

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,
Petitioner

v.

PACT XPP SCHWEIZ AG,
Patent Owner

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

Case IPR2020-TBD

**PETITION FOR *INTER PARTES* REVIEW
UNDER 35 U.S.C. § 312 AND 37 C.F.R. § 42.104**

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	MANDATORY NOTICES	2
A.	37 C.F.R. § 42.8(b)(1): Real Party-In-Interest	2
B.	37 C.F.R. § 42.8(b)(2): Related Matters	2
C.	37 C.F.R. §§ 42.8(b)(3) and (4): Counsel and Service Information	2
III.	PAYMENT OF FEES UNDER 37 C.F.R. § 42.103	3
IV.	CERTIFICATION OF GROUNDS FOR STANDING	3
V.	OVERVIEW OF CHALLENGE AND RELIEF REQUESTED	3
A.	37 C.F.R. § 42.104(b)(1): Claims for Which IPR is Requested	3
B.	37 C.F.R. § 42.104(b)(2): Grounds for Challenge	4
C.	37 C.F.R. § 42.104(b)(4): How the Claims are Unpatentable	5
D.	37 C.F.R. § 42.104(b)(5): Evidence Supporting Challenge	5
VI.	OVERVIEW OF THE TECHNOLOGY	5
VII.	OVERVIEW OF THE '605 PATENT	7
A.	The Claims	7
B.	Summary of the Specification	7
C.	Summary of the Prosecution History	8
VIII.	LEVEL OF ORDINARY SKILL IN THE ART	10
IX.	37 C.F.R. § 42.100(B): CLAIM CONSTRUCTION	10
X.	OVERVIEW OF THE PRIMARY PRIOR ART	12
A.	Nicol	12

B.	Kling	16
C.	DeHon.....	19
D.	Bhatia.....	20
XI.	SPECIFIC GROUNDS FOR PETITION	23
A.	Ground I: Claims 1–6 Are Obvious in View Nicol In Combination with Kling Under Patent Owner’s Construction	23
1.	Reasons to Combine Nicol and Kling.....	23
2.	Independent Claim 1	25
3.	Dependent Claim 2	45
4.	Dependent Claim 3	46
5.	Dependent Claim 4	48
6.	Dependent Claim 5	50
7.	Dependent Claim 6	50
B.	Ground II: Claims 1–6 Are Obvious in View of Nicol in Combination with Kling and DeHon Under Petitioner’s Constructions.....	52
1.	Reasons to Combine Nicol, Kling, and DeHon	53
2.	Independent Claim 1	58
C.	Ground III: Claims 1-6 Are Obvious In View of Nicol in Combination with Kling, DeHon, and Bhatia Under Both Parties’ Constructions	63
XII.	CONCLUSION.....	66

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>Cooler Master Co., Ltd. v. Aavid Thermalloy LLC</i> , No. IPR2019-00337, Paper 7 (P.T.A.B. June 12, 2019)	11
<i>In re Dance</i> , 160 F.3d (Fed. Cir. 1998)	59
<i>PACT XPP Schweiz AG v. Intel Corp.</i> , Case No. 1:19-cv-1006-UNA (D. Del.).....	2, 10
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005)	10
Statutes	
35 U.S.C. § 101	8
35 U.S.C. § 102	4
35 U.S.C. § 112(a)	9
35 U.S.C. § 282(b)	10
35 U.S.C. § 315(e)(1).....	3
Rules	
37 C.F.R. § 1.68	5
37 C.F.R. § 42.8(b)	2
37 C.F.R. § 42.10(b)	3
37 C.F.R. § 42.15(a)(1)	3
37 C.F.R. § 42.100(b)	10
37 C.F.R. § 42.103	3
37 C.F.R. § 42.104	3, 4, 5

PETITIONER’S EXHIBIT LIST¹

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

¹ Unless otherwise specified, citations are to the original page, column, and line numbers in exhibits. Brackets ([]) are used to refer to the sequential page numbers added to exhibits pursuant to 37 C.F.R. § 42.63(d)(2)(i).

Petition for *Inter Partes* Review of U.S. Patent No. 9,075,605

Exhibit No.	Description
1018	U.S. Patent No. 6,704,877 (“Cline”)
1019	RUDOLF F. GRAF, MODERN DICTIONARY OF ELECTRONICS (7th ed. 1999) (“Graf”)
1020	<i>Advanced Configuration and Power Interface</i> (ACPI), Revision 1.0 (December 22, 1996)
1021	Michael Peter Kennedy & Leon O. Chua, <i>Hysteresis in Electronic Circuits: A Circuit Theorist’s Perspective</i> , in 19 INT. JOURNAL OF CIRCUIT THEORY AND APPLICATIONS, 471–515 (1991)
1022	John P. Hayes, <i>Computer Architecture and Organization</i> (3rd ed. 1998)
1023	Reserved
1024	Reserved
1025	Declaration of Rachel J. Watters in Support of Public Availability of <i>Hysteresis in Electronic Circuits: A Circuit Theorist’s Perspective</i> , in 19 INT. JOURNAL OF CIRCUIT THEORY AND APPLICATIONS, 471–515 (1991) (Kennedy)
1026	U.S. Patent No. 6,397,343 (“Williams”)

Petitioner Intel Corporation (“Intel” or “Petitioner”) respectfully requests *inter partes* review (“IPR”) of claims 1–6 (“the Challenged Claims”) of U.S. Patent No. 9,075,605 (the “’605 patent”) (Ex. 1003).

I. INTRODUCTION

The ’605 patent claims a method of operating a multiprocessor comprising a plurality of data processing units (DPUs), whereby a clock frequency is dynamically adjusted based on the number of pending operations, or processing load, in addition to processor voltage, available power, a user setting, and a temperature threshold to prevent an over-temperature condition.

But the claimed invention is unpatentable. At bottom, it claims nothing more than well-known power and thermal management of a multiprocessor system taught in numerous prior art patents and publications, including Nicol and Kling relied upon in this petition. Nicol describes a multiprocessor system in which a controller dynamically allocates tasks to the PEs and dynamically adjusts the supply voltage and clock frequency of the PEs during runtime based on the processing load of the PEs. Kling teaches ways of measuring power consumption in a multiprocessor system, and also teaches reducing power whenever it exceeds a preset threshold by selectively reducing the processor clock frequencies. Because the ’605 patent merely claims well-known power and thermal management techniques, the Challenged Claims should be canceled.

II. MANDATORY NOTICES

A. 37 C.F.R. § 42.8(b)(1): Real Party-In-Interest

Intel is the real party-in-interest for Petitioner.

B. 37 C.F.R. § 42.8(b)(2): Related Matters

PACT XPP Schweiz AG (“Patent Owner”) has asserted the ’605 patent in *PACT XPP Schweiz AG v. Intel Corp.*, Case No. 1:19-cv-1006-UNA (D. Del.) (“District Court Action”).

C. 37 C.F.R. §§ 42.8(b)(3) and (4): Counsel and Service Information

Lead Counsel	Back-Up Counsel
Kevin Bendix Reg. No. 67,164 kevin.bendix@kirkland.com <u>Postal and Hand-Delivery Address:</u> KIRKLAND & ELLIS LLP 555 South Flower Street Los Angeles, CA 90071 Telephone: (213) 680-8400 Facsimile: (213) 680-8500	Robert A. Appleby, P.C. Reg. No. 40,897 robert.appleby@kirkland.com <u>Postal and Hand-Delivery Address:</u> KIRKLAND & ELLIS LLP 601 Lexington Avenue New York, NY 10022 Telephone: (212) 446-4800 Facsimile: (212) 446-4900 Gregory S. Arovas, P.C. Reg. No. No. 38,818 garovas@kirkland.com <u>Postal and Hand-Delivery Address:</u> KIRKLAND & ELLIS LLP 601 Lexington Avenue New York, NY 10022 Telephone: (212) 446-4800 Facsimile: (212) 446-4900

Intel concurrently submits a Power of Attorney, 37 C.F.R. § 42.10(b), and consents to electronic service directed to the following email address: Intel_PACT_IPR@kirkland.com.

III. PAYMENT OF FEES UNDER 37 C.F.R. § 42.103

The undersigned authorizes the Office to charge the fee set forth in 37 C.F.R. § 42.15(a)(1) for this Petition to Deposit Account No. 506092. The undersigned further authorizes payment of any additional fees for that may be due in connection with this Petition to be charged to the above-referenced deposit account.

IV. CERTIFICATION OF GROUNDS FOR STANDING

Intel certifies pursuant to Rule 42.104(a) that the '605 patent is available for IPR and that Intel is not barred or estopped from requesting an IPR of the Challenged Claims on the grounds identified in this Petition. Intel certifies: (1) Intel is not the owner of the '605 patent; (2) Intel (or any real party-in-interest) has not filed a civil action challenging the validity of any claim of the '605 patent; (3) Intel files this Petition within one year of the date it was served with a complaint asserting infringement of the '605 patent; (4) estoppel provisions of 35 U.S.C. § 315(e)(1) do not prohibit this IPR; and (5) this Petition is filed after the '605 patent was granted.

V. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED

A. 37 C.F.R. § 42.104(b)(1): Claims for Which IPR is Requested

Intel challenges claims 1–6 of the '605 patent.

B. 37 C.F.R. § 42.104(b)(2): Grounds for Challenge

Intel's challenge is based on the following prior art references, none of which were before the Patent Office during prosecution of the '605 patent:

1. U.S. Patent No. 6,141,762 ("Nicol") (Ex. 1009); filed August 3, 1998; issued October 31, 2000. Nicol is prior art under 35 U.S.C. § 102(b).²
2. U.S. Patent No. 6,367,023 ("Kling") (Ex. 1010); filed December 23, 1998; issued April 2, 2002. Kling is prior art under 35 U.S.C. § 102(e).
3. U.S. Patent No. 6,052,773 ("DeHon") (Ex. 1012); filed June 6, 1995; issued April 18, 2000. DeHon is prior art under 35 U.S.C. § 102(b).
4. U.S. Patent No. 6,535,798 ("Bhatia") (Ex. 1011); filed December 3, 1998; issued March 18, 2003. Bhatia is prior art under 35 U.S.C. § 102(e).

Petitioner requests cancellation of the Challenged Claims on the following grounds:

² All citations refer to pre-AIA 35 U.S.C. § 102. Nicol came up during an examiner's search (Ex. 1004 at 411), but was not discussed or listed as cited art.

Ground	Claim(s)	Proposed Statutory Rejection
1	1–6 under <i>Patent Owner’s</i> claim interpretation	Obvious under §103 in view of Nicol in combination with Kling.
2	1–6 under <i>Petitioner’s</i> claim interpretation	Obvious under §103 in view of Nicol in combination with Kling and DeHon.
3	1-6 under <i>Petitioner’s</i> and <i>Patent Owner’s</i> claim interpretations	Obvious under §103 in view of Nicol in combination with Kling, DeHon, and Bhatia.

C. 37 C.F.R. § 42.104(b)(4): How the Claims are Unpatentable

A detailed explanation of how the Challenged Claims are unpatentable is provided in Section XI.

D. 37 C.F.R. § 42.104(b)(5): Evidence Supporting Challenge

A list of exhibits is provided at the beginning of this Petition. The relevance of this evidence and the specific portions supporting the challenge is provided in Section XI. Intel submits a declaration of Dr. Carl Sechen, (Ex. 1001), in support of this Petition under 37 C.F.R. § 1.68.

VI. OVERVIEW OF THE TECHNOLOGY

The ’605 patent is directed to a method of operating a multiprocessor system that includes a plurality of processing elements. The patent, titled “methods and devices for treating and processing data,” discusses well known techniques of

adjusting a clock frequency in an effort to conserve power and adjusts the frequency based on a number of pending operations, which is equivalent to the processor load. Ex. 1003, Title, 1:29–31. The patent also discusses the well-known method of lowering a clock frequency when the temperature of the system reaches a certain threshold to avoid over-temperature. Ex. 1003, 7:25–30.

Power consumption is not new and plays a significant role in mobile computer systems, as battery life impacts the quality of the system. It is well known that to reduce power consumption, one can reduce the amount of time the processor is consuming excess power in a high-power state and to adjust the clock frequency based on the load or demand on the processor at a given time. Ex. 1001, ¶¶32–34. Reducing power consumption can be done in a variety of ways including stopping or lowering the clock frequency of the microprocessor, halting operations or instructions within the microprocessor, and lowering the supply voltage. Ex. 1001, ¶34. Often times this is done when the system is approaching its maximum power dissipation or a thermal operating limit. Ex. 1001, ¶35.

Operating temperature is another variable that is often monitored in processing systems. Ex. 1001, ¶50. If a processor overheats, it can potentially cause irreparable harm to the processing chip. Ex. 1001, ¶50. Thus, systems can monitor the operating temperature of a processor, and if it approaches a thermal threshold,

the system can lower the clock frequency to reduce the operating temperature. Ex. 1001, ¶¶52–53.

VII. OVERVIEW OF THE '605 PATENT

The '605 patent issued from U.S. App. No. 13/653,639, filed on October 17, 2012, claiming priority, through multiple divisional and continuation patents, to a foreign application filed on March 5, 2001. Ex. 1003, Cover.³

A. The Claims

The '605 patent has 6 claims, including one independent claim which is claim 1. Claims 1–6 are the Challenged Claims.

B. Summary of the Specification

The specification of the '605 patent describes a method of operating a multiprocessor system having a reconfigurable architecture where “power consumption may be reduced and/or optimized.” Ex. 1003, 2:41–42. It is important to note that the specification is focused on “reconfigurable processors” and “reconfigurable architecture[s],” despite the claims not containing the phrase “reconfigurable.” *Id.* at 1:31–39.

