, the frequency is slowly reduced to 100 MHz, but does not go back to the starting frequency of 70 MHz until much further delay." *Id.* Intel contends that "[t]his illustrates the inability to retrace the reverse swing, which is characteristic of hysteresis." *Id.* (citing Ex. 1055 ¶ 23).

We do not find Intel's arguments persuasive. First, although Nicol's annotated Figure 3 shows that there is a delay between a change in the processor load (i.e., when a new tasks begins executing) and the Clk-L signal (broken line), we agree with PACT that this delay is not a nonlinear *response* by the system, but a deliberate, pre-determined choice by Nicol's RTOS. *See* PO Resp. 34–35 (citing Ex. 2027 ¶ 94); *see also supra* part III.C.3 (disagreeing with Intel's contention that Nicol adjusts the clock frequency according to an actual state of the multiprocessor system or of any particular PE).

Second, we agree with PACT that Intel has failed to show that Figure 3 shows an *inability* (rather than a deliberate choice) to retrace a reverse swing of input and output conditions. *See* PO Resp. 37–40 (citing Ex. 1008, 5:30–31; Ex. 1016, 473, 496, Fig. 7; Ex. 2027 ¶¶ 99–100). Although Nicol's Figure 3 begins at 70 MHz and ends at 100 MHz, Intel has not pointed to any teaching in Nicol indicating that Nicol's controller could not have continued to decrease to 70 MHz after the task ended, depending on its chosen control algorithm, since Nicol's RTOS predetermines the Clk and Clk-L signals at each step, based on the anticipated load. *See* Ex. 1008, 3:15–18, 4:3–5, 5:23–30, 5:66–6:2.

Finally, Intel has not persuasively shown that in Nicol's control system there is a need to respond to small workload fluctuations that might

unnecessarily trigger a step change in the frequency or voltage. This is because Intel has not shown that Nicol's RTOS responds to actual, unanticipated workload fluctuations. Rather, the RTOS predetermines the system's step changes based on the anticipated workload (based on the number of tasks on the "to-do" list), and calculates the optimal voltage based on the selected frequency. *See* Ex. 1008, 5:3–30.

In its Reply, Intel raises two new theories for why Nicol teaches hysteresis. First, Intel contends that Nicol's calibration circuit 120 (shown in Figures 2 and 4), determines the proper supply voltage by "operat[ing] two copies of a PE's critical path circuit," one of the copies operating at frequency Clk, and the other operating at a slightly lower frequency, so that at the optimal supply voltage, the first copy operates correctly (passes) and the slower copy fails. Pet. Reply 12 (citing Ex. 1008, 4:38–5:18; Ex. 1054;[16] Ex. 1055 ¶ 27). According to Intel, a person of ordinary skill in the art "would have understood that this calibration circuit necessarily involves hysteresis, otherwise any minute environmental factor change would cause the two copies to constantly switch from the Pass/Fail 'sweet spot' state to a Pass/Pass state, which triggers a voltage reduction, or to a Fail/Fail state, which triggers a voltage increase." *Id.* at 12–13 (citing Ex. 1055 ¶ 27; Ex. 1051, 118:22–120:1).

We do not find Intel's argument persuasive. Even if we accept Intel's argument that Nicol's calibration circuit 120 sets the supply voltage using a hysteresis characteristic, Intel has only established that supply voltage (i.e.,

[16] Tadahiro Karoda et al., *Variable Supply-Voltage Scheme for Low-Power High-Speed CMOS Digital Design*, 33 IEEE J. Solid-State Circuits, 454 (1998) (cited in Nicol, Ex. 1008, 5:5–6).

V_{dd}-local) is set in accordance with a hysteresis characteristic—not clock frequency (e.g., Clk or Clk-L) as recited in claims 7 and 27.

Second, Intel raises the argument in its Reply that Nicol uses hysteresis "to avoid excessive [voltage or frequency] switching due to reasons other than a change of workload, such as changes in device temperature." Pet. Reply 9 (citing Ex. 1055 ¶ 28). According to Intel, "Nicol discusses using temperature measurements to improve its voltage and frequency settings," and "[j]ust like thermostat control, a [person of ordinary skill in the art] would have understood that Nicol uses hysteresis to avoid minute changes in temperature from triggering unnecessary changes of voltage and frequency." *Id.* at 13 (citing Ex. 1008, 4:11–16; Ex. 1055 ¶ 28).

We find this argument unpersuasive. We agree with PACT that the passage Intel cites for this argument has no apparent relation to hysteresis. *See* PO Sur-reply 11. The passage states that "the operating system can reduce the supply voltage even further by tracking temperature and process variations," and then it gives an example of selecting a frequency along a "nominal" line on a plot of frequency vs. voltage in a prior art chip (Nicol's Figure 1). Ex. 1008, 4:11–16; *see also id.* at 1:9–18, 2:37–39. Intel has not explained how selecting a frequency in this way involves hysteresis or anything like thermostat control. Figure 1 does not reference any control loop involving temperature, or even explicitly reference temperature at all. *See id.* at Fig. 1. Further, as with the previous argument, even if we accept Intel's arguments, Intel only establishes that supply voltage—not clock frequency—is set in accordance with a hysteresis characteristic.

For the above reasons, we determine that Intel has not shown, by a preponderance of the evidence, that claims 7 or 27 is unpatentable as obvious over Nicol.

7. Claims 18 and 33

Claim 18 depends from claim 1 and further recites "wherein the plurality of data processing units are disposed within a mobile device." Ex. 1003, 16:58–59. Similarly, claim 33 depends from claim 20 and further recites "wherein the mobile device is a mobile computer." *Id.* at 18:11–12.

For claim 18, Intel argues that "Nicol discloses this claim element," and in particular, "Nicol discloses power minimization of a multiprocessor DSP chip used in consumer products such as personal computers, which include laptop computers already ubiquitous by the late 1990s." Pet. 49 (citing Ex. 1008, 2:54–64, 6:39–41). Intel also argues that "Nicol's power minimization is particularly beneficial to mobile devices, such as laptop computers, to extend operational time between battery charges." Pet. 49–50 (citing Ex. 1001 ¶ 215 (opining that a person of ordinary skill in the art would have understood "that the primary application inherently targeted by Nicol was mobile devices"); Ex. 1017 ("Cline"[17]), 1:13–17 (teaching that the use of low-power states in a computer system is especially useful to reduce power consumption in battery-powered devices)). Intel relies on the same arguments for claim 33. Pet. 52 (citing Ex. 1001 ¶¶ 228–229).

In response, PACT argues that Intel's argument should be understood as an inherency argument, rather than one based on obviousness, because of the structure of the argument, because Intel and Dr. Sechen do not use the

[17] Cline et al., US 6,704,877 B2, issued Mar. 9, 2004 (Ex. 1017).

word "obvious," and because Dr. Sechen uses the term "inherently." PO Resp. 40–41 (citing Ex. 2027 ¶ 102). But according to PACT, Intel has not adequately established that a mobile device is inherent within Nicol's disclosure. *Id.* at 41 (citing Ex. 2027 ¶ 103). PACT also argues that Nicol does not mention or suggest mobile devices, that Intel did not cite Cline (Ex. 1017) as part of Intel's obviousness ground, and that Intel did not provide a motivation to combine Cline with Nicol. *Id.* at 41–42.

In its Reply, Intel asserts that its argument is based on obviousness, rather than inherency. Pet. Reply 13. Intel also contends that Dr. Mangione-Smith, after being asked during cross-examination whether "the benefits of minimizing power consumption are applicable to mobile computers," admitted that "the benefits of minimizing power consumption are valuable to a greater or less[er] degree in any sort of computing system." *Id.* at 13–14 (alteration in original) (emphasis omitted) (quoting Ex. 1051, 145:2–12) (citing Ex. 1055 ¶ 31).