The '605 patent purports to improve power consumption through use of different parameters such as the state of the multiprocessor system, the state of the processing element or other elements, processor load/pending number of operations,

³ Petitioner reserves the right to challenge this or any other alleged priority date.

or temperature measurements to jointly establish the clock-determining state, and thus the clock frequency of the processors of the system. *Id.*, 4:54–63; Ex. 1001, ¶73.

C. Summary of the Prosecution History

U.S. Application No. 13/653,639 (the “’639 application”)—which issued as the ’605 patent—was filed on October 17, 2012, claiming priority, through multiple divisional and continuation patents, to a German Application DE 101 10 530 filed on March 5, 2001. Ex. 1003, Cover.

Claim 1 of the ’639 application was rejected in a non-final office action on December 19, 2012, as being anticipated by Williams, U.S. Patent No. 5,774,704 (“Williams ’704”). Ex. 1004, 239. The applicant responded to the office action on June 17, 2013, traversing the rejection under Williams, arguing that Williams failed to identically disclose “setting a clock frequency of at least part of the multiprocessor system to a minimum in accordance with a first processor load.” *Id.*, 307. The applicant traversed this rejection by arguing: (1) Williams ’704 does not relate to a “*processing load minimum*” but rather only a clock frequency decrease from a “maximum”; and (2) cooling of the system based not on setting the clock frequency to a minimum, but by utilizing “cooling steps, not frequency adjustment steps.” Ex. 1004 at 308.

Claim 1 was rejected in a subsequent office action on September 20, 2013 under 35 U.S.C. § 101 as an abstract idea. *Id.*, 318–23. In response to the second rejection, the applicant submitted an amendment to the claims on October 23, 2013, amending claim 1 and adding 45 new claims. *Id.*, 324–34. Notably, the 45 new claims are almost identical to claims 1–45 of the now divisional patent of the '605 patent, U.S. Patent No. 9,552,047. *See id.*, 324–34; '047 Patent, [Claims 1–45] (Ex. 1007). On February 5, 2014, the Patent Office issued a restriction requirement, and the applicant responded on March 19, 2014, electing Group 1, and filing on the same day U.S. Application No. 14/219,945, now U.S. Patent No. 9,552,047. Ex. 1004, 341–44, 353; Ex. 1007, Cover.

On July 2, 2014, claim 1 was rejected under 35 U.S.C. § 112(a) for failing to comply with the enablement requirement. ' Ex. 1004, 357–60. In response, on October 30, 2014, the applicant amended claim 1, substituting “a number of pending operations” for the “load” limitation, and adding new claims, 47–64. *Id.*, 365–82. The applicant noted that for enablement, the term “minimum” was disclosed in the specification where “. . . one or a plurality of cells are activated, if necessary, *on a very low level, i.e., very slow clocking . . .*” *Id.* at 378.⁴ In support of the hysteresis characteristic limitation, the applicant pointed to the specification that noted “it

⁴ All emphases and annotations added unless otherwise noted.

should be pointed out that a hysteresis characteristic may be provided in determining the clock rates, when a temperature-sensitive change of the clock frequencies is to be performed, for example.” *Id.* at 379. A notice of allowance was issued on April 1, 2015. *Id.* at 399.

VIII. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in the Art (“POSITA”) at the time of the ’605 patent would have been a person having at least a Master’s degree in Electrical Engineering (or equivalent experience), and at least three years of experience with processor design and memory architecture. Ex. 1001 ¶¶83–85. Additional education might compensate for a deficiency in experience, and vice-versa.

IX. 37 C.F.R. § 42.100(b): CLAIM CONSTRUCTION

Claim terms “shall be construed 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) (2018); *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). The table below lists the only disputed claim term of the ’605 Patent and their proposed constructions from the parties’ joint claim construction chart in the District Court Action. Ex 1016 (JCC), 10, 14, 15.

Claim Terms in Dispute in D. Del. Action	Intel’s Proposed Construction (“Petitioner’s Constructions”)	PACT’s Proposed Construction (“Patent Owner’s Constructions”)
“data processing unit” and/or “data processing unit . . . adapted for sequentially processing data” Claim 1	<u>“data processing unit” and/or “data processing unit . . . adapted for sequentially processing data”</u> : “reconfigurable and sequential data processors where the data results from one processor are fed to another, for each processor to perform a separate computation”	<u>“sequentially processing data”</u> : “passing results onto one or more other data processing units which are subsequently processing data”
“to a minimum” Claim 1	“to a level no more than is required for the preservation of memory contents or the like”	Plain and ordinary meaning. No construction necessary.

As explained in Section XI, the prior art renders the Challenged Claims obvious under any of the proposed constructions. Ex. 1001 ¶¶9. Ground I explains the unpatentability of the claims under *Patent Owner’s* proposed construction, Ground II explains the unpatentability of the claims under *Petitioner’s* proposed construction, and Ground III explains the unpatentability of the claims under either party’s constructions. Ex. 1001 ¶¶9. Thus, the Board does not need to construe any claim term for purposes of evaluating the prior art in this Petition. *Cooler Master Co., Ltd. v. Aavid Thermalloy LLC*, No. IPR2019-00337, Paper 7 at 11 (P.T.A.B.

June 12, 2019) (“we need not construe these claim terms to decide whether to institute review” because Petitioner’s “proposed unpatentability ground applies equally regardless of whether [Petitioner’s or Patent Owner’s] claim construction is applied”).

X. OVERVIEW OF THE PRIMARY PRIOR ART

A. Nicol

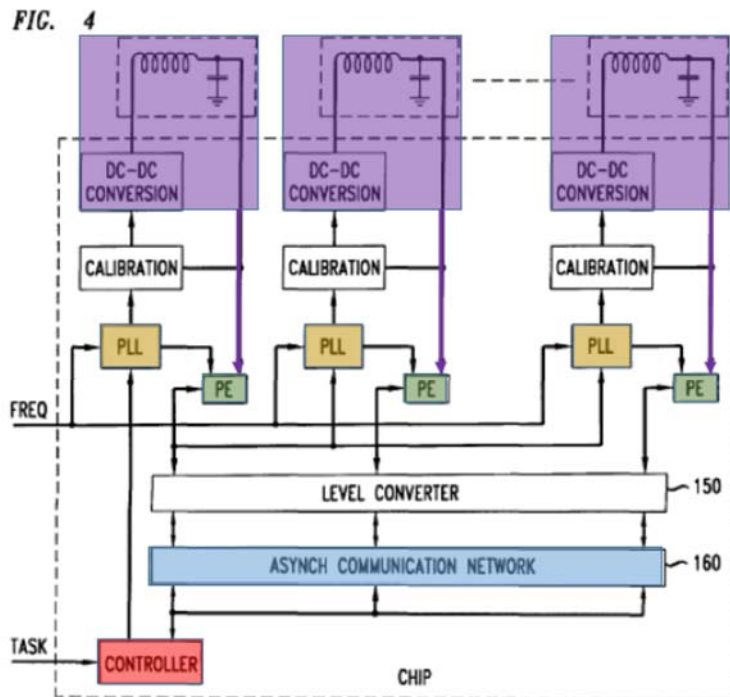
Nicol, titled “Power reduction in a multiprocessor digital signal processor based on processor load,” is directed to reducing power consumption of a multiprocessor system or chip through dynamic control of the processing load, clock frequency and supply voltage of the processing elements (PEs) of the system or chip:

Improved performance of multi-processor chips is achieved by dynamically controlling the *processing load* of chips [] so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors ..., lowers the clock frequency on the chip to as low as possible while assuring proper operation, and finally reduces the supply voltage.

Ex. 1009 (Nicol), 2:18–27. Nicol also teaches estimating processing loads based on the number of instructions of each of the “to do” tasks to be allocated to the PEs:

If the *number of instructions* that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the tasks to the available processors in order to balance the computation.


Id., 3:13–18; Ex. 1001, ¶93. In Nicol’s multiprocessor system or chip, the PEs—which are ***not restricted*** in their number or type—can be configured to support “a mix of many digital signal processing applications.” Ex. 1009, 2:53–56; Ex. 1001, ¶94. Nicol Figure 4 is an exemplary embodiment of a multiprocessor chip having multiple PEs (green) configured to “exchange data” over a communication network (blue). Ex. 1009, 6:27–33; Ex. 1001, ¶95. Each PE has its own power supply (purple) and clock generation circuitry (tan). This distributed architecture allows a controller PE (red) to dynamically allocate tasks to selected PEs (green) and selectively program “the lowest operating frequency and corresponding lowest supply voltage for each PE.” Ex. 1009, 5:23–28, 5:66–6:2.



Id., Figure 4 (annotated).

While Figure 4 shows a system-on-chip (SoC) solution, Nicol explains that its power control scheme is “extensible to other system arrangements,” including any system made of “a plurality of *separate* processor elements that operate at different frequencies and operating voltages, as well as components that are not typically thought of as processor elements.” Ex. 1009, 6:34–39.

Nicol Figure 3 illustrates the runtime clock frequency and supply voltage adjustments of the multiprocessor system as “to do” tasks are created, executed and then completed. Ex. 1009, 4:54–5:2; Ex. 1001, ¶98.



Ex. 1009, Figure 3 (annotated and corrected).⁵

In this example, the multiprocessor system is initially in a low-activity or idle state (left), and operates at a low frequency (70 MHz) and low voltage (1.8v) to conserve power. Ex. 1001, ¶99. As new “to do” tasks appear, the controller determines the required completion time based on estimated processor load, allocates the tasks to the PEs, raises the supply voltage of the PEs to a high voltage (2.7v), and also raises their clock frequency (to 140 MHz), at which point the PEs transition to a high-performance state (middle), allowing them to execute the allocated tasks and complete them within the required completion time. *Id.* After

⁵ The sequence of task creation, task start and task end in the original Figure 3 is wrong and has been corrected in red. Ex. 1001 ¶¶97–98. Nicol explains in discussing Figure 3:

FIG. 3 demonstrates the timing associated with increasing Clk, Clk-L and V_{dd}-local ***when a new task is created*** [depicted as “task created” on corrected figure] ... Specifically, it shows the system operating at 70 MHz from a 1.8V supply when the load is increased in three steps to 140 MHz. When the ***2.7V supply*** is stable, as shown by the supply voltage plot, ***the new task is enabled for execution*** [depicted as “task started” on corrected figure]. Sometime thereafter according to FIG. 3, ***a task completes, which reduces the load on the multiprocessor*** [depicted as “task ended” on corrected figure]. The reduced load permits lowering the clock frequency to ***100 MHz*** and lowering the supply voltage to 2.1 V.

Ex. 1009, 4:58–66.

completion of the tasks, the clock frequency and supply voltage of the PEs are progressively reduced to minimize power consumption (right). *Id.* Nicol teaches that inactive PEs can also be selectively clock-gated and turned off to save power. Ex. 1009, 5:43–45.

B. Kling

Kling describes a method and apparatus “for managing the power consumed in a computer system.” Ex. 1010 (Kling), 1:50–51. Kling explains that computer systems typically include one or more processors and that modern computer systems are designed to operate at increasingly higher clock speed for higher performance. *Id.*, 1:16–21. Kling explains that, “as processor speed increases, the power consumed by the processor tends to increase as well.” *Id.*, 1:21–23. Increased power consumption could destroy devices due to over-temperature, and could significantly reduce battery life for mobile devices. Ex. 1001, ¶102. Kling discusses conventional techniques to address those issues. In particular, to address critical over-temperature, a thermal sensor is typically used to monitor the temperature of a processor. Ex. 1001, ¶103. “If the processor temperature exceeds a particular threshold, the processor is placed into a low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.” Ex. 1010, 1:38–46.

Kling's Figure 1 illustrates a multiprocessor computer system that includes processors 100 and 101 (green) and a power controller (red), as well as other integrated circuits (ICs) such as audio/video accelerators and chipsets. *Id.*, 2:65–3:19.

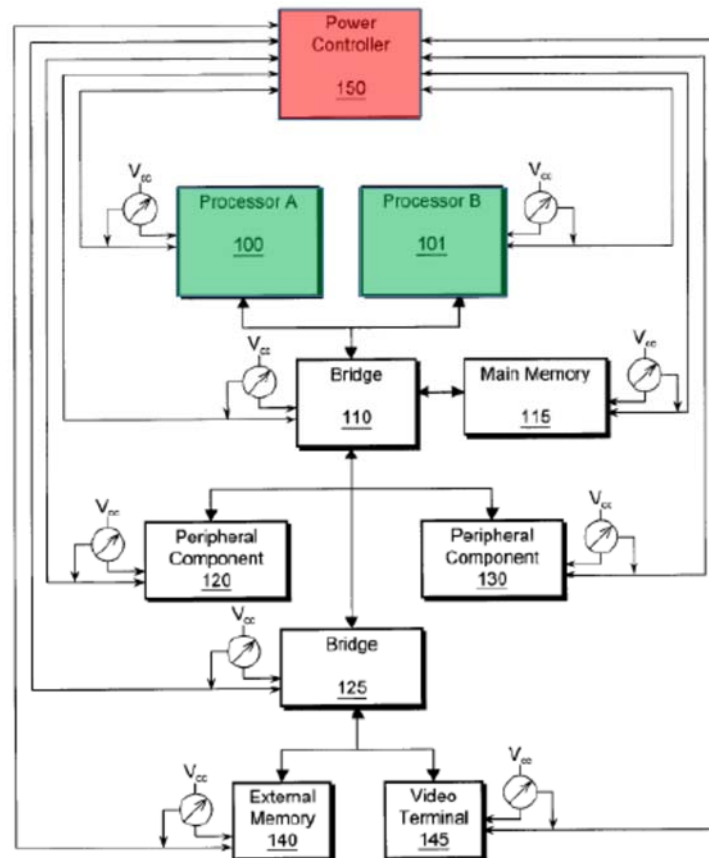


Figure 1

Id., Figure 1 (annotated).

The computer system of Figure 1 “includes one or more independent power supplies, each of which provides power to one or more ICs.” *Id.*, 3:24–28. The system can “monitor the power consumed by a single IC or any number of ICs.” *Id.*,

3:28–29. If the measured power exceeds a threshold, the power controller (red) sends a “throttle signal” to “one or both of the processors 100 or 101” or other ICs of the computer system. *Id.*, 4:50–5:8. The ICs that receive the throttle signal must reduce their power consumption. *Id.*, at 2:45–47; Ex. 1001, ¶107. Notably, a processor “*reduces its core frequency*” in response to receiving the throttle signal, as shown in the flow chart Figure 5. Ex. 1010, 2:47–51.

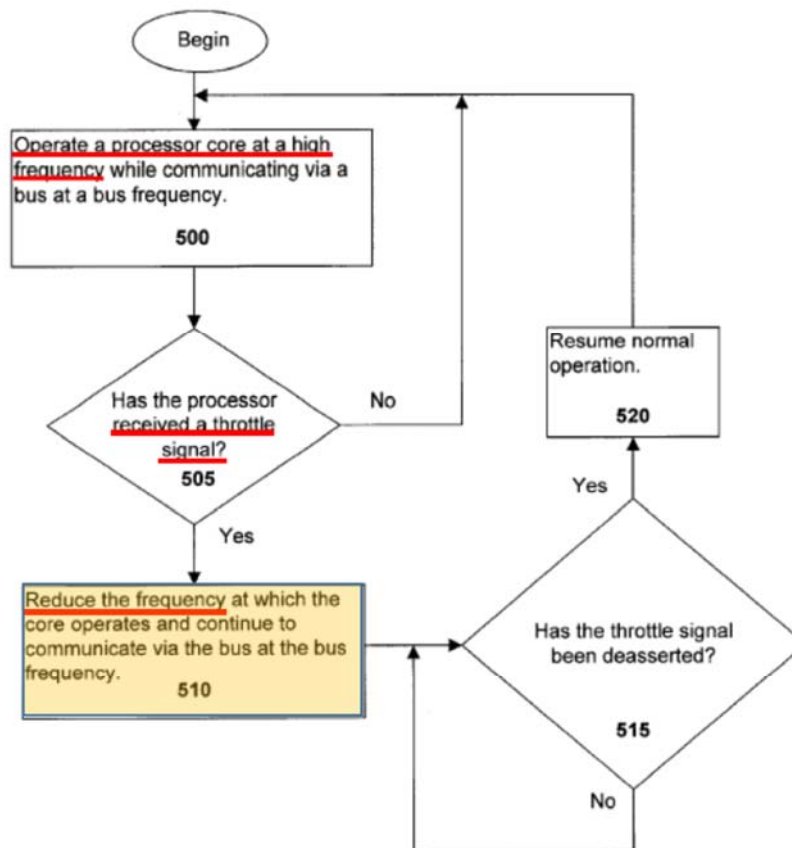


Figure 5

Ex. 1010, Figure 5 (annotated).

C. DeHon

DeHon is directed to a reconfigurable microprocessor chip, called a DGPA-coupled microprocessor chip. Ex. 1012 (DeHon), Title. As illustrated in DeHon's Figure 1, the chip "includes a ***fixed logic processor*** [core cell 10, red] as is common to perform logic operations." Ex. 1012, 2:10–11. "Minimally, the processor cell 10 includes an arithmetic logic unit (ALU), possibly a fixed processing unit (FPU), register files, instruction caches, and data caches." *Id.*, 5:42–44.

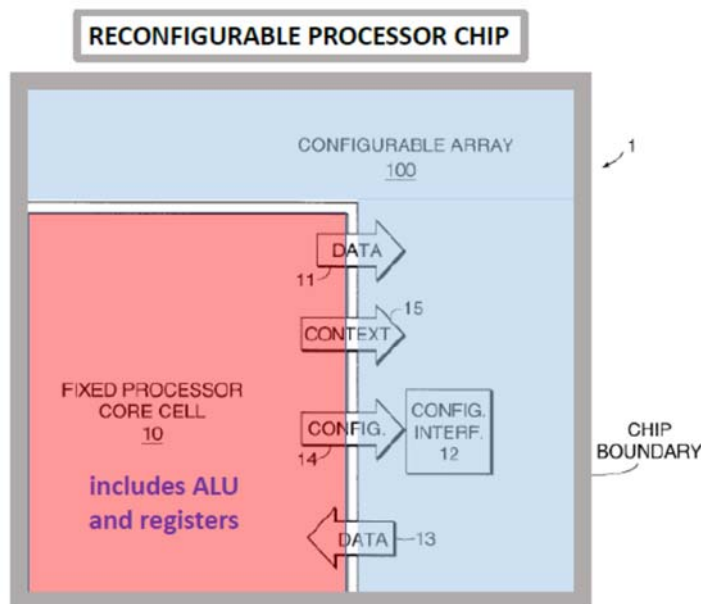


FIG. 1

Id., Figure 1 (annotated).

DeHon explains that "[a] portion of the chip silicon real-estate, however, is dedicated to a ***programmable gate array*** [100, blue]. This feature enables ***application-specific configurations***." *Id.*, 2:12–15. Similar to the '605 patent,

DeHon's microprocessor chip is dynamically adjustable at runtime and "enables application-specific configurations" where the gate array is "configurable to adapt to the particular time-changing demands of the microprocessor to provide the functionality required to best serve those demands." *Id.*, 2:13–17; Ex. 1001, ¶112. Further, the configurable array initially receives configuration programming at power up or from the processor cell and, similar to the '605 patent, "[o]nce the proper configuration has been programmed into the array 100, the processor cell 10 feeds input data 11 to and receives the data output 13 from the configurable array 100." Ex. 1012, 5:45–53; Ex. 1001, ¶113. "For example, the operating system and applications specialize the configurable array to serve as a co-processor for compute-intensive functions such as compression, encryption, signal processing, or protocol processing." Ex. 1012, 5:63–67. Different DGPA-coupled chips can thus be advantageously configured in different ways to handle a variety of compute-intensive functions "such as compression, encryption, signal processing, or protocol processing." *Id.*, 5:63–67. Another advantage is that the configurable DGPA-coupled chips can replace or be deployed alongside traditional fixed processors. *Id.*, 7:20–25; Ex. 1001, ¶114.

D. Bhatia

Bhatia is directed to the thermal management of a heterogeneous system having processors of different types. Ex. 1011 (Bhatia), 2:31–36 ("Example

components may include general or special-purpose processors ... application-specific integrated circuits (ASICs), programmable gate arrays (PGAs), peripheral device controllers, and other types of devices.”). The system (or parts of it) operates at different performance states, such as a high performance (HP) state and a low performance (LP) state, that are associated with different core clock frequency and voltage level settings. *Id.*, 2:41–48. In normal operation, the system may mostly operate in the HP state—at high voltage and frequency settings—for maximum performance. Ex. 1001, ¶116. However, when an over-temperature condition is detected, Bhatia teaches “throttling” back by cycling between the HP and LP states to reduce power consumption and lower system temperature below a target temperature. *Id.* at 2:51–55; Ex. 1001, ¶117. Thus, Bhatia’s thermal management involves dynamic voltage and frequency control of processors. Ex. 1001, ¶117. Bhatia’s Figure 6 illustrates the effectiveness of reducing temperature through dynamic voltage and frequency control of the processors. Ex. 1011, Figure 6.




FIG. 6

Id., Figure 6 (annotated).

Bhatia also discusses temperature sensing in detail, and discusses reducing the performance state “when an over-temperature condition is detected.” *Id.*, Abstract. Although Bhatia is primarily directed to thermal management of a system or chip, Bhatia also discusses power management. *Id.*, 12:26–37; Ex. 1001, ¶120. Indeed, power and thermal management techniques are commonly implemented in computer systems; also, many power and thermal management features are defined in the *Advanced Configuration and Power Interface* (ACPI) industry standard specification. Ex. 1011, 1:7–18 (citing the 1996 published version of ACPI). Notably, ACPI specifies “*low activity state[s]* (e.g., deep sleep, stop grant, C1, C2, or C3 state)” commonly used by power management systems to reduce power. *Id.*, 12:58–64; Ex. 1001, ¶121. In the C3 state for example, power is minimized because

“no activity is performed by the processor except maintenance of the stored data in the processor’s internal cache.” Ex. 1011, 12:58–64.

XI. SPECIFIC GROUNDS FOR PETITION

A. Ground I: Claims 1–6 Are Obvious in View Nicol In Combination with Kling Under Patent Owner’s Construction

Ground I explains the unpatentability of claims 1–6 under *Patent Owner’s* proposed construction. *See* Section VIII.

1. Reasons to Combine Nicol and Kling

A POSITA would have been motivated to combine the complementary teachings of Nicol and Kling, as explained below and supported by the Declaration of Dr. Sechen. Ex. 1001, ¶122.

First, a POSITA would have found it obvious to combine the teachings of Nicol and Kling because both are directed to interrelated thermal and power management aspects that are critical and ubiquitous in processor-based systems. Ex. 1009, Title (“Power Reduction in a Multiprocessor Digital Signal Processor Based on Processor Load”); Ex. 1010, 1:7–10 (“The present invention relates to computer systems and more particularly to limiting the power consumed in a computer system”); Ex. 1001 ¶123, citing Ex. 1011 (Bhatia), 1:6–8 (“Different types of power and thermal management techniques have been implemented in system.”)

Second, as detailed under Section X.A (Overview of Nicol) and Section X.B (Overview of Kling), Nicol and Kling describe the same flexible and distributed

multiprocessor architecture in which processors have their own power supplies, and in which software residing in a controller is programmed to selectively control the power consumption of each processor, notably by dynamically adjusting the processor's clock frequency during operation. Ex. 1001 ¶124. Nicol teaches that the controller software is programmed to adjust the clock frequency and supply voltage of the processors (PEs) based on their processing loads. Ex. 1009, 2:18–27. Kling discloses controller software programmed to adjust clock frequency of the processors as well, although based on sensor measurements that reflect processing loads, such as power, voltage and current. Ex. 1010, 3:50–57. The software differences are not material to the combination, although a POSITA would have found it obvious to program Nicol's controller to handle the additional software functionalities discussed in Kling, yielding an improvement with predictable results. Ex. 1001 ¶126.

Third, Nicol is focused on minimizing power consumption and does not discuss over-temperature protection. Because over-temperature protection is critical for high-performance systems such as Nicol's multiprocessor system and SoC, a POSITA seeking to improve Nicol would have been motivated to consult Kling, which discusses well-known techniques of providing over-temperature protection, as well as additional improvements. Ex. 1001 ¶127. One such improvement described in Kling uses simple power measurement, such as measuring current, to

detect excessive power consumption (when it is above a threshold) and use that information to “throttle” back processors to reduce power consumption. Ex. 1010, 5:38–59. A POSITA would have been motivated to implement this functionality in Nicol’s system for over-temperature protection and, more generally, to detect excessive power consumption and reduce it, which is consistent with Nicol’s objective of minimizing power consumption. Ex. 1009, 2:18–22; Ex. 1001 ¶129. Implementing this added functionality would not require undue experimentation, has a good likelihood of success, and would yield predictable results because the added functionality is primarily add-on software to Nicol’s existing controller software. *Id.*, 3:9–20; Ex. 1001 ¶129.

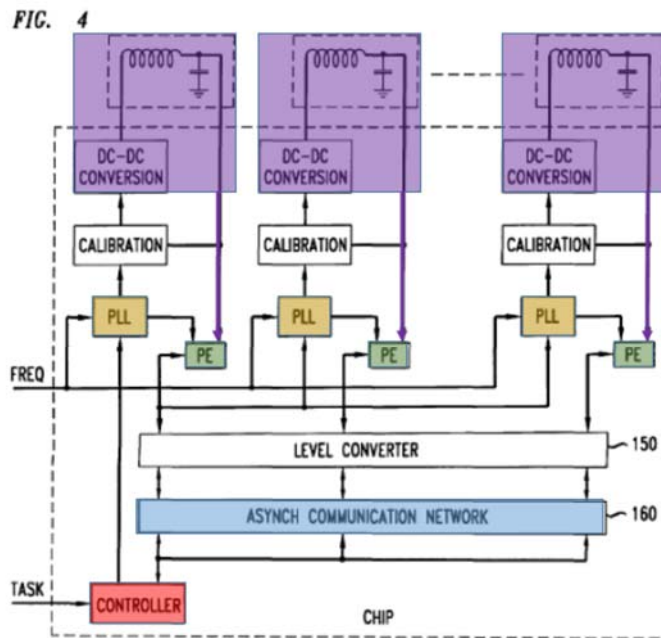
2. Independent Claim 1

a) “A method for operating a multiprocessor system comprising”

As discussed above in Section X.A (Overview of Nicol), Nicol discloses a method of operating a multiprocessor system, such as the single-chip multiprocessor system illustrated in the Figure 4 embodiment (below, annotated) that includes a plurality of processing elements (PEs) (green).⁶ Ex. 1001 ¶130. Each PE has its

⁶ Although Figure 4 embodiment illustrates a multiprocessor system integrated in a single chip, Nicol explains that its method of operating a multiprocessor system

own power supply (purple) and clock generation circuitry (tan). Ex. 1009, 5:66–6:13. The PEs exchange data over a communication network (blue). *Id.*, 6:27–33 (explaining that the level converter 150 and asynchronous communication network 160 in Figure 4 resolve “the issue of synchronizing the exchange of data among the PEs of [the] multiprocessor chip.”).