We find Intel's argument persuasive. If a technique "has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way," then it would have been obvious to apply the technique according to that person's ordinary skill. *KSR*, 550 U.S. at 417. Intel has persuasively shown that a person of ordinary skill in the art would have recognized that Nicol's power-reduction features, as applied to a personal computer, would have improved similar well-known devices, such as mobile laptop computers, in the same way, by reducing their power consumption. As Intel points out, PACT "does not contest that mobile devices were well-known to a [person of ordinary skill in the art], []or that a laptop (a mobile computer) was a type of 'personal

computer' disclosed by Nicol." Pet. Reply 13 (citing Ex. 1008, 2:54–62, 6:39–41; Ex. 1051, 144:13–18).

We also find, given the relatively high level of ordinary skill in the art (e.g., a Master's degree in Electrical Engineering with three years of experience with processor design, *see supra* part III.A), that a person of ordinary skill would have had the ability to implement Nicol's teachings in a laptop computer.

Although Intel did not cite Cline as a prior art reference in its ground, Cline (in conjunction with Dr. Sechen's testimony and Dr. Mangione-Smith's admission, both of which we credit) is persuasive evidence directed to the background knowledge and perspective of a person of ordinary skill in the art at the time of the claimed invention. *See Yeda Research v. Mylan Pharm. Inc.*, 906 F.3d 1031, 1041 (Fed. Cir. 2018) ("[T]he Board can rely on evidence other than just prior art," including opinions and other evidence "indicating the level of ordinary skill in the art . . . and how one with ordinary skill in the art would have understood a prior art disclosure." (quoting *Dominion Dealer Sols., LLC v. AutoAlert, Inc.*, IPR2014-00684, 2014 WL 5035359, at *5 (PTAB Oct. 6, 2014))).

Thus, we determine that Intel has shown, by a preponderance of the evidence, that claims 18 and 33 would have been obvious to a person of ordinary skill in the art, in light of the teachings of Nicol, and are therefore unpatentable.

8. Claim 20

Claim 20 depends from claim 19 and further recites "a battery." Ex. 1003, 17:14–15. Relying on its arguments for claim 18 and testimony by Dr. Sechen, Intel argues that "Nicol discloses using multiprocessor chips in

consumer products such as laptop computers," and contends that "[a] laptop computer is a mobile device that includes a battery." Pet. 50 (citing Ex. 1001 ¶¶ 220–221).

PACT does not contest Intel's arguments. We credit Dr. Sechen's supporting testimony and find Intel's arguments persuasive for the reasons presented in the Petition, and for the reasons we discuss above for claims 18 and 33. *See supra* part III.C.7. Thus, we determine that claim 20 is unpatentable as obvious over Nicol.

9. Conclusion

Having considered the arguments and evidence presented in this *inter partes* review, we find Intel's arguments for its first ground persuasive as to claims 1, 5, 6, 17, 18, 19, 20, 22, 23, and 33. Thus, we determine that these claims are unpatentable. But we do not find persuasive Intel's arguments as to claims 2–4, 7, or 27; thus, we determine that Intel has not shown that these claims are unpatentable as obvious over Nicol.

D. GROUND BASED ON NICOL AND BHATIA

Intel alleges that claims 8–16, 21, 24, 28, and 30 are unpatentable under 35 U.S.C. § 103 as obvious over Nicol in view of Bhatia. Pet. 7. For the reasons below, we find Intel's arguments persuasive as to claims 8, 12, 14–16, 21, 24, 28, and 30, but not as to claims 9–11, or 13.

1. Overview of Bhatia

Bhatia describes a system including a processor "with a clock and a thermal management controller that monitors a temperature of the system"

and "varies the [processor] between different performance states." Ex. 1009, code (57).

As part of the background for describing performance states, Bhatia refers to "the Advanced Configuration and Power Interface (ACPI) Specification, Rev. 1.0, published on Dec. 22, 1996, which provides an interface between the operating system of a system and hardware devices to implement power and thermal management." Ex. 1009, 1:13–18. According to Bhatia, the ACPI specification defines two execution states for a CPU: a full-speed state and a throttled state in which the CPU clocks are active only for a programmed percentage of the time. *Id.* at 1:39–42. But Bhatia teaches that there is a disadvantage to operating in the throttled state, since it "may reduce overall system performance if the CPU remains in the throttled state for an extended period of time for system temperature control." *Id.* at 1:42–46.

Thus, Bhatia describes a "thermal management technique that is capable of maintaining system performance at a relatively high level." Ex. 1009, 1:47–49. Bhatia's system includes a high performance (HP) state, and a lower performance (LP) state in which the processor's core clock frequency and voltage are at a lower setting. *Id.* at 2:37–51. The system clock and thermal management controller "vary the [processor] between performance states." *Id.* at 1:54–58. The thermal management scheme is designed to increase the amount of time the processor operates in the HP state. *Id.* at 2:49–51.

Figure 6 of Bhatia, reproduced below, is a graph showing the relationship between temperature and power dissipation levels during thermal management:

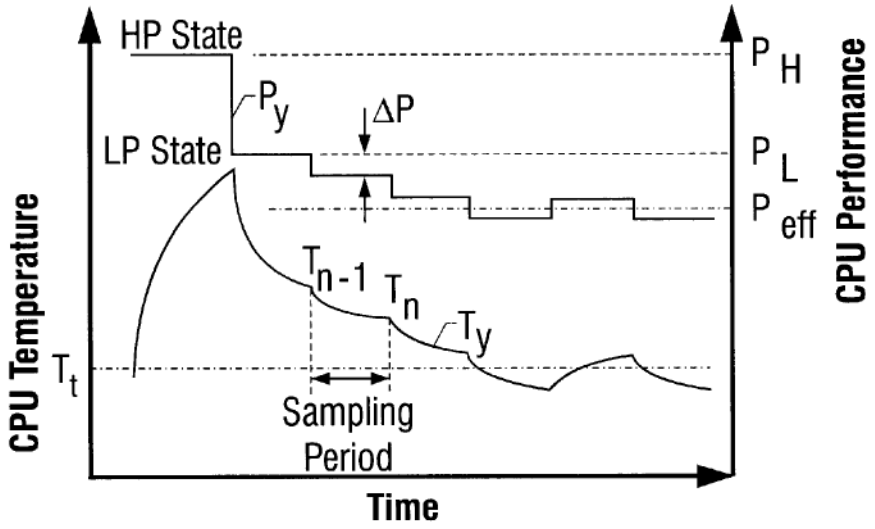


FIG. 6

Ex. 1009, 2:2–5. Figure 6, above, plots both CPU temperature T_y and CPU performance P_y versus time. While the processor is in the HP state, P_y is at a high power level P_H, while T_y gradually rises. When the processor transitions down to the LP state, P_y immediately decreases to low power level P_L, while T_y gradually decreases. Next, the processor is throttled down from the LP state, and P_y decreases stepwise, in decrements or increments of ΔP calculated each sampling period according to a control equation. *Id.* at 9:41–45; *see also id.* at 8:40–67. The value of P_y settles around an effective power dissipation level P_{eff}, which corresponds to target temperature T_t. *See id.* at 9:32–40.

2. Reason to combine Nicol and Bhatia

Intel argues that a person of ordinary skill in the art "would have been motivated to combine the complementary teachings of Nicol and Bhatia" for the following four reasons: (1) both references "are directed to interrelated thermal and power management aspects that are critical and ubiquitous in processor-based systems," Pet. 52–53 (citing Ex. 1008, code (54); Ex. 1009,

1:6–8; Ex. 1001 ¶¶ 231–233); (2) the references "both describe a flexible multiprocessor architecture in which software residing in a controller is programmed to selectively control the power supply of each processor, notably by dynamically adjusting the processor's clock frequency during operation," Pet. 53–54 (citing Ex. 1008, 2:18–27, 3:13–18; Ex. 1009, 1:57–59, 2:28–31, 3:23–28, 3:33–38, 12:26–37; Ex. 1001 ¶¶ 231–233); (3) since Nicol "does not discuss over-temperature protection," which is "critical for high-performance systems such as Nicol's multiprocessor system," a person of ordinary skill in the art "would have been motivated to consult Bhatia, which discusses well-known techniques of providing over-temperature protection," Pet. 54–55 (citing Ex. 1008, 3:9–20, 4:11–13; Ex. 1009, 2:18–22, 3:33–46; Ex. 1001 ¶¶ 231–233); and (4) a person of ordinary skill in the art "would have been motivated to consult Bhatia [to implement] low activity states in Nicol's system as . . . means to further minimize power consumption," Pet. 55 (citing Ex. 1009, 12:58–64; Ex. 1001 ¶¶ 231–233).