Id., Figure 4 (annotated).

is “extensible to other system arrangements,” including any system made of “a plurality of *separate* processor elements that operate at different frequencies and operating voltages, as well as components that are not typically thought of as processor elements.” Ex. 1009, 6:34–39.

A controller (red) “assigns the tasks given to the multi-processor chip of FIG. 4 among the PEs.” Ex. 1009, 6:11–13. The operating system in the controller can “independently program the lowest operating frequency and corresponding lowest supply voltage for each PE.” *Id.*, 5:66–6:2. Nicol teaches that the operating system can also adjust the clock frequency and supply voltage to all PEs in unison. *Id.*, 6:67–7:3. In that case, more explicitly described in the context of the similar but simpler Figure 2 embodiment, the operating system “needs to ascertain the required completion time [of the tasks], divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency [common to all PEs] to insure that the most heavily loaded PE carries out its assigned task within the required completion time.” *Id.*, 3:66–4:5.

As discussed above in Section X.B (Overview of Kling), Kling also discloses a method of operating a multiprocessor system, such as the computer system illustrated in Kling Figure 1 (below) that includes a plurality of processors (green). Ex. 1010, Figure 1; Ex. 1001 ¶¶136–37. Power consumption of the system is monitored and, if it exceeds a threshold, the system’s power controller (red) sends a “throttle signal” to command a reduction in power consumption. Ex. 1010, 3:28–29 and 4:50–5:8. A processor “*reduces its core frequency*” in response to receiving the throttle signal. *Id.*, 2:47–51.

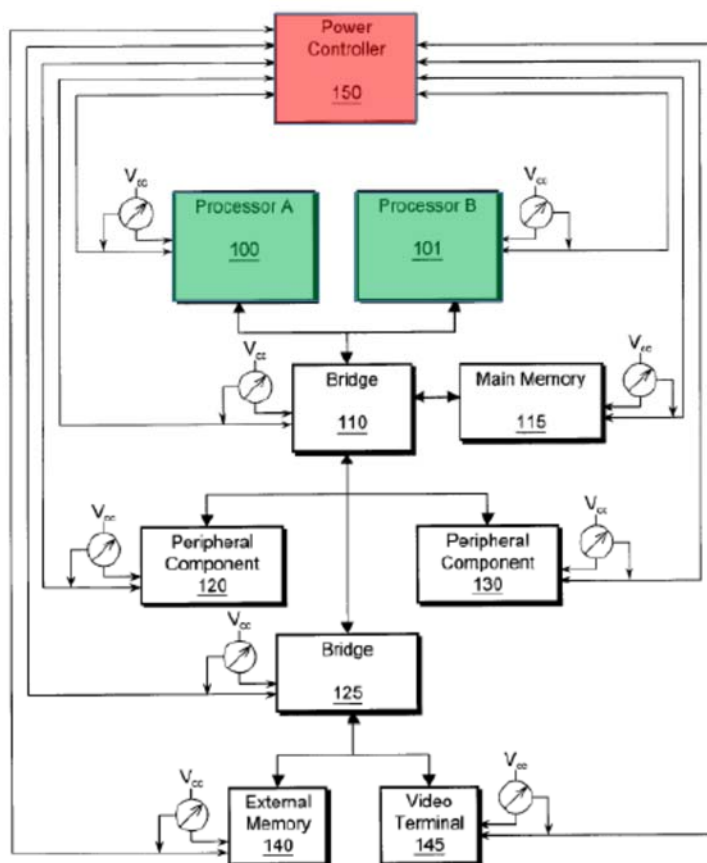


Figure 1

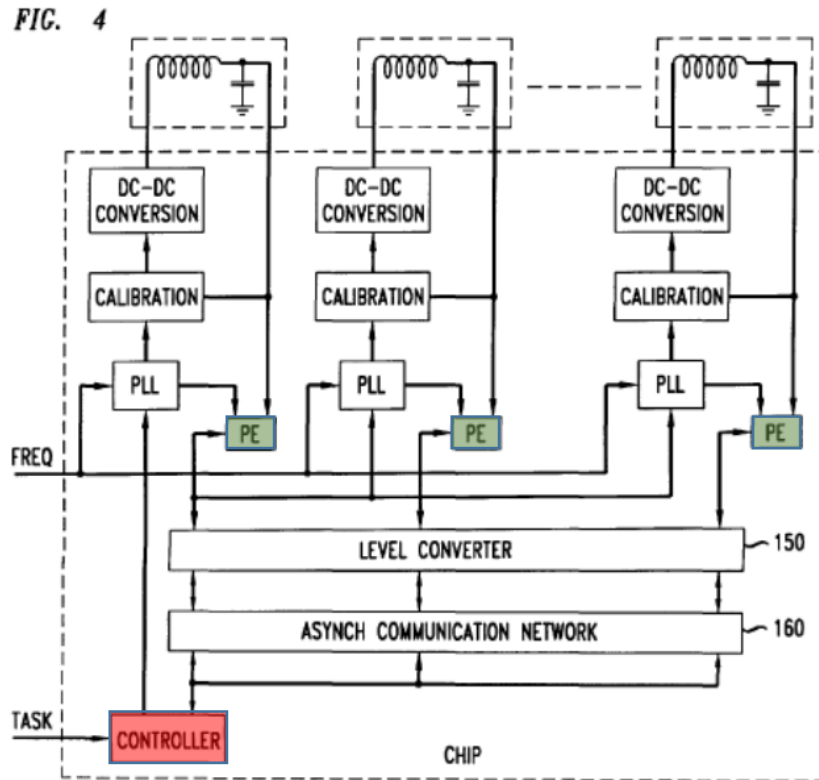
Id., Figure 1 (annotated).

b) “a plurality of data processing units,”

Nicol discloses this claim element. Ex. 1001 ¶139; Ex. 1016.

Nicol’s multiprocessor chip, illustrated in both Figure 4 (below) and Figure 2, contains a plurality of processing elements (PEs) (green) that are data processing units (DPUs) because they each process data. Ex. 1009, 2:50–53 (explaining that the multiprocessor chip “contains *processing elements (PEs)* . . . and each PE contains a central processing unit (CPU) and a local cache memory (not shown).”). The multiprocessor chip also includes a controller (red) that is a PE “that assigns the

tasks given to the multi-processor chip among the PEs.” *Id.*, 3:10–20; Ex. 1001, ¶141.



Ex. 1009, Figure 4 (annotated).

- c) “each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit”*

Nicol renders this element obvious. Ex. 1001 ¶144. Nicol discloses multi-processor chips containing processing elements (PEs), whereby “each PE contains a central processing unit (CPU) and a local cache memory (not shown).” Ex. 1009, 2:50–53. A POSITA would have understood that a PE having a CPU and cache includes arithmetic logic unit(s) (ALU) and registers. Ex. 1001, ¶145, citing Ex.

1022 (Hayes) at 33 and Figure 1.17. Hayes Figure 1.17 (annotated and reproduced below) shows a PE (labeled “data-processing unit DPU”) (green) with multiple ALUs (magenta) and associated registers (purple). *Id.*

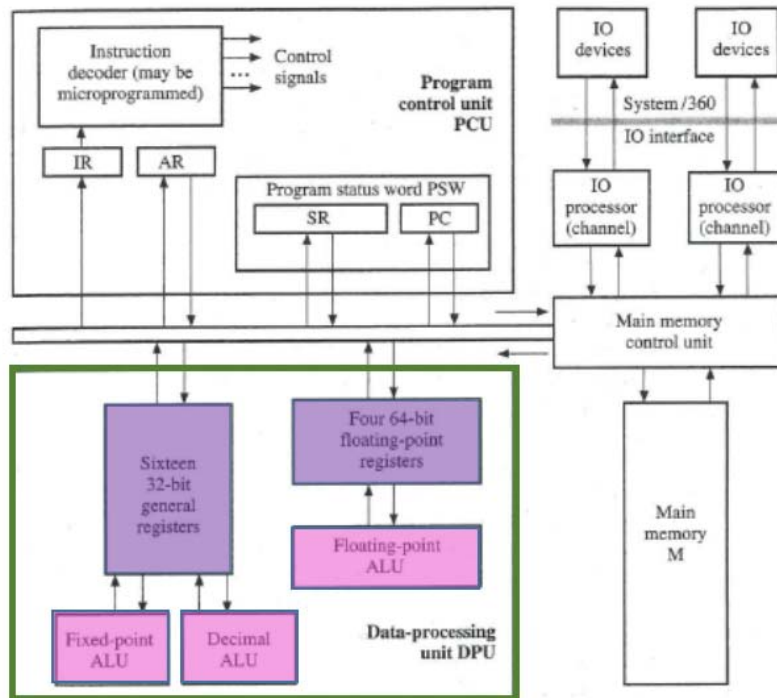


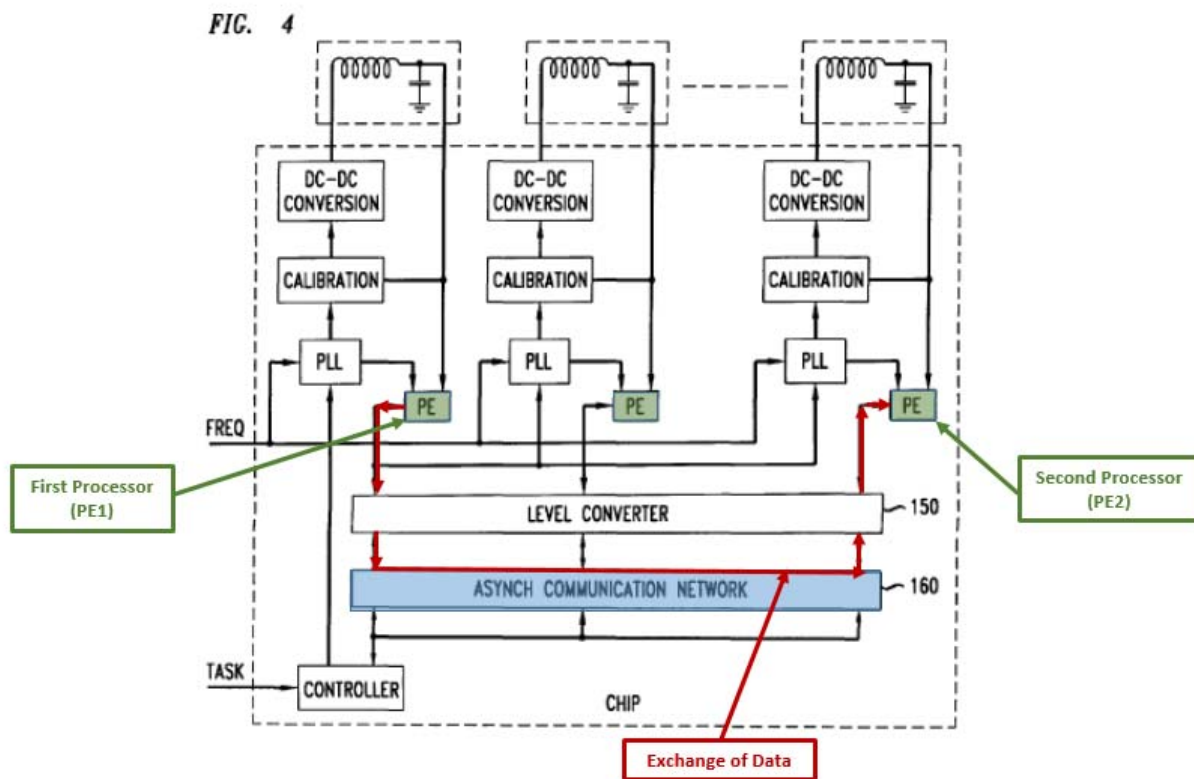
Figure 1.17
Structure of the IBM System/360.

Id., Figure 1.17

d) “and (b) being adapted for sequentially processing data, the method comprising:”

Nicol renders obvious this limitation under Patent Owner’s construction of the term “sequentially processing data” to mean “passing results onto one or more other data processing units which are subsequently processing data.” Ex. 1001 ¶147; Ex. 1016 (JCC) at 11 and 17. Specifically, Nicol teaches the need for “*synchronizing the exchange of data* among the PEs [green] of a multiprocessor

chip” through an “asynch[ronous] communication network” 160 (highlighted in blue in Figure 4). Ex. 1009, 6:26–28, 6:14–18, Figure 4. For example, Nicol’s annotated Figure 4 below depicts the passing of data (red) from PE1 (green) through the asynchronous communication network (blue) to PE2 (green) for further processing. Ex. 1001 ¶148; Ex. 1009, Figure 4.



Id., Figure 4 (annotated).

Asynchronous communication networks were well-known in the art. Ex. 1001 ¶150, citing Ex. 1022 (Hayes) at 491–93. A POSITA 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).” Ex. 1009, 6:19–21; Ex. 1001 ¶151. Thus, Nicol’s multiprocessor system is arranged with an asynchronous communication network (160 blue) *specifically* for the purpose of “passing results onto one or more other data processing units which are subsequently processing data.” Ex. 1016 (JCC) at 11 and 17; Ex. 1001 ¶151. Such a communication network would have been unnecessary—and Nicol’s multiprocessor system would have been of little practical use—if one PE did not have the ability to pass its results to other PEs for further processing, as was commonly done in prior art multiprocessor systems. Ex. 1001, ¶152, *citing* Ex. 1022 (Hayes), 159, Figure 3.14 (below) (describing a multiprocessor system in which a processor (red) sends data to a floating-point co-processor (green) over a 32-bit data bus (blue) for further processing).

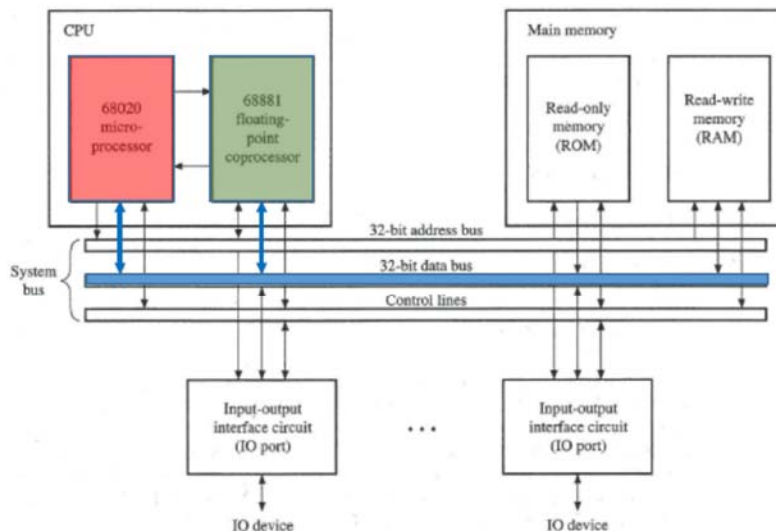


Figure 3.14
68020-based microcomputer with floating-point coprocessor.

Ex. 1022 (Hayes), Figure 3.14 (annotated).

Dr. Sechen explains that, in Nicol’s multiprocessor system, the PEs “exchange data” over a communication network because the system is intended to handle “a mix of many digital signal processing [DSP] applications,” and DSP applications typically involve complex operations performed in a pipeline fashion, whereby partial results from one PE are passed to another PE of the pipeline for further processing. Ex. 1001, ¶155; Ex. 1009, 2:55–56. A well-known example is graphics processing, where graphical data (such as an image or pixels) is processed in successive stages through a graphics pipeline, illustrated in the figure below. Ex. 1001, ¶156, citing Ex. 1026 (Williams), Figure 4. In Williams’ Figure 4, data generated by the “Pixel Ops” processor are sent to the “Rasterization” processor for further processing, then to the “Per Fragment Ops” processor for additionally processing, and finally to the “Frame Buffer.” Ex. 1001, ¶156, citing Ex. 1026 (Williams), 10:56–67.

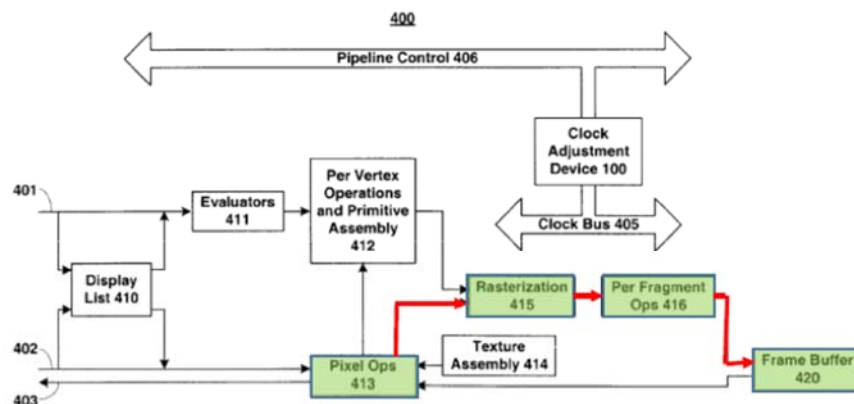


Fig. 4

Ex. 1026, Figure 4 (annotated).

Moreover, a POSITA 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. Ex. 1001, ¶158, citing Ex. 1014 (Kung, H.T., *Why Systolic Architectures?*, Computer Vol. 15 (Pub. January 1982)) ("Kung"), 45 (describing "numerous systolic designs [] known" in the 1980s), 37 ("In a systolic system, data flows from the computer memory in a rhythmic fashion, *passing through many processing elements before it returns to memory*, much as blood circulates to and from the heart."). Kung's Figure 1 below illustrates PEs that are arranged in a systolic array and pass data from one PE to the next, with each PE performing a separate processing operation on the received data. Ex. 1014 (Kung), Figure 1; *see also* Ex. 1022 (Hayes textbook), at 290 ("Closely related conceptually to arithmetic pipelines are the data-processing circuits called *systolic arrays* [] formed by interconnecting a set of identical data-processing cells in a uniform manner. Data words flow synchronously from cell to cell, with each cell performing a small step in the overall operation ... Systolic processors have been designed to implement various complex arithmetic operations ...").

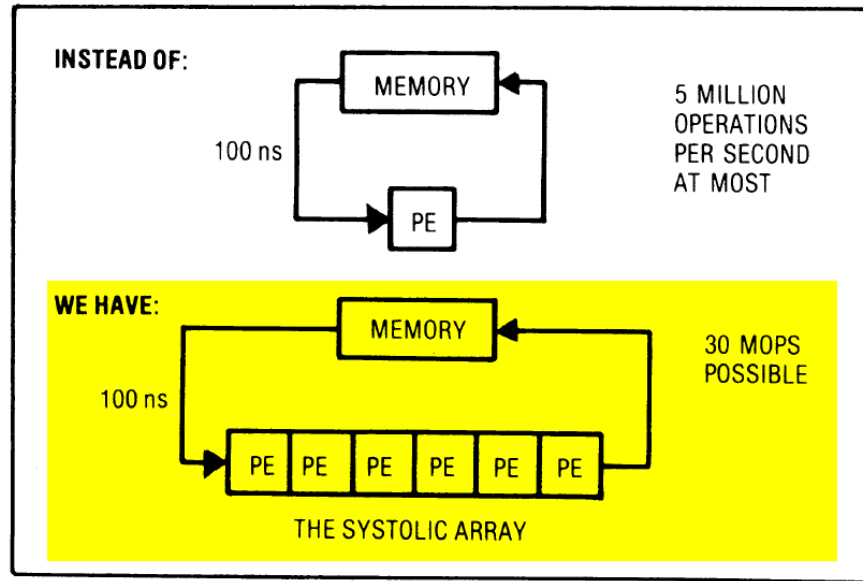
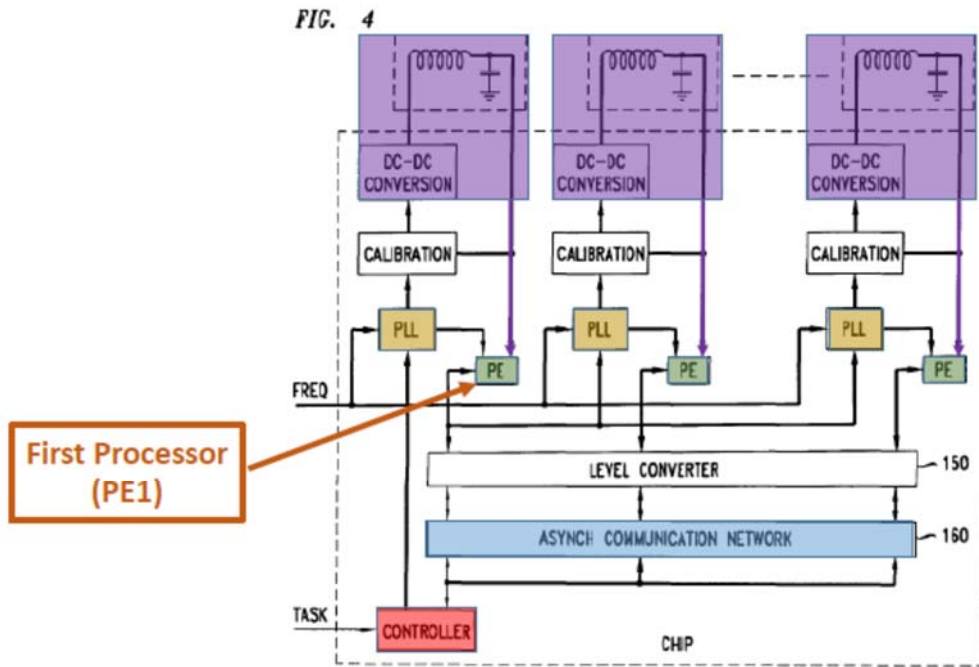


Figure 1. Basic principle of a systolic system.

Ex. 1014, Figure 1 (annotated).

- e) “the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;”*

Nicol discloses this claim element. Ex. 1001 ¶161. Nicol explains that the controller (red) in its multiprocessor system sets a clock frequency to a minimum according to a number of pending operations (instructions) of each of the PEs of the system, including the number of pending operations of PE1 (“first processor”) in the annotated Figure 4 below. Ex. 1001 ¶162; Ex. 1009, Figure 4.

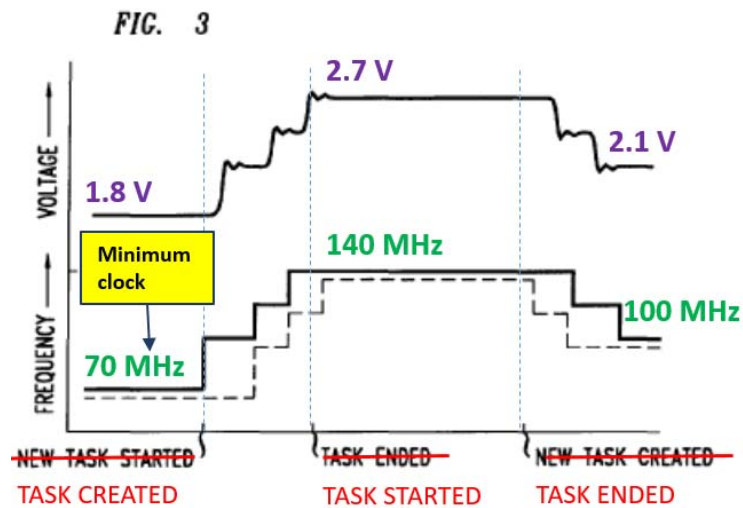


Ex. 1009, Figure 4 (annotated).

Indeed, Nicol explains that the controller “allocates tasks to the individual processors [PEs, including PE1] to equalize processing load among the [PEs, including PE1], then the controller *lowers the clock frequency* on the chip to *as low a level as possible* while assuring proper operation, and finally reduces the supply voltage.” Ex. 1009, 2:23–27. Nicol further explains that the controller takes into account the number of pending operations (instructions) of each PE—including of PE1 (“first processor”)—in its determination of processing load and task allocation. Ex. 1009, 3:13–18 (“If *the number of instructions* that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the

tasks to the available processors in order to balance the computation.”); Ex. 1001, ¶165.⁷

Figure 3 of Nicol illustrates setting the clock frequency of the multi-processor system to a minimum (70 MHz in Figure 3) when the PEs, including PE1, have a light load with few (or no) pending operations, and a new “to do” task is not yet assigned to the PEs for execution. Ex. 1001 ¶166; Ex. 1009, 3:31–33, 4:54–62, 5:23–28.



Although Figure 3 illustrates a minimum clock frequency of 70 MHz, Nicol also teaches setting the clock frequency to zero if none of the PEs are needed because

⁷ Patent Owner also represented to the patent office during prosecution of the '605 patent that “a number of operations pending [] is equivalent to the processor load.” Ex. 1004 at 377.

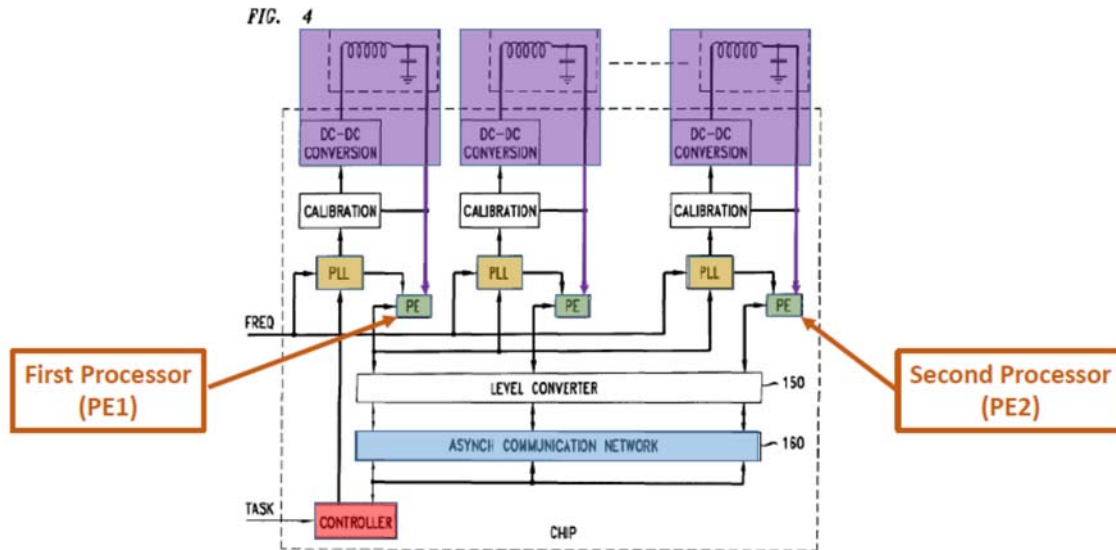
all tasks have been completed and there are no new tasks or pending operations for any PE. Ex. 1009, 5:42–44 (“the operating system employs gated clocking techniques to ‘shut down’ PEs that are not needed.”); Ex. 1001, ¶168.