In its Response, PACT argues that Intel's four arguments "amount to little more than a suggestion that Nicol and Bhatia manage aspects of a computer system and it would be beneficial to combine these aspects," which, "without any showing of actual motivation is plainly insufficient," and "evidences application of impermissible hindsight bias." PO Resp. 43–44 (citing Ex. 2027 ¶ 110; *Polaris Indus., Inc. v. Arctic Cat, Inc.*, 882 F.3d 1056, 1068 (Fed. Cir. 2018)).

PACT also argues that Intel "makes 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." PO Resp. 44 (citing Ex. 2027 ¶ 111). Citing

Dr. Mangione-Smith's testimony, Patent Owner contends that a person of ordinary skill in the art "might spend their entire career developing intricate systems for managing processor temperature and clock frequency among a plurality of processors, and [Intel] dismisses it as a mere triviality." *Id.* at 44–45 (citing Ex. 2027 ¶ 111).

Lastly, PACT contends that "combining Nicol and Bhatia would frustrate the purpose of each," because "Nicol is designed to minimize power consumption by *lowering the clock frequency* and therefore the voltage demanded by the processors," whereas "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." *Id.* at 45 (citing Ex. 1008, code (57), 2:22–27; Ex. 1009, 1:47–49, 3:14–23, 6:41–51; Ex. 2027 ¶ 112).

In its Reply, Intel contends that PACT's arguments against combining Nicol and Bhatia are meritless and that they contradict Dr. Mangione-Smith's expert testimony. Pet. Reply 14–17. First, Intel argues that PACT is incorrect that the Petition merely describes the "benefits" of combining Nicol and Bhatia without showing actual motivation to combine. *Id.* at 14. According to Intel, the Petition explains, for example, that "a [person of ordinary skill in the art] seeking to improve Nicol would have been motivated to consult Bhatia, which discusses well-known techniques of providing over-temperature protection" and "low activity states . . . to further minimize power consumption." *Id.* at 14–15 (quoting Pet. 54, 55) (alteration in original). Intel contends that PACT failed to address these reasons and that, in fact, PACT's expert Dr. Mangione-Smith testified that

Nicol explicitly contemplates the need for and ability to address thermal management. *Id.* at 15 (citing Ex. 1052, 134:16–135:9).

Moreover, Intel argues, "[t]he Petition and Dr. Sechen explain how Bhatia discloses well-known techniques and thermal sensors to protect against over-temperature and thus addresses the exact thermal-management need undisputedly called for in Nicol." Pet. Reply 15–16 (citing Pet. 54–55; Ex. 1001 ¶¶ 230–233). Intel also argues that "there is no dispute that a [person of ordinary skill in the art] understood how to use such thermal sensors to manage temperature in Nicol's multiprocessor, as Dr. Mangione-Smith "*admitted* that a [person of ordinary skill in the art] 'concerned about overheating' Nicol's system would 'know of ways that he could do temperature management using a temperature sensor.'" *Id.* at 16 (quoting Ex. 1052, 129:23–130:9). According to Intel, "[t]he well-known benefits of combining the teachings of Nicol and Bhatia explained in the Petition, which stand unrebutted, provide ample reasons to combine." *Id.* (citing *E.I. du Pont De Nemours & Co. v. MacDermid Printing Sols., LLC*, 657 F. App'x 1004, 1014 (Fed. Cir. 2016); *Intermec, Inc. v. Alien Tech., LLC*, 664 F. App'x 962, 966 (Fed. Cir. 2016); Ex. 1055 ¶¶ 34–35).

Responding to PACT's contentions that the Petition does not "explain how a system combining Nicol and Bhatia would function" (PO Resp. 44), Intel argues that "[t]his is irrelevant, and [PACT] fails to cite any authority for the proposition that [Intel] must show *how* a combined system *would function* for a motivation to combine analysis. Pet. Reply 16–17 (citing *Facebook, Inc. v. Windy City Innovations, LLC*, 973 F.3d 1321, 1343 (Fed. Cir. 2020); *Elbrus Int'l Ltd. v. Samsung Elecs. Co.*, 738 F. App'x 694, 698–99 (Fed. Cir. 2018)). "Regardless," Intel argues, a person of ordinary skill in

the art "would have found it obvious to program Nicol's controller to handle the additional functionalities taught by Bhatia, yielding . . . predictable results as the implementation is primarily add-on software." *Id.* at 17 (citing Pet. 54–55; Ex. 1001 ¶ 233).

Finally, responding to PACT's contentions that the combination would "frustrate the purpose of each" reference and that the Petition fails to address hypothetical issues that would arise with such a combination (PO Resp. 45), Intel argues, "Not so. Bhatia's over-temperature protection system would not interfere with Nicol's task scheduling under normal operation," and according to Intel, a person of ordinary skill in the art "would have understood that Nicol's task scheduling would need to be altered or suspended when the multiprocessor system overheats, in which case Bhatia's software would take control to avoid device failure." Pet. Reply 17 (citing Ex. 1055 ¶ 36).

In its Sur-reply, PACT argues that the proposed combination of Nicol and Bhatia would not work and that a person of ordinary skill would not have had a reasonable expectation of success. PO Sur-reply 11. According to PACT, Nicol "describes a simplistic method for setting the clock frequency based on the anticipated load, by just determining a frequency so tasks will be completed within the required completion time," without considering "any potential overtemperature condition (or anything that interrupts task execution)," and "no adjustment of the clock frequency after the task is started." *Id.* at 12 (citing Ex. 1001 ¶¶ 100–101; Ex. 2032, 95:2–96:2). PACT contends that "Bhatia uses a different technique for determining clock frequency, instead cycling between a discrete set of states." *Id.* (citing Ex. 2032, 109:7–21, 115:17–116:8, 116:19–24).

In response to Intel's argument that Bhatia's overtemperature production system would not interfere with Nicol's task scheduling under normal operation and that Bhatia would take over once the overtemperature condition is met (Pet. Reply 17), PACT alleges that under Intel's proposal, "if Nicol's task execution were interrupted, the task would not finish on time." PO Sur-reply 12. According to PACT, a person of ordinary skill in the art "would not be motivated to combine Nicol and Bhatia in a way that failed to achieve Nicol's goal of completing tasks within a required completion time." *Id.* at 12–13.

Next, PACT contends that "[Intel]'s described system combining Nicol and Bhatia only works if we make assumptions counter to the stated motivation to combine, namely the need for overtemperature protection." PO Sur-reply 13. In particular, according to PACT, "[t]he only way the combined system would work properly is if the overtemperature condition, which is the motivation to combine, were not met and Nicol's task scheduling could proceed without any alteration or interference." *Id.*

PACT next argues that, "[i]n power management, there are many factors to consider in creating a system with a reasonable expectation of success," including a "need[] to consider the tradeoffs involved in selecting the right combination." PO Sur-reply 13–14 (citing Ex. 1015, 8-145 (ACPI specification); Ex. 2032, 97:20–98:7). PACT repeats from its Response that a person of ordinary skill in the art "might spend their entire career developing systems for managing processor temperature and clock frequency among a plurality of processors." *Id.* at 14 (citing Ex. 2022 ¶ 57); *see also* PO Resp. 28 (same).