Consistent with such gating techniques described in Nicol, a POSITA would have known that Nicol’s power management system with dynamic frequency control, would set the clock frequency of processing elements that are not needed (or had a low number of pending operations) to a “minimum.” Ex. 1001, ¶169. For example, the well-known industry standard, “the Advanced Configuration and Power Interface (ACPI) Specification” was commonly used and referred to by a POSITA. Ex. 1001, ¶170; Ex. 1020 (ACPI Spec.); *see also* Ex. 1011 (Bhatia), 1:6–17 (explaining “the Advanced Configuration and Power Interface (ACPI) Specification, Rev. 1.0., published on Dec. 22, 1996, that provides an interface between the operating system of a system and hardware devices to implement power and thermal management.”). Such a “minimum” clock frequency teaching is discussed in the ACPI Specification, and it was thus well-known in the art to set the clock frequency to a minimum in a low power mode. Ex. 1001, ¶171; Ex. 1011, 12:58–65 (“the C3 state defined under 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.*”).

Thus, for example, a POSITA would have understood that Nicol’s system could have referenced the teachings of the ACPI, and implemented a similar low power mode in Nicol’s system, thus setting the clock rate of inactive PEs to a “minimum” clock rate. Ex. 1001, ¶172. A POSITA 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 system’s transitioning back to full operation upon “waking up” from a low power mode. Ex. 1001, ¶172.

f) ***“the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor;”***

Nicol discloses this element. Ex. 1001 ¶173. As discussed above in Section X.A (Overview of Nicol) and in the preceding section XI.A.2.e, the controller in Nicol’s multiprocessor system adjusts the clock frequency of the system according to a number of pending operations of each of the PEs—including the number of pending operations of PE2 (“second processor”)—identified in Figure 4 (annotated) below. *Id.*; Ex. 1009, 2:23–27 (discussing allocating tasks to the PEs and setting an appropriate clock frequency for task execution) and 3:13–18 (discussing using ***“the number of instructions*** that need to be executed for each task” for task allocation).

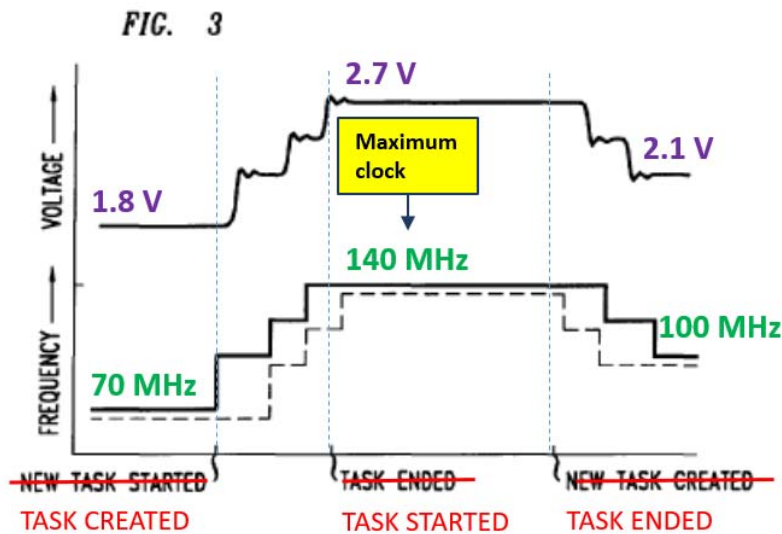


Ex. 1009, Figure 4 (annotated).

Notably, the controller takes into account the number of pending instructions for each PE—including PE1 (“first processor”) and PE2 (“second processor”)—because it sets the clock frequency based on identifying the processing requirement of “the most heavily loaded PE.” Ex. 1001, ¶176; Ex. 1009, 3:66–4:5 (explaining that the operating system “needs to ascertain the required completion time [of the tasks], divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and *adjust the clock frequency* [common to all PEs] to insure that *the most heavily loaded PE* carries out its assigned task within the required completion time.”).

Nicol Figure 3 illustrates the multiprocessor system increasing the clock frequency of the PEs of the multiprocessor system from 70 MHz (minimum clock

frequency) to 140 MHz (maximum clock frequency) prior to the execution of a high load, *i.e.*, involving a high number of pending operations for the PEs. Ex. 1001 ¶177; Ex. 1009, 3:31–33, 4:54–62, 5:23–28. Nicol also explains that each chip has a “*maximum [frequency]* that the chip can support” (Ex. 1009, 3:63–64), as would have been well-known to a POSITA, and that “[f]or a *maximally utilized system*, all of the available processors [PEs] ought to be operating at full speed [maximum frequency] when satisfying the maximum load encountered by the system.” *Id.*, 2:65–3:4; Ex. 1001, ¶178.



Ex. 1009, Figure 3 (annotated).

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

Nicol in view of Kling renders this claim element obvious. Ex. 1001 ¶180. Kling explains that processing systems must be able to rapidly detect and mitigate over-temperature situations. Ex. 1010, 1:38–46 (“To address the thermal issue [in modern computer systems] . . . processor packages include a thermal sensor to monitor the temperature of the processor. *If the processor temperature exceeds a particular threshold*, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.”). Thus, as discussed in Section XI.A.1 (Reasons to Combine Nicol and Kling), a POSITA would have found it obvious, *and in fact necessary*, to include over-temperature control in the multiprocessor system of Nicol. Ex. 1001 ¶182.

Notably, a POSITA would have found it obvious to detect over-temperature in Nicol’s multiprocessor system by comparing the system’s measured temperatures against a temperature threshold, as required by the claim. Ex. 1001 ¶183. Kling explains that this was in fact conventionally done in prior art systems. Ex. 1010, 1:43–45 (“If the processor temperature *exceeds a particular threshold*, the processor is placed into low power mode until it cools off.”); *see also* Ex. 1011, 1:6–12 (same). Moreover, a POSITA would have been motivated to use this conventional technique because Nicol’s multiprocessor system already includes on-chip temperature sensor(s) for “tracking temperature.” Ex. 1009, 4:11–13 (“the operating system [in the controller] can reduce the supply voltage even further by *tracking temperature*

...”); Ex. 1001 ¶185.⁸ Although Nicol discusses using the tracked temperatures for on-chip voltage scaling, the same temperature measurements can obviously and advantageously be used to also detect over-temperature conditions, so to improve the safety and reliability of Nicol’s multiprocessor chips. Ex. 1001 ¶186.

As Kling explains and as would have been well-understood by a POSITA, once over-temperature is detected, the overheated chip or system needs to be “placed into low power mode until it cools off.” Ex. 1010, 1:44–45; Ex. 1001 ¶187. Placing Nicol’s multiprocessor system in a low power mode requires *reducing the clock frequency* of its PEs to reduce power consumption, as also taught by Nicol and Kling and as would have been well-known to a POSITA. Ex. 1001 ¶188; Ex. 1009, 1:26–28 (“The power consumption of a CMOS circuit increases linearly with operating frequency and quadratically with supply voltage.”); Ex. 1010, 2:43–52 (explaining that when a processor receives a “throttle signal” requesting it to reduce its power consumption, the processor “reduces its core frequency”).

⁸ See also Ex. 1009, 1:41–43 (“Voltage scaling as a function of temperature has been incorporated into the Intel Pentium product family as a technique to achieve high performance at varying operating temperatures and process corners. ... The approach uses an *on-chip temperature sensor* ...”).

Kling also discusses the use of hysteresis to improve robustness and “reduce the occurrence of power oscillation,” as would also have been well-known in the art. Ex. 1010, 5:50–58; Ex. 1001 ¶189. Hysteresis is commonly implemented by comparing the same signal against two thresholds, as shown in Kling Figure 2A (below) and as would have been well-understood by a POSITA. Ex. 1010, 5:51–54 (“In accordance with the present invention, the lower threshold may be set to a predetermined value *to provide hysteresis* to reduce the occurrence of power oscillation between the upper and lower thresholds.”); Ex. 1001 ¶190, citing Ex. 1021 (Kennedy, M. & Chua, L., *Hysteresis in Electronic Circuits: A Circuit Theorist’s Perspective*, International Journal of Circuit Theory and Applications, Vol. 19, (1991)) at 471 (explaining that circuits with hysteresis characteristics were “indispensable” and commonly used “in noisy environments and to guard against false triggers”),.




Figure 2A

Id., Figure 2A (annotated).

A POSITA would have found it obvious to apply the well-known two-threshold method (providing hysteresis) from Kling in implementing an over-temperature monitoring system in Nicol for the same well-known reason—to reduce unwanted oscillations, and thus to improve the reliability and performance of Nicol’s power management system. Ex. 1001 ¶192.

3. Dependent Claim 2

a) “The method of claim 1 wherein:”

See supra Section XI.A.2.

b) “said clock frequency is determined in part by processor voltage.”

Nicol in view of Kling renders this claim element obvious. Ex. 1001 ¶193. As discussed above in Section XI.A.1 (Reasons to Combine Nicol and Kling), a POSITA would have found it obvious to add Kling’s power reduction based on measured power (“PRBMP”) functionality to Nicol’s system. *Id.*

Kling discloses an implementation of PRBMP whereby the “a parameter that is approximately proportional to the consumed power, *such as voltage* . . .” is tracked. Ex. 1010, 2:31–37. Further, Kling discloses an implementation of PRBMP whereby current consumed by a processor is tracked by measuring “a voltage across a resistor” of the processor as would have been well-known in the art. Ex. 1010, 3:57–59; Ex. 1001 ¶195. If the voltage exceeds a threshold—indicating excessive power consumption—the controller sends a “throttle signal” to the processor to reduce its clock frequency, which reduces power consumption. Ex. 1010, 2:48–51 (explaining that “a processor . . . reduces its core frequency” upon receiving a throttle signal sent by the controller); Ex. 1001, ¶196. Thus, Kling discloses the design incentive of reducing a clock frequency based in part on a voltage in a processor (“processor voltage”), and a POSITA would have understood that PRBMP can also be implemented in Nicol’s system with a reasonable expectation of success to yield predictable results. Ex. 1001 ¶197.

4. Dependent Claim 3

a) ***“The method of claim 1 wherein:”***

See supra Section XI.A.2.

b) “said clock frequency is determined in part by available power.”

Nicol in view of Kling discloses this claim element. Ex. 1001 ¶198. As discussed above in Section XI.A.1 and XI.A.3, a POSITA would have found it obvious to add Kling’s power control functionality to Nicol’s system. Ex. 1001 ¶199. Kling teaches measuring the power (based on the current or voltage across an appropriate resistor) of a processor and comparing the measured power against an “upper limit” threshold, as shown in Figure 2A (annotated) below. Ex. 1001, ¶200; Ex. 1010, Figure 2A.




Figure 2A

Ex. 1010, Figure 2A (annotated).

The upper limit in Figure 2A is the available power to the processor. Ex. 1010, 5:21–25 (explaining that “the upper limit is a constant value ... associated with an approximately maximum power that can be reliably consumed by the computer system ...”); Ex. 1001 ¶¶202–203. If the measured power reaches the upper limit, the controller sends a “throttle signal” to the processor to instruct it to reduce its clock frequency to reduce power consumption. Ex. 1010, 5:38–40, 2:48–51. A POSITA would have understood that this functionality can be readily implemented in Nicol’s system with reasonable expectation of success. Ex. 1001 ¶204. The functionality involves relatively simple add-on software to Nicol’s more complex task scheduler software residing in the controller, and relies on well-known measurement techniques, such as measuring a voltage across an appropriately placed resistor, and would have yielded predictable results. Ex. 1001, ¶¶205–206. In addition, a POSITA would have been motivated to add the functionality to Nicol’s system to implement the design incentives of protecting against over-temperature by controlling power consumption, consistent with Nicol’s objective. Ex. 1001, ¶206.

5. Dependent Claim 4

a) “The method of claim 1 wherein:”

See supra Section XI.A.2.

b) “said clock frequency is determined in part by a user setting.”

Nicol in view of Kling renders obvious this claim element. Ex. 1001 ¶207. As discussed above in Section XI.A.1, XI.A.3 and XI.A.4, a POSITA would have found it obvious to add Kling’s power control functionality to Nicol’s system. Ex. 1001 ¶208. Kling teaches measuring the power (current or voltage across an appropriate resistor) of a processor and comparing the measured power against an “upper limit” threshold as well as a “lower limit” threshold, as shown in Figure 2A (annotated) below. The thresholds are user settings. Ex. 1010, 5:57–59 (“Alternatively, *the lower threshold may be modifiable by a user* or automatically adjusted within the computer system to, for example, reduce power oscillation.”).




Figure 2A

Id., Figure 2A (annotated).

If the measured power reaches the upper limit, the controller sends a “throttle signal” to the processor to instruct it to reduce its clock frequency to reduce power consumption. Ex. 1010, 5:38–40, 2:48–51; Ex. 1001, ¶211. “Once [the] lower threshold is reached, the throttle signal may be deasserted, and, in response, the IC resumes normal operation.” *Id.*, 5:48–50. In this system, the clock frequency is determined in part by user settings, *e.g.*, the upper and lower power limits. Ex. 1001 ¶213.

6. Dependent Claim 5

a) “The method of claim 1 wherein:”

See supra Section XI.A.2.

b) “said clock frequency is determined in part by available power and a user setting.”

First, Nicol in combination with Kling discloses a system wherein the clock frequency is determined in part by available power. *See supra* Section XI.A.4; Ex. 1001 ¶215.

Second, Nicol in combination with Kling discloses a system wherein the clock frequency is determined in part by a user setting. *See supra* Section XI.A.5; Ex. 1001 ¶216.

7. Dependent Claim 6

a) “The method of claim 1 wherein:”

See supra Section XI.A.2.

b) “said number of pending operations is determined at by the fill level of one or more buffers.”

Nicol renders obvious this claim element. Ex. 1001 ¶217. Nicol teaches estimating processing loads based on the number of instructions of each of the “to do” tasks to be allocated to the PEs:

If the ***number of instructions*** that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the tasks to the available processors in order to balance the computation.