Next, PACT argues that Intel's "proposed combination runs into the same problem for the low-activity states, C1, C2 and C3 states in the ACPI, referenced in Bhatia, because as noted in the ACPI and acknowledged by [Intel]'s expert, no instructions are executed in these states." PO Sur-reply 14 (citing Ex. 1015, 3-25, 8-145; Ex. 2032, 110:5–20). And PACT contends that the C2 and C3 states are only "optionally supported" in the ACPI, so a person of ordinary skill in the art "may not want to include all features in the ACPI." *Id.* at 14–15 (citing Ex. 1015, 8-146).

Finally, PACT argues that "Nicol only uses [its temperature sensors] for temperature tracking, not overtemperature. Temperature tracking for voltage-frequency control and overtemperature are different thermal management functions and the presence of one does not make the other one obvious." PO Sur-reply 15.

Considering the parties' arguments and evidence, we find that Intel has shown persuasively that a person of ordinary skill in the art would have been motivated to combine Nicol and Bhatia in the general manner that Intel proposes, and would have had a reasonable expectation of success in doing so. As Intel points out, Nicol expressly contemplates that temperature tracking would be desirable in its power management system, at least for the purpose of adjusting the supply voltage. *See, e.g.*, Ex. 1008, 4:11–13 (explaining that "the operating system can reduce the supply voltage even further by tracking temperature and process variations"); *see also* Ex. 1052, 134:25–135:9 (Dr. Mangione-Smith recognizing the same).

Although Nicol tracks temperature for a different purpose than Bhatia, Nicol teaches that temperature is a variable of interest in controlling a multiprocessing system. We credit Dr. Sechen's logical, clear and reasoned

testimony, supported by the record evidence, that a person of ordinary skill in the art would have been motivated to consult Bhatia and implement the disclosed low activity states in Nicol's system, for preventing an overtemperature condition. *See* Ex. 1001 ¶¶ 230–233). *See, e.g.*, *KSR*, 550 U.S. at 420 ("[A]ny need or problem known in the field . . . can provide a reason for combining the elements in the manner claimed."). We also find persuasive Intel's argument that it is undisputed in the testimonial evidence that a person of ordinary skill in the art would have understood how to manage temperature in Nicol's multiprocessor system. *See* Pet. Reply 16 (citing Ex. 1052, 129:23–130:9 (Dr. Mangione-Smith agreeing that a person of ordinary skill in the art would have known of ways to use a temperature sensor in implementing Nicol "[a]t least to the extent that the Pentium disclosed it, for example")).

Despite PACT's arguments to the contrary (*see* PO Resp. 45), we disagree that Nicol and Bhatia are incompatible in any relevant way. *See MCM Portfolio LLC v. Hewlett-Packard Co.*, 812 F.3d 1284, 1294 (Fed. Cir. 2015) ("[T]he test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference." (quoting *In re Keller*, 642 F.2d 413, 425 (CCPA 1981)); *Lear Siegler, Inc. v. Aeroquip Corp.*, 733 F.2d 881, 889 (Fed. Cir. 1984) (The ordinary artisan is not "compelled to adopt every single aspect of [a reference's] teaching without the exercise of independent judgment."). Contrary to PACT's arguments (*see* PO Sur-reply 12–13), we do not understand Intel to contend that implementing the low activity states of Bhatia would entail cycling between states with different clock frequencies in the normal operation of Nicol's system, but rather that such cycling would

occur, if at all, only when the multiprocessor system is determined to be overheating. *See* Pet. Reply 17 (citing Ex. 1055 ¶ 36).

We also disagree with PACT's arguments that an interruption of Nicol's task execution for overtemperature protection would cause the task not to finish on time, and that "[t]he only way the combined system would work properly is if the overtemperature condition, which is the motivation to combine, were not met and Nicol's task' scheduling could proceed without any alteration or interference," and that this would have dissuaded a person of ordinary skill in the art from combining Nicol and Bhatia. *See* PO Sur-reply 12–13. Notwithstanding PACT's arguments, it is not credible that a person of ordinary skill in the art would have understood either Nicol or Bhatia to prefer device failure over any alteration or suspension of task scheduling, even if the latter potentially would result in tasks not being completed "within the required times." Rather, we credit Dr. Sechen's testimony that a person of ordinary skill in the art would have understood that Nicol's task scheduling would need to be altered or suspended to avoid chip failure due to overheating. Ex. 1050 ¶ 36; *see* Pet. Reply 17.

Finally, notwithstanding Dr. Mangione-Smith's testimony that a person of ordinary skill in the art "might spend their entire career developing systems for managing processor temperature and clock frequency among a plurality of processors" (Ex. 2027 ¶ 111), the challenged claims recite only modifying clock frequencies based on temperature measurement, not a perfect optimization of such parameters. *See* Ex. 1003, 16:15–24, 16:28–39. Thus, we disagree with PACT's arguments, based on that testimony, that a person of ordinary skill in the art would not have had a reasonable expectation of success in implementing the combination of Nicol and Bhatia,

or that the combination would have required undue experimentation and would not have yielded predictable results. *See* PO Resp. 44–45; PO Sur-reply 13–14. As the Federal Circuit has explained, "[t]he reasonable expectation of success requirement refers to the likelihood of success in combining references to meet the limitations of the claimed invention." *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367 (Fed. Cir. 2016).

For the above reasons, we are persuaded that a person of ordinary skill in the art would have had a reason to combine Nicol with Bhatia, and would have reasonably expected that the combination would achieve the claimed invention as recited in claims 8–16, 21, 24, 28, and 30.

3. Claims 8 and 24

Claims 8 and 24 depend from claims 1 and 23, respectively, and further recites "a heterogeneous plurality of clocked data processing units." Ex. 1003, 16:10–11, 17:25–26. For claim 24, Intel relies on its arguments for claim 8. *See* Pet. 71 (citing Ex. 1001 ¶ 293).[18]

Intel makes two alternative arguments. First, Intel argues that "Nicol alone renders obvious this claim element." Pet. 55 (citing Ex. 1001 ¶¶ 234–241). Intel argues that "nothing in Nicol would require [the controller PE and the other PEs] to be all of the same type," and a person of ordinary skill in the art would have understood that the PEs could be heterogeneous, because Nicol teaches that the controller PE "must be suitable to support a real time operating system (RTOS) and a task scheduler while the other PEs do not have such requirements and can be simpler PEs." Pet. 55–56 (citing

[18] PACT does not appear to address claim 24, specifically, in its Response.

Ex. 1008, 3:9–18, 5:66–6:12; Ex. 1001 ¶¶ 235–241). Intel also points out that in Nicol's Figure 4, the controller PE "has more interfaces to the asynchronous communication network 160 compared to other PEs, and it drives multiple PLLs, unlike the other PEs." *Id.* (citing Ex. 1008, Fig. 4; Ex. 1001 ¶¶ 235–237).

Alternatively, Intel argues that a person of ordinary skill in the art would have looked to Bhatia for this limitation, because "Bhatia discloses a multiprocessor system that includes processors of different types, including general purpose processors, special-purpose processors, and controllers such as bridge controllers and peripheral controllers." Pet. 57–58 (citing Ex. 1009, 3:25–33, 3:60–65; Ex. 1001 ¶¶ 231–232).

In response, PACT characterizes Intel's first argument, based on Nicol alone, as simply showing a "possibility" that Nicol's processors might be heterogeneous "by the lack of disclosure of a prohibition" of there being heterogeneous processors. PO Resp. 46 (citing Ex. 2027 ¶ 114; *PersonalWeb Techs., LLC v. Apple, Inc.*, 917 F.3d 1376, 1382 (Fed. Cir. 2019)). PACT also makes the following arguments: (1) that all the PEs in Nicol's Figure 2 operate at the same clock frequency, *Id.* (citing Ex. 2027 ¶ 114); (2) that Nicol's task scheduler "is a software component and is part of the operating system," which can be run on any sufficient PE, *id.* at 46–47 (citing Ex. 1008, 6:7–8; Ex. 2027 ¶ 115); (3) that operating at different frequencies and operating voltages, or the use of "other system arrangements," such as memory arrangements, does not mean the processors are heterogeneous, *id.* at 47 (citing Ex. 1008, 6:52–56; Ex. 2027 ¶ 118); and (4) that even though Hayes (Ex. 1011) is a heterogeneous system, Intel is not citing Hayes as prior art in combination with Nicol, *id.* at 47–48 (citing Ex. 2027 ¶ 119).

As to Intel's argument combining Nicol with Bhatia, PACT contends that Intel has failed to demonstrate a motivation to combine, or that an ordinarily skilled artisan would have had a reasonable expectation of success in making the combination. PO Resp. 48; *see also id.* at 49–50 (citing Ex. 2027 ¶¶ 119, 121–122, 125; *TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1066 (Fed. Cir. 2016)) (arguing that Intel has not specifically tailored its motivation-to-combine argument for claim 8). PACT characterizes Intel's arguments as merely, and insufficiently, arguing that Nicol "could" be used in a heterogeneous system. *Id.* at 48–49 (citing Pet. 55–57; Ex. 2027 ¶ 120; *Belden Inc. v. Berk-Tek LLC*, 805 F.3d 1064, 1073 (Fed. Cir. 2015); *In re Stepan Co.*, 868 F.3d 1342 (Fed. Cir. 2017)); *see also id.* at 50–51 (citing Ex. 2027 ¶ 126).

PACT also contends that Intel and Dr. Sechen "cite no evidence or underlying data as to what Nicol's system arrangements are, and do not explain why Bhatia's so-called 'heterogeneous multiprocessor system' falls within such an arrangement," so Dr. Sechen's "testimony should be given no weight." PO Resp. 50 (citing 37 C.F.R. § 42.65(a)).

In its Reply, Intel argues that "Nicol makes clear that the controller of the system needs not be a processor of the same type as the computational PEs," and contends that during cross-examination, PACT's expert Dr. Mangione-Smith admitted that a person of ordinary skill in the art "would have recognized that the controller and the computational PEs in Nicol's system have 'fundamentally different' roles and thus, while a 'general purpose' processor can be used as a controller, a 'customized DSP' is better suited as [the] computational PE." Pet. Reply 18 (quoting Ex. 1051, 147:24–148:17) (citing Ex. 1055 ¶¶ 39–40).

Responding in its Sur-reply, PACT contends that Dr. Mangione-Smith's "admissions" were "no more than answers to hypotheticals about what is possible in light of the very bare disclosure." PO Sur-reply 15 (citing Ex. 1008, 2:52–53).

Considering the arguments on both sides, and the submitted evidence, we find Intel's argument persuasive that a person of ordinary skill in the art, at least in light of Bhatia's teachings of a heterogeneous multiprocessing system, would have had reason to use a different type of processor for Nicol's controller PE than for the other PEs. In particular, we agree with Intel that a person of ordinary skill in the art would have understood that Nicol's statement that the controller PE "may be just another processing element" suggests using a controller PE that differs from the other PEs. *See* Ex. 1008, 6:7–9. We also find Intel's argument for combining Bhatia to be sufficiently specific and persuasive, given the relatively high level of ordinary skill in the art. *See KSR*, 550 U.S. at 417 ("[I]f a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill."); *see also id.* at 416 ("[W]hen a patent claims a structure already known in the prior art that is altered by the mere substitution of one element for another known in the field, the combination must do more than yield a predictable result."). We credit Dr. Sechen's testimony that implementing a heterogeneous system was within the level of ordinary skill in the art at the time of the claimed invention, and that Nicol's PEs could be altered by substituting with different processor types in a predictable way. *See* Ex. 1055 ¶¶ 38–40. And supporting his testimony, we find it notable that that

inventors of the ’047 patent evidently did not consider it necessary to teach an ordinarily skilled artisan any significant details about how to incorporate a set of heterogeneous processors within the claimed invention. *Id.* ¶ 38.

For the above reasons, we determine that Intel has shown, by a preponderance of the evidence, that claims 8 and 24 would have been obvious over Nicol to a person of ordinary skill in the art, at least in light of the teachings of Bhatia, and thus the claims are unpatentable. Because this is dispositive, we do not need to determine whether the claims would have been obvious over Nicol alone.

4. Claims 9, 11, and 13

Claims 9, 11, and 13 depend from claims 8, 10, and 12, respectively, and further require that the system of the parent claim is “adapted to subsequently reduce the clock frequency in accordance with a hysteresis characteristic.” Ex. 1003, 16:12–14, 16:25–27, 16:40–42.

Intel relies on the same arguments it made for claim 7. *See* Pet. 58, 61, 66; Pet. Reply 8–13. And in response, PACT also relies on the same arguments it made for claims 7 and 27. *See* PO Resp. 32–40; PO Sur-reply 9–11. We address these arguments above. *See supra* part III.C.6. Also, claim 11 depends from claim 10 and, as we discuss below, we determine that Intel has not shown that claim 10 is obvious over Nicol in view of Bhatia. *See infra* part III.D.5.

For the reasons given above for claims 7 and 27 and below for claim 10, we determine that Intel has not shown, by a preponderance of the evidence, that claims 9, 11, or 13 are unpatentable as obvious over Nicol in view of Bhatia.

5. Claim 10

Claim 10 is as follows:

> 10. The system of claim 1, further comprising:
>
> a plurality of temperature sensors and a heterogeneous plurality of clocked data processing units; and wherein
>
> the multiprocessor having a plurality of regions, a temperature sensor being provided for each of said plurality of regions to measure the temperature of said specific region; and
>
> the clock frequencies of said data processing units being dynamically adjustable in accordance with the sensed temperatures.

Ex. 1003, 16:15–24.

Intel contends that each of these limitations would have been obvious over Nicol in view of Bhatia. Pet. 58, 60 (citing Ex. 1001 ¶¶ 244, 251–253). First, Intel relies on its arguments (addressed above) that claim 8 teaches "a heterogeneous plurality of clocked data processing units." Pet. 58 (citing Ex. 1001 ¶¶ 244–250). PACT does not further address this part of claim 10 in its Response. *See* PO Resp. 51–53. Thus, we find Intel's arguments for this portion of claim 10 persuasive for same reasons as our analysis above regarding claim 8. *See supra* part III.D.3.

For the limitation "a plurality of temperature sensors" and the other recited limitations, Intel argues that Bhatia discloses a multiprocessor system in which "[o]ne or more temperature sensor units . . . monitor system temperature in one or more corresponding thermal zones," and where each sensor is "capable of issuing an interrupt . . . or some other notification when a sensed temperature rises above a preset target temperature T_t or falls below

the target temperature T_t." Pet. 59 (alterations in original) (emphasis omitted) (quoting Ex. 1009, 3:32–39). According to Intel, Bhatia also teaches that "thermal management in the system . . . may be performed independently *for multiple thermal zones*." Pet. 60 (quoting Ex. 1009, 3:64–65). And Intel argues that Bhatia discloses "dynamically adjusting the clock frequency of a processing unit in accordance with the sensed temperatures." Pet. 60–61 (citing Ex. 1009, code (57), 1:7–13, 2:24–28, 12:5–10).

Although Intel concedes that "Nicol does not expressly mention the use of a plurality of thermal sensors in its chips," Intel contends that Nicol discloses the use of temperature tracking for setting the optimal voltage. Pet. 59 (citing Ex. 1008, 1:41–48, 4:11–13). Citing testimony of Dr. Sechen, Intel also argues that "[t]he use of temperature sensor(s) for thermal and power management was in fact well-known in the art." *Id.* (citing Ex. 1009, 1:6–17; Ex. 1001 ¶¶ 245–249).