Ex. 1009, 3:13–18. A POSITA would have understood that the instructions that need to be executed are stored in an instruction queue or buffer, and that the fill level of the buffer corresponds to the number of instructions. Ex. 1001 ¶219, citing Ex. 1022 at 155 and Figure 3.11 (showing an “instruction queue”). Indeed, the applicant stated during the prosecution of the ’605 patent that “one of ordinary skill in the art would understand that ***the fill level of a data buffer [is] representative of the number of waiting operations or processor load.***” Ex. 1004 at 377.

A POSITA would also have known that Nicol detects the number of pending operations of the PEs by inspecting the fill level of buffers; it was well known in the art that the number of pending operations could be determined by monitoring queues, such as the fill level of buffers that feed work to the processors. Ex. 1001 ¶222, citing Ex. 1018 (Cline), 4:57–60 (“Usage can be monitored in various ways,

such as by taking an average of recent workload and assuming the average represents the near-term workload, *or by monitoring queues that feed work to the device.*”), Ex. 1017 (Hong), 1704 (explaining dynamic supply voltage can be adjusted based on “processing load as measured by *the number of data samples queued in the input (or output) buffer.*”). Thus, a POSITA would have known that the number of pending operations or load of the PEs in Nicol are determined by the fill level of a buffer. Ex. 1001 ¶223.

B. Ground II: Claims 1–6 Are Obvious in View of Nicol in Combination with Kling and DeHon Under Petitioner’s Constructions

As explained above in Section IX (“Claim Construction”), Patent Owner proposed the construction of “passing onto one or more other data processing units which are subsequently processing data” for the term “*sequentially processing data*” and proposes a construction of plain and ordinary for the term “*to minimum,*” in the parties co-pending District of Delaware action. Ex. 1016, 10, 15. In Section XI.A, Petitioner explains how the Challenged Claims are unpatentable under Patent Owner’s proposed constructions of the terms (Ground I) “to a minimum” (as plain and ordinary meaning) and “sequentially processing data” (“passing results onto one or more other data processing units which are subsequently processing data”). See *supra* Section XI.A. In this section, Petitioner explains how the Challenged Claims are unpatentable under *Petitioner’s proposed constructions* of the terms (Ground

II): “to a minimum” to mean “to a level no more than is required for the preservation of memory contents or the like,” and “data processing unit” and/or “data processing unit . . . adapted for sequentially processing data” to mean “reconfigurable and sequential data processors where the data results from one processor are fed to another, for each processor to perform a separate computation.”

Under Ground II, Petitioner adds the DeHon reference to satisfy Petitioner’s claim construction requiring that the data processing unit be reconfigurable. DeHon teaches a reconfigurable data processing chip, as detailed under X.C (Overview of DeHon). Ex. 1001 ¶226. Section XI.B.1 below explains the reasons to combine DeHon with Nicol, and thus with Nicol and Kling. Ex. 1001, ¶¶229–45.

Petitioner’s proposed constructions only affects the “data processing unit” and/or “data processing unit . . . adaptable for sequentially processing data” and “to a minimum” claim elements in claim 1. Section XI.B.2 below explains that Nicol in view of Kling and DeHon renders obvious these claim elements. For all other claim elements, Nicol in view of Kling and DeHon satisfies the claim elements the same way as Nicol in view of Kling satisfies them as set forth under Ground I.

1. Reasons to Combine Nicol, Kling, and DeHon

As discussed above, a POSITA would have been motivated to combine Nicol and Kling. *See* Section XI.A.1; Ex. 1001 ¶229. Further, a POSITA would have been motivated to combine the system of Nicol and Kling with Dehon. Ex. 1001 ¶230.

First, POSITA would have found it obvious to combine the Nicol and Kling system with DeHon because Nicol itself provides a motivation to combine its teachings with DeHon. Ex. 1001, ¶231. Nicol explains that its multiprocessor system is intended for use in a variety of digital signal processing applications. Ex. 1009, 2:53–56 (“A real-time operating system resides in PE 100 and allocates tasks to the other PEs *from a mix of many digital signal processing applications.*”). A POSITA would have been motivated to consult and apply DeHon, because DeHon teaches an “inventive microprocessor [that] can serve as the principle ***building block for a wide range of applications*** including personal computers, embedded systems, application-specific computers, and general and special-purpose multiprocessors” Ex. 1012, 2:18–24; Ex. 1001, ¶232. DeHon explains that such an arrangement “yields application acceleration and in system-specific functions.” Ex. 1012, Abstract. For similar reasons, a POSITA would have been motivated to apply DeHon’s teachings of a flexible, reconfigurable, microprocessor chip, as PEs in Nicol’s system, knowing that adding a reconfigurable processor would improve application acceleration in Nicol’s system. Ex. 1001, ¶233. DeHon Figure 1 (annotated, below) illustrates DeHon’s reconfigurable processor chip. *See* Section X.C.

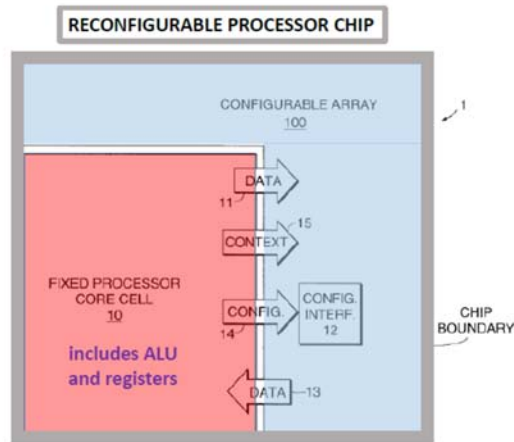
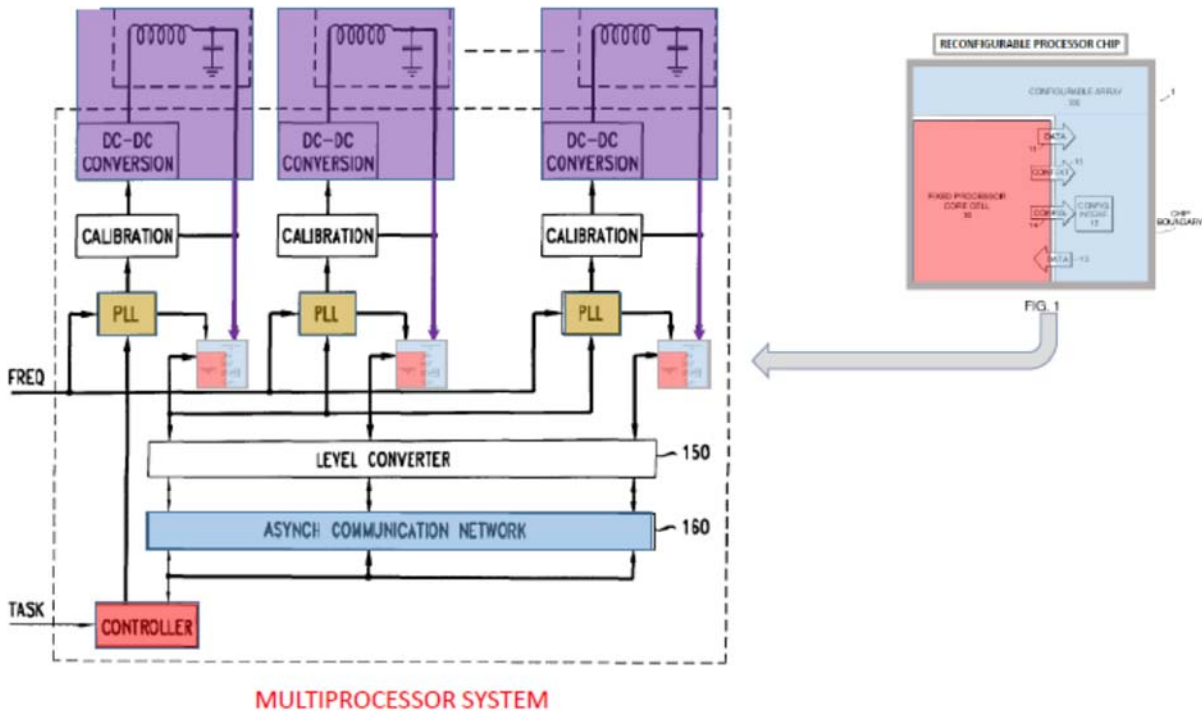


FIG. 1

Ex. 1012, Figure 1 (annotated).

While Nicol’s preferred embodiment integrates all PEs on a multiprocessor chip, Nicol expressly discusses using its dynamic power management scheme in a multiprocessor system having discrete processor chips. Ex. 1009, 6:34–39 (“The principles disclosed above for a multiprocessor is extendible to other system arrangements. This includes systems with a plurality of *separate* processor elements ...”); Ex. 1001, ¶235. A POSITA would have viewed understood the flexibility of Nicol’s system and that it encompasses the use of reconfigurable technology. Ex. 1001, ¶236. It would have been obvious to a POSITA that the PEs disclosed in Nicol’s system can be reconfigurable. Ex. 1001, ¶236. And because a POSITA would have known that Nicol’s PEs could be reconfigurable, it would have been obvious to a POSITA to replace each of Nicol’s PEs with DeHon’s reconfigurable microprocessor chip, illustrated in DeHon’s Figure 1. Ex. 1001, ¶237; Ex. 1012, FIG. 1. The figure below combines Figure 4 of Nicol with Figure 1 of DeHon to

illustrate Nicol's multiprocessor system in which the PEs are DeHon's reconfigurable processor chips. Ex. 1001, ¶238.



Ex. 1009, Figure 4 (left, annotated); Ex. 1012, Figure 1 (right, annotated).

A POSITA would have been motivated to substitute fixed PEs with DeHon's reconfigurable PEs because the resulting reconfigurable multiprocessor system has significantly improved capacity to handle any "mix of many digital signal processing applications" through rapid reconfiguration. Ex. 1009, 2:55–56; Ex. 1012, Abstract ("[Reconfiguration] enables application-specific configurations to allow adaptation to the particular time-changing demands of the microprocessor and provide the functionality required to best serve those demands."); Ex. 1001, ¶240. Such a

substitution would have had a reasonable likelihood of success because Nicol teaches a flexible power management scheme that is agnostic to the types of PEs—whether fixed or configurable, homogeneous or heterogeneous. Ex. 1001, ¶241; Ex. 1009, 6:34–35 (“The principles disclosed above for a multiprocessor is *extendable* to other system arrangements.”).

Second, a POSITA would have found it obvious to combine Nicol and DeHon because DeHon itself provides a motivation to combine its teachings with Nicol. Ex. 1001, ¶242. DeHon explains that a portion of its reconfigurable chip “is dedicated to a *programmable gate array*. This feature enables *application-specific configurations*.” Ex. 1012, 2:12–14, which a POSITA would have connected with Nicol’s “mix of many digital signal processing applications.” Ex. 1009, 2:53–56; Ex. 1001, ¶242. DeHon’s reconfigurable chips can thus be advantageously configured in different ways to handle a variety of compute-intensive functions “such as compression, encryption, signal processing, or protocol processing.” Ex. 1012, 5:63–67. Another advantage is that the configurable DGPA-coupled chips can replace or be deployed alongside traditional fixed processors. *Id.*, 7:20–25. A POSITA would have been motivated to use DeHon’s reconfigurable chips in Nicol’s multiprocessor system because DeHon’s reconfigurable chips are designed to replace or be deployed alongside traditional fixed processors, which minimizes integration risks. Ex. 1001, ¶243.

Third, a POSITA would have been motivated to combine Nicol and DeHon because both relate to the same field of invention, optimizing processor systems by dynamically controlling processors. Ex. 1009, Abstract, [Claim 45]; Ex. 1012, 5:61–67; Ex. 1001, ¶244. “When the references are in the same field as that of the applicant’s invention, knowledge thereof is presumed.” *In re Dance*, 160 F.3d at 1343 (Fed. Cir. 1998). Specifically, DeHon relates to dynamically reconfiguring the programmable array dependent on the particular application. Ex. 1012, 5:61–63 (“Each application program configures the reconfigurable array in a manner most beneficial to that application.”). Similarly, Nicol’s controller dynamically allocates tasks to the individual PEs based on the application to be performed, like the configuration of the array in DeHon, and dynamically programs the respective clock frequencies and supply voltages for each PE. Ex. 1009, Abstract; Ex. 1001, ¶244.

2. Independent Claim 1

The two claim elements with claim terms in dispute are discussed below.

a) “and (b) being adapted for sequentially processing data, the method comprising:”

Under Petitioner’s construction of the term: “data processing unit” and/or “data processing unit . . . adaptable for sequentially processing data,” DeHon in combination with Nicol and Kling discloses this limitation. *See supra* Section XI.A.1.d. DeHon in combination with Nicol and Kling discloses “reconfigurable

and sequential data processors where the data results from one processor are fed to another, for each processor to perform a separate computation.” Ex. 1001 ¶247.

First, the combination of Nicol, Kling, and DeHon discloses reconfigurable processors. Ex. 1001 ¶248. For example, DeHon depicts in Figure 1 a microprocessor with a fixed core 10 (red) connected to a reconfigurable array 100 (blue). Ex. 1012, Figure 1; *see also id.*, 2:25–27 (“In specific implementations, the programmable gate array is configurable by the processor rather than some off-chip configurer.”); [Claim 19].

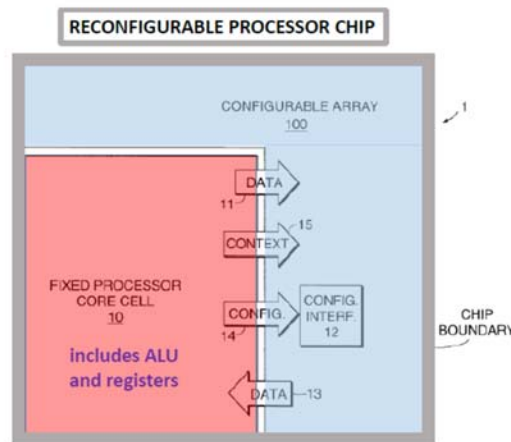
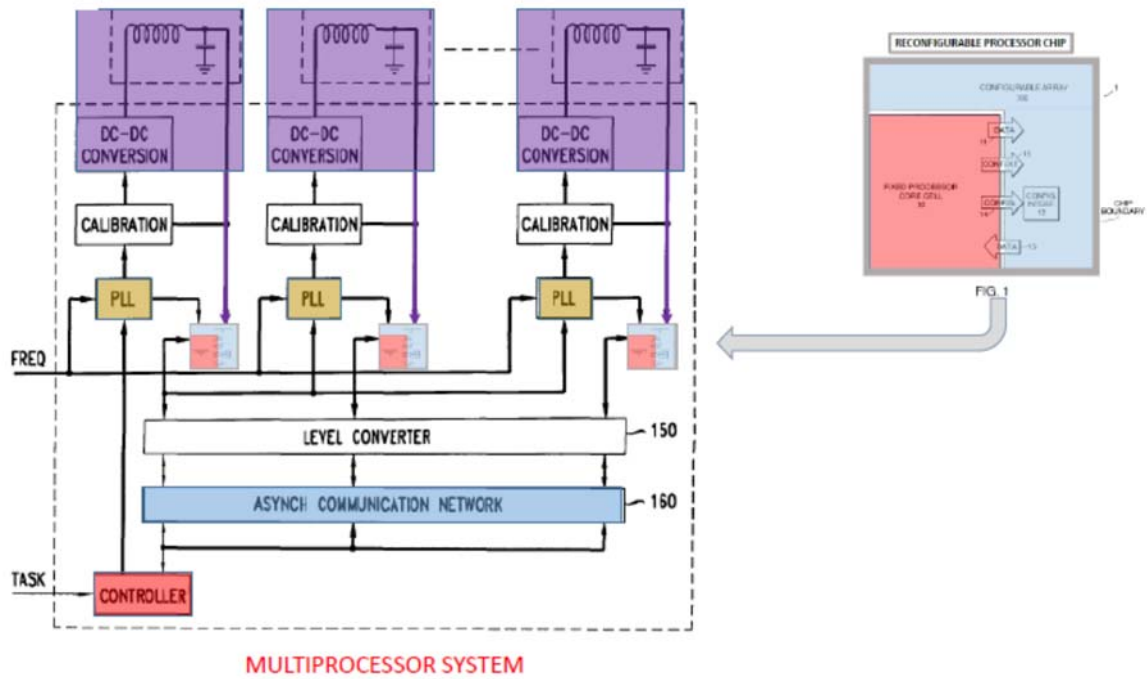


FIG. 1

Ex. 1012, Figure 1 (annotated).

Second, in Nicol’s multiprocessor systems that use DeHon reconfigurable PEs, illustrated below, the DeHon reconfigurable PEs can be configured as “sequential data processors where the data results from one processor are fed to another, for each processor to perform a separate computation.” Ex. 1013, 9; Ex. 1001 ¶251.



Ex. 1009, Figure 4 (left, annotated); Ex. 1012, Figure 1 (right, annotated).

Nicol teaches that the PEs in its multiprocessor system—which are DeHon’s reconfigurable PEs in this combination—are interconnected through a communication network over which they “exchange data” to perform “a mix of many digital signal processing systems.” Ex. 1009, 2:55–56, 6:26–33, Figure 4; Ex. 1001, ¶252. Such data exchange allows data to flow from one PE (the source PE) to another PE (destination PE) for further processing, as was common in digital signal processing systems. Ex. 1001, ¶253. Notably, a POSITA would have known that the DeHon reconfigurable PEs can be configured to sequentially process data in a manner similar to the well-known systolic architecture. Ex. 1001, ¶252, citing Ex. 1014 (Kung, H.T., *Why Systolic Architectures?*, Computer Vol. 15 (Pub. January

1982)) (“Kung”), 45 (describing “numerous systolic designs [] known” in the 1980s), 37 (“In a systolic system, data flows from the computer memory in a rhythmic fashion, *passing through many processing elements before it returns to memory*, much as blood circulates to and from the heart.”). Kung Fig. 1 below illustrates PEs that are arranged in a systolic array and pass data from one PE to the next, with each PE performing a separate operation on the received data.

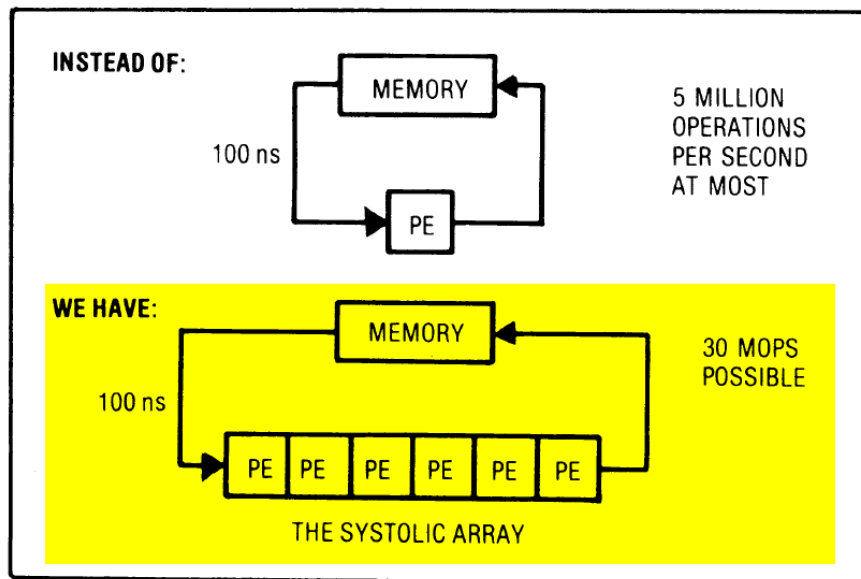


Figure 1. Basic principle of a systolic system.

Id., Figure 1 (annotated).

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

As previously discussed, Nicol discloses this claim element under Patent Owner’s Construction. *See supra* Section XI.A.2.e.

Nicol renders obvious this claim element under Petitioner’s construction of “to a minimum” to mean “to a level no more than is required for the preservation of memory contents or the like.” Ex. 1001 ¶258. Nicol teaches minimizing power consumption of a chip, whereby the controller PE can “independently program the lowest operating frequency and corresponding lowest supply voltage for each PE.” Ex. 1009, 5:66–6:2; Ex. 1001, ¶259. Nicol further discloses that its operating system “employs gated clocking techniques to ‘shut down’ PEs that are not needed.” *Id.*, 5:42–45.

Because Nicol is directed to power management, a POSITA would have been motivated to consult the power and thermal requirements specified in “the Advanced Configuration and Power Interface (ACPI) Specification.” Ex. 1001, ¶261; Ex. 1020 (ACPI Spec.); *see also* Ex. 1011 (Bhatia), 1:6–17 (explaining “the Advanced Configuration and Power Interface (ACPI) Specification, Rev. 1.0., published on Dec. 22, 1996 [] provides an interface between the operating system of a system and hardware devices to implement power and thermal management.”). The ACPI was (and still is) an industry standard specification, and would have been well-known to a POSITA. Ex. 1001, ¶¶261–62.

Notably, the ACPI specification teaches a low power mode, defined as the C3 state. Ex. 1001, ¶262; Ex. 1011, 12:58–65 (“the C3 state defined under 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.*”). A POSITA 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 C3 state is a low power sleep mode that reduces power consumption when the computer system is inactive and in a sleep mode, yet because it preserves cache memory contents, allows the computer system to regain full functionality quickly upon wakeup. Ex. 1001, ¶264. In essence, the C3 state is akin to a “light” sleep state. A POSITA would also have been motivated to conform to ACPI and implement ACPI's C3 state in any computer system, including Nicol's system. Ex. 1001, ¶¶264-65.

C. Ground III: Claims 1-6 Are Obvious In View of Nicol in Combination with Kling, DeHon, and Bhatia Under Both Parties' Constructions

In Section XI.B *supra*, Petitioner explains that Nicol in view of Kling and DeHon renders obvious Challenged Claims 1–6 under Petitioner's claim constructions (Ground II). With regard to the claim 1 limitation requiring “the multiprocessor system setting a clock frequency . . . to a minimum . . .” (“to a minimum” limitation), Petitioner explains that the limitation is met under Petitioner's construction of “setting a clock frequency . . . to a minimum” to mean “setting the frequency to a level no more than is required for the preservation of memory contents or the like” even though Nicol, Kling and DeHon do not expressly

disclose such a setting. *See* Section XI.B.2 *supra*. This is because setting a clock frequency “to a level no more than is required for the preservation of memory contents” corresponds to a power-saving mode that was not only well-known and prevalent in the computer industry, it was specified in the 1996 *Advanced Configuration and Power Interface* (ACPI) industry specification and described in detail in other patents and publications such as Bhatia. Ex. 1001, ¶267; Ex. 1011, 1:6–17. Underlying Ground II is Petitioner’s view that, even though the combination of Nicol in view of Kling and DeHon does not expressly disclose “setting the frequency to a level no more than is required for the preservation of memory contents or the like,” common knowledge in the art and the prevalent use of such a low power mode setting renders the “to a minimum” limitation obvious. Ex. 1001, ¶268.

In any event, the Challenged Claims are also rendered obvious under the instant Ground III. The only difference between Ground II, discussed under Section XI.B *supra* and the instant Ground III is that, under Ground III, Petitioner ***explicitly combines*** Nicol, Kling and Dehon ***with Bhatia*** for the limited reason that Bhatia describes the ACPI’s C3-state, a low power mode setting that satisfies the “to a minimum” limitation under Petitioner’s construction, and thus also satisfying Patent Owner’s construction of plain and ordinary meaning:

[In] the C3 state defined under the ACPI specification, 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***

Ex. 1011 (Bhatia), 12:61–64; Ex. 1001, ¶269. Bhatia teaches “setting the frequency to a level no more than is required for the preservation of memory contents or the like” (as required under Petitioner’s construction of “to a minimum”) because, under the C3-state specified in ACPI and discussed in Bhatia, the processor clock is disabled and the processor does not perform any activity except maintenance of the cache memory contents. Ex. 1001, ¶270. The C3-state also satisfies Patent Owner’s plain meaning construction of “to a minimum” because the clock frequency is zero (clock disabled). Ex. 1001 ¶270. Thus, in light of the analysis set forth under Grounds I and II, the Ground III combination of Nicol, Kling, DeHon and Bhatia renders obvious all Challenged Claims under both parties’ claim constructions. Ex. 1001, ¶271.

A POSITA would have been motivated to consult Bhatia because Nicol, Kling and Bhatia are all directed to interrelated thermal and power management aspects that are critical and ubiquitous in processor-based systems, and the techniques can all be used to control the operation of processors such as DeHon’s processors. Ex. 1001 ¶272; Ex. 1011 (Bhatia), 1:6–7 (“Different types of power and thermal management techniques have been implemented in systems.”). In addition, Bhatia describes in detail “low activity state[s] (e.g., deep sleep, stop grant, C1, C2, or C3 state)” commonly used by power management systems to reduce power, but not

described in much detail in Nicol or Kling. Ex. 1011, 12:58–64 (citing the *Advanced Configuration and Power Interface* (ACPI) industry standard specification); Ex. 1001 ¶273. A POSITA would have been motivated to consult Bhatia and implement those low activity states in Nicol’s system as additional means to minimize power consumption. Ex. 1001 ¶274. A POSITA would have had a reasonable expectation of success in doing so because Bhatia explains that those low activity states were commonly used in the industry. Ex. 1001 ¶274.

XII. CONCLUSION

Petitioner Intel requests institution of an *inter partes* review and cancellation of the Challenged Claims.

Date: February 7, 2020

Respectfully Submitted,

/s/ Kevin Bendix

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

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

CERTIFICATE OF COMPLIANCE

This Petition complies with the type-volume limitations as mandated in 37 C.F.R. § 42.24, totaling 11,249 words. Counsel has relied upon the word count feature provided by Microsoft Word.

/s/ Kevin Bendix
Kevin Bendix (Reg. No. 67,164)

CERTIFICATE OF SERVICE

The undersigned hereby certifies that a copy of the foregoing document with accompanying Exhibits 1001 – 1026, as listed in Petitioner’s Exhibit List on pages iv - v of this document, was served on February 7, 2020 via electronic mail and FTP, upon agreement under 37 CFR § 42.105, to the attorney of record for the patent owner at the following address:

Aaron Grunberger
Norton Rose Fulbright LLC
1301 Avenue of the Americas
New York, New York 10019-6022
aaron.grunberger@nortonrosefulbright.com

A copy was also served via electronic mail and FTP on the attorneys of record listed below for the following related matters:

PACT XPP Schweiz AG v. Intel Corporation, Civil Action No. 19-1066 (JDW)
(D. Del).

Brian E. Farnan, Esquire
Michael J. Farnan, Esquire
FARNAN LLP
919 North Market Street, 12th Floor
Wilmington, DE 19801
(302) 777-0300
bfarnan@farnanlaw.com
mfarnan@farnanlaw.com

Frederick A. Lorig, Esquire
Danielle L. Gilmore, Esquire
Pushkal Mishra, Esquire
Nima Hefazi, Esquire
QUINN, EMANUEL, URQUHART