According to Intel (citing Dr. Sechen), a person of ordinary skill in the art would have understood that a thermal zone is the same as a "region" recited in claim 10. Pet. 60 (citing Ex. 1001 ¶¶ 252–253). Intel argues that a person of ordinary skill in the art "would have understood that Bhatia's approach of partitioning a system or chip into multiple thermal zones, each monitored by its own thermal sensor, is beneficial to Nicol's multiprocessor chip because the different PEs of the chip can have different workloads and thus different temperature profiles." Pet. 59 (citing Ex. 1008, 5:41–44; Ex. 1001 ¶ 249). According to Intel, "[p]lacing multiple thermal sensors, for example one sensor on or near each of the PEs, would improve the accuracy of temperature tracking and improve voltage scaling as a function of temperature discussed in Nicol." Pet. 59–60 (citing Ex. 1001 ¶ 250).

In response, PACT argues that, "[w]hile Bhatia does discuss 'thermal zone[s] of [a] computer system,' there is no suggestion of thermal zones and thermal sensors to monitor zones of a chip." PO Resp. 52–53 (alterations in original) (quoting Ex. 1009, 12:9) (citing Ex. 1009, code (57), 15:46–48, Fig. 1). According to PACT, "the task of utilizing multiple temperature sensors on a single chip is significant and not to be dismissed as trivial or obvious." *Id.* at 53 (citing Ex. 2027 ¶ 129). Citing testimony of Dr. Mangione-Smith, PACT contends that "[m]onitoring the temperature of zones of a chip would require specialized temperature sensors that Bhatia does not disclose," and there are numerous other issues such as sizing, precision, and geometry, for which Intel has not accounted," and for which "the cost would be significant." *Id.* at 52 (citing Ex. 2027 ¶ 128).

In its Reply, Intel clarifies that it "does not suggest taking Bhatia's physical temperature sensors and placing them on Nicol's system—nor is such a combination required to render this claim limitation obvious." Pet. Reply 19. Rather, according to Intel, a person of ordinary skill in the art "would have understood that Bhatia's approach of partitioning a system or chip into multiple thermal zones, each monitored by its own thermal sensor, is beneficial and readily[] applicable to Nicol's multiprocessor system because the different PEs (or groups of PEs) can have different workloads and thus different temperature profiles." *Id.* at 19–20 (citing Pet. 59; Ex. 1009, 3:29–39; Ex. 1051, 160:21–161:13; Ex. 1052, 129:23–130:9; Ex. 1055 ¶ 42).

In its Sur-reply, PACT points out that Intel did not, in its Reply, directly address the issue of whether a person of ordinary skill in the art would have been motivated and enabled to include multiple temperature

sensors on a single chip, with a reasonable expectation of success. PO Sur-reply 16–17.

We agree with PACT that Intel has not persuasively shown that a person of ordinary skill in the art would have modified Nicol by adding a plurality of temperature sensors on Nicol's single chip. Neither Nicol nor Bhatia teaches multiple sensors on a single chip. Rather, Dr. Sechen merely opines that it would be beneficial to monitor multiple regions of Nicol's chip as taught by Bhatia. Ex. 1055 ¶ 42. But Dr. Sechen fails to direct us to where Bhatia describes multiple temperature sensors on a chip as alleged. Moreover, Dr. Sechen fails to direct us to supporting evidence that "a POSITA would have understood that the temperature regions of Nicol's chip would differ" such that multiple temperature sensors on a single chip "would be beneficial" as alleged. Dr. Sechen's testimony in that regard is unsupported by record evidence.

For the reasons given above, we determine that Intel has not shown, by a preponderance of the evidence, that claim 10 is unpatentable as obvious over Nicol in view of Bhatia.

6. Claim 12

Claim 12 is as follows:

> 12. The system of claim 1, adapted to:
>
> (i) setting a clock frequency of at least a part of the multiprocessor system of a value in accordance with a first processor load, the value being less than a maximum value;
>
> (ii) subsequently increasing the clock frequency of the at least a part of the multiprocessor system to a maximum in accordance with a second processor load; and

> (iii) subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with an operating temperature threshold for preventing overtemperature.

Ex. 1003, 16:28–39.

For claim 12, Intel relies, as it did for claim 1, on Nicol's Figure 4 embodiment in light of Figure 3, and Nicol's teaching that the operating system may "independently program the lowest operating frequency and corresponding lowest supply voltage for each PE." *See* Pet. 61–62 (quoting Ex. 1008, 5:67–6:2) (citing Ex. 1008, 5:66, 6:14–18, 6:63–66, Fig. 3).

Referring to annotated Figure 3, Intel identifies limitation (i) as when the controller sets the clock frequency of 70 MHz before a newly created task begins to execute (i.e., the "first processor load"), which is less than the "maximum value" of 140 MHz. Pet. 62–63 (citing Ex. 1008, 2:53–64, 3:13–18, 3:54–4:5, Fig. 3; Ex. 1001 ¶¶ 257–262). Intel identifies limitation (ii) as when the controller increases the frequency to the maximum of 140 MHz after the task begins execution (i.e., the "second processor load"). Pet. 63–64 (citing Ex. 1008, 2:65–67, 3:54–65, Fig. 3; Ex. 1001 ¶¶ 269–270).

For limitation (iii), Intel relies on Bhatia rather than Nicol's Figure 3, and also relies on its general motivation to combine Nicol with Bhatia in the context of overtemperature protection (discussed above, *supra* part III.D.2), as well as the more specific argument that, in response to an overtemperature condition, a person of ordinary skill in the art would have reduced the clock frequency from its 140 MHz maximum by causing the processor to enter its LP (low power) state. Pet. 64–65 (citing Ex. 1008, code (19), code (57); Ex. 1009, 1:19–21, 2:24–28, 4:34–37, 12:6–10; Ex. 1001 ¶¶ 267–270).

In response, PACT argues, first, that Intel relies on Nicol's Figure 4 embodiment for claim 1, but on the Figure 2 embodiment for claim 12. *See* PO Resp. 54–55, 58–59 (citing Ex. 2027 ¶¶ 132–133, 138–139); *see also* PO Sur-reply 17–18 (arguing that, as with claim 1, Intel has failed to combine Nicol's Figure 2 and Figure 4 embodiments). We disagree. As Intel clarifies in its Reply, Intel relies on the Figure 4 embodiment for both claims 1 and 12, and also relies on Nicol's teachings relating to Figure 3, which apply generally to each of Nicol's embodiments. *See* Pet. Reply 20 (citing Pet. 23, 61; Ex. 1008, 4:54–66). As we discuss in the context of claim 1, we find Intel's argument for applying the teachings of Figure 3 to the Figure 4 embodiment persuasive. *See supra* part III.C.2(h).

Second, PACT argues that Intel has failed to identify a "second processor" that is different from a "first processor." PO Resp. 56–57 (citing Ex. 2027 ¶¶ 134–136). This assumes that when claim 12 refers to a "first processor load" and a "second processor load," that the words "first" and "second" modify the word "processor," rather than the word "load." This is an unnatural and illogical interpretation of claim 12. If we were to accept PACT's proposed interpretation, there would be no antecedent basis in either claim 12 or its parent claim 1 for the terms "first processor" and "second processor" before the claim refers to their "loads." (Claim 1 refers to "data processing units," not "processors." *See* Ex. 1003, 15:39.)

The manipulations to clock frequency in claim 12 occur on "at least a part of the multiprocessor system"—and this is the same part associated with both the "first processor load" and the "second processor load." *See* Ex. 1003, 16:29–30, 16:33–34. By its literal terms, a "part of the multiprocessor system" could include a single processor. We disagree with PACT's

contention that Intel's straightforward interpretation of claim 12 reads out the word "processor," because the word does not function in claim 12 as a noun identifying a specific processor, but as an adjective specifying the type of "load." Thus, we disagree with PACT's claim interpretation argument.

Alternatively, in its Reply, Intel argues that even under PACT's interpretation of "first processor load" and "second processor load," Nicol teaches this limitation because as Dr. Mangione-Smith admitted in cross-examination, in Nicol's "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." Pet. Reply 21 (quoting Ex. 1051, 69:7–16) (citing Ex. 1055 ¶ 46). We find Intel's argument persuasive. Thus, even under PACT's interpretation, there would be two processors whose loads meet the limitations of claim 12 in Nicol's "middle ground" embodiment.

We also disagree with PACT's argument, in the Sur-reply (PO Sur-reply 18), that Intel's alternative argument is untimely, because in its Petition, Intel gave fair notice of its contention that the teachings associated with Nicol's Figure 3 apply to Nicol's other embodiments, including groups of processors as in Nicol's "middle ground." *See, e.g.*, Pet. 15 (Figure 3 illustrates "adjustments of one or more PEs (or of the entire chip)."), Pet. 17 ("the controller . . . allocates [a task] to one or more available PEs [and] raises the supply voltage of those PEs"); Pet. 37 ("Nicol teaches a range of control strategies in which adjustments can be done selectively for each PE, can be common to a group of PEs, or can be common to all PEs of a chip.").

Third, PACT argues that Nicol "teaches away from using the maximum frequency of the chip that [Intel] relies on," because Nicol

provides "multiple examples . . . for running at frequencies below the maximum (70 MHz instead of 140 MHz; 80 MHz instead of 160 MHz) and does not disclose or suggest 'increasing the clock frequency . . . to a maximum in accordance with a second processor load.'" PO Resp. 57 (quoting Ex. 1003, 16:33–35) (citing Ex. 1008, 3:64–65; Ex. 2027 ¶ 137).

We disagree. As Intel correctly argues, Nicol teaches that a PE may operate at "full speed." Pet. Reply 21 (citing Ex. 1008, 2:65–67 ("For 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. 1055 ¶ 47). Intel also notes, correctly, that PACT "fails to support its allegation of 'teaching away' because [PACT] has not provided any evidence that Nicol discredits or criticizes running PEs at their maximum frequency." *Id.* (citing *Livanova, Inc. v. Neuro and Cardiac Techs., LLC*, IPR2019-00264, Paper 32 at 38–39 (PTAB May 18, 2020)).

Finally, with regard to limitation (iii) of claim 12, PACT argues that Intel "fails to provide sufficient motivation to suggest Nicol and Bhatia can be combined," and also that "Nicol discusses voltage scaling based on temperature variations, but not clock frequency changes. A reduction in clock frequency due to temperature would prevent tasks from being completed on time." PO Resp. 59 (citing Ex. 2027 ¶ 140). We address PACT's arguments against combining Nicol with Bhatia above, including the argument that Nicol's reference to temperature is not in the context of setting the clock frequency and the argument that Bhatia's teachings would interrupt a task and prevent it from completing on time, and we disagree here for the same reasons. *See supra* part III.D.2.

Thus, we determine that Intel has shown, by a preponderance of the evidence, that claim 12 would have been obvious over Nicol, in combination with the teachings of Bhatia, and that claim 12 is therefore unpatentable.

7. Claims 14 and 28

Claims 14 and 28 depend from claims 1 and 20, respectively, and further recite "wherein a low power operation mode is provided such that a first data processing unit is active while a plurality of other data processing units are deactivated." Ex. 1003, 16:43–46, 17:53–56.

Intel argues that Nicol teaches this limitation (i.e., keeping a first data processing unit while deactivating a plurality of other data processing units) because the operating system on Nicol's controller PE "employs gated clocking techniques to '*shut down*' PEs . . . that are not needed." Pet. 66 (quoting Ex. 1008, 5:42–45) (citing Ex. 1008, 10:34–36 ("comprising means within said controller for turning supply power off to at least one of said N processing elements")); *see also* Pet. 72 (citing Ex. 1001 ¶ 295) (claim 28, relying on the arguments for claim 14). Citing uncontested testimony of Dr. Sechen, which we credit, Intel argues that "[d]eactivating inactive processing units to reduce power was well-known in the art and commonly implemented in processing systems." Pet. 67 (citing Ex. 1001 ¶ 278).

Intel alternatively relies on Bhatia for teaching this limitation by disclosing that a processor can enter a reduced power mode in which "an internal clock of the processor is stopped and most functions are disabled." Pet. 67 (emphasis omitted) (quoting Ex. 1009, 4:20–21) (citing Ex. 1009, 4:16–21, 12:58–65; Ex. 1001 ¶¶ 267–268, 274–279).

In response, PACT contends that "Nicol only states that PEs can be clock gated, but does not state whether there is an active PE when other PEs

are clock gated," which relies on a "lack of disclosure of an opposite scenario" to establish a claim limitation. PO Resp. 60–61 (citing *PersonalWeb Techs.*, 917 F.3d at 1382; Ex. 2027 ¶ 142).

We disagree with PACT's argument. As Intel points out in its Reply, Nicol "expressly states that 'the operating system [in the controller] employs gated clocking techniques to 'shut down' PEs *that are not needed*.'" Pet. Reply 22 (alteration in original) (quoting Ex. 1008, 5:43–45); *accord* Pet. 66. We agree with Intel that a person of ordinary skill in the art would have understood that shutting down PEs "that are not needed" also suggests keeping the "needed" PEs active. *See id.* We also credit Dr. Sechen's testimony that a person of ordinary skill in the art "would have understood that the controller remains activated so that it can process and distribute the new tasks entering its 'to-do' list and 'wake up' the clock-gated PEs should the need arise." Ex. 1055 ¶ 50; *accord* Pet. Reply 22–23.

Thus, we determine that Intel has shown, by a preponderance of the evidence, that claims 14 and 28 would have been obvious over Nicol alone, and that claims 14 and 28 are thus unpatentable. Because this is dispositive, we do not need to address Intel's alternative argument relying on the combination of Nicol and Bhatia.

1. Claim 15

Claim 15 depends from claim 14 and further recites "wherein the first data processing unit is adapted to activate at least some of the other data processing units for data processing in response to an event." Ex. 1003, 16:47–50.

Intel argues that Nicol discloses this limitation because when some PEs are clock-gated, "[t]he controller PE ('the first data processing unit') is

adapted to re-activate the clock-gated PEs when it receives new tasks ('event') that require additional processing resources." Pet. 68–69 (quoting Ex. 1008, 2:65–67) (citing Ex. 1001 ¶¶ 281–285).

PACT does not specifically contest the limitation of claim 15. *See* PO Resp. 59–61. We find Intel's arguments for claim 15 persuasive for the reasons given in the Petition, and based on Dr. Sechen's supporting testimony, which we credit. *See* Pet. 68–69; Ex. 1001 ¶¶ 281–285.

Thus, we determine that Intel has shown, by a preponderance of the evidence, that claim 15 would have been obvious over Nicol alone, and is therefore unpatentable. Because this is dispositive, we do not need to address Intel's alternative argument (based on parent claim 14, *see supra* part III.D.7) relying on the combination of Nicol and Bhatia.

2. Claims 16 and 30

Claims 16 and 30 depend from claims 1 and 28, respectively, and further recite "wherein in the low power mode the minimum clock rate of the plurality of other units is set in accordance [sic] clock rate necessary for preserving the memory contents." Ex. 1003, 16:51–54, 17:60–63.

Intel argues that Nicol teaches independently minimizing power consumption for each PE, and employing gated clocking techniques to "shut down" PEs that are not needed. Pet. 69 (citing Ex. 1008, 5:42–45, 5:66–6:2). Intel also relies on Bhatia's similar teaching of placing processors into a low power mode, such as the C3 state as defined in the ACPI specification, whereby "the external clock BCLK to the processor is disabled so that no activity is performed by the processor *except maintenance of the stored data in the processor's internal cache*." Pet. 70 (quoting Ex. 1009, 12:58–65) (citing Ex. 1009, 1:6–17). Given that ACPI was a well-known standard, Intel

contends that a person of ordinary skill in the art would "have been motivated to implement a low power mode such as ACPI's C3 state in any computer system, including Nicol's system, because such a mode preserves memory contents, which serves to speed up the computer's system transitioning back to full operation upon 'waking up' from a low power mode." Pet. 70 (citing Ex. 1001 ¶ 290); *see also* Pet. 72 (citing Ex. 1001 ¶¶ 288–290) (claim 30, relying on arguments for claim 16).

In response, PACT argues, first, that when Nicol's PEs are inactive, "V_{dd}-local [is] reduced [to a point when] the circuits start to fail." PO Resp. 62 (alterations in original) (quoting Ex. 1008, 5:42–45). Thus, according to PACT, a person of ordinary skill in the art "would have understood that when the circuits start to fail, the voltage is too low to 'preserv[e] the memory contents' as recited in claims 16 and 30." *Id.* (citing Ex. 2027 ¶ 145) (alteration in original).

In its Reply, Intel contest's PACT's argument that Nicol's clock-gating is directed to avoiding circuit failure. Pet. Reply 24. According to Intel, "The language [PACT] cites 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)." *Id.* (citing Ex. 1008, Fig. 1; Ex. 1055 ¶ 56). Thus, Intel argues that "when a PE is idling, reducing its V_{dd} supply voltage below its minimum voltage to save power is not appropriate and the PE should instead by clock-gated. *Id.* (citing Ex. 1055 ¶ 56).

We find Intel's argument and explanation persuasive that Nicol does not teach reducing V_{dd}-local to the point of processor failure, and we credit

the supporting testimony of Dr. Sechen (Ex. 1055 ¶ 56). Thus, we disagree with PACT's first argument.

Second, PACT argues that Dr. Sechen's opinion that that a person of ordinary skill in the art would have been motivated to implement a low power mode, such as ACPI's C3 state in any computer system, including Nicol's system, amounts to stating that "Nicol must meet this industry standard." PO Resp. 62. PACT contends that this is a conclusory statement, and that a person of ordinary skill in the art "would have known that this standard is not required for all computers," including many computer systems existing at the time of the claimed invention. *Id.* at 62–63 (citing Ex. 2027 ¶ 146); *see also* PO Sur-reply 19 (arguing that in the ACPI specification, the C3 state is optional, and that not every personal computer would implement it).

We disagree. Intel's argument does not depend on the ACPI C3 state being implemented in *every* personal computer at the time of the claimed invention. As Intel argues in its Reply, its argument is that the ACPI C3 state was well known in the art at the time of the claimed invention, and that a person of ordinary skill in the art would have had reason to select it for the low-power state as Bhatia suggests. *See* Pet. Reply 23–24 (citing Ex. 1008, 6:41; Ex. 1015, 1-1; Ex. 1052, 143:19–24 (Dr. Mangione-Smith stating that ACPI specification was well known); Ex. 1055 ¶ 55). We find Intel's argument persuasive.

Based on the evidence of record, including the supporting testimony of Dr. Sechen, which we credit, we determine that Intel has shown, by a preponderance of the evidence, that claims 16 and 30 would have been

obvious over Nicol in light of the teachings of Bhatia, and are thus unpatentable.

3. Claim 21

Claim 21 depends from claim 20 and further recites "a heat sink for the multiprocessor system." Ex. 1003, 17:16–17.

According to Intel, "Bhatia explains that a high performance state of the processor may need to be reduced because of 'failure of one or more fans or other *heat dissipation devices in the system*.'" Pet. 71 (citing Ex. 1009, 6:59–60). Citing testimony by Dr. Sechen, Intel argues that a person of ordinary skill in the art "would have understood that a heat sink is a type of heat dissipation device which is well known in the art." *Id.* (citing Ex. 1001 ¶ 292).

PACT does not contest Intel's arguments. We credit Dr. Sechen's supporting testimony that a heat sink was a well-known heat dissipation device, and find Intel's arguments persuasive for the reasons presented in the Petition. *See KSR*, 550 U.S. at 416 ("[W]hen a patent claims a structure already known in the prior art that is altered by the mere substitution of one element for another known in the field, the combination must do more than yield a predictable result."). Thus, we conclude that claim 21 is unpatentable as obvious over Nicol in view of Bhatia.

4. Conclusion

Having considered the arguments and evidence presented in this *inter partes* review, we find Intel's arguments for its second ground persuasive as to claims 8, 12, 14–16, 21, 24, 28, and 30, and thus we determine that Intel has shown, by a preponderance of the evidence, that these claims are

unpatentable as obvious over Nicol in light of Bhatia. But we determine that Intel has not shown, by a preponderance of the evidence, that claims 9–11, or 13 are unpatentable as obvious on this ground.

E. GROUNDS BASED ON NICOL, BHATIA, AND DEHON

Grounds 3 and 4 mirror grounds 1 and 2, respectively, except that Intel has added DeHon as an additional reference, based on a claim construction that Intel proposed in the co-pending litigation. *See* Pet. 13. In its Reply, Intel agreed "to narrow the issues before the Board" by withdrawing its Grounds 3 and 4 in light of the district court's claim construction in the co-pending litigation. Pet. Reply 24–25.

Based on Intel's representation that it has withdrawn Grounds 3 and 4, we do not need to address those grounds.

IV. CONCLUSION[19]

For the reasons above, Intel has shown by a preponderance of the evidence that 1, 5, 6, 8, 12, 14–24, 28, 30, and 33 of the '047 patent are unpatentable, but has not shown that claims 2–4, 7, 9–11, 13, or 27 are unpatentable, as summarized in the following table:

Claims	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not Shown Unpatentable
1–7, 17–20, 22, 23, 27, 33	103(a)	Nicol	1, 5, 6, 17–20, 22, 23, 33	2–4, 7, 27
8–16, 21, 24, 28, 30	103(a)	Nicol, Bhatia	8, 12, 14–16, 21, 24, 28, 30	9–11, 13
1–7, 17–20, 22, 23, 27, 33	103(a)	Nicol, DeHon[20]		
8–16, 21, 24, 28, 30	103(a)	Nicol, Bhatia, DeHon		
Overall Outcome			1, 5, 6, 8, 12, 14–24, 28, 30, 33	2–4, 7, 9–11, 13, 27

[19] Should PACT wish to pursue amendment of claims in a reissue or reexamination proceeding after this decision, we draw PACT's attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. *See* 84 Fed. Reg. 16,654 (Apr. 22, 2019). If PACT chooses to file a reissue application or a request for reexamination of the challenged patent, we remind PACT of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. *See* 37 C.F.R. § 42.8(a)(3), (b)(2) (2020).

[20] As we explain above, we do not reach either of the two withdrawn grounds that include DeHon as a reference. *See supra* part III.E.

V. ORDER

In consideration of the foregoing, it is

ORDERED that claims 1, 5, 6, 8, 12, 14–24, 28, 30, and 33 of the '047 patent are unpatentable;

FURTHER ORDERED that claims 2–4, 7, 9–11, 13, and 27 of the '047 patent have not been shown to be unpatentable; and

FURTHER ORDERED that parties to this proceeding seeking judicial review of our decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

For PETITIONER:

Kevin Bendix
Robert A. Appleby
Gregory S. Arovas
KIRKLAND & ELLIS LLP
kevin.bendix@kirkland.com
robert.appleby@kirkland.com
greg.arovas@kirkland.com

For PATENT OWNER:

Ziyong Li
Nima Hefazi
Joseph M. Paunovich
QUINN EMANUEL URQUHART & SULLIVAN LLP
seanli@quinnemanuel.com
nimahefazi@quinnemanuel.com
joepaunovich@quinnemanuel.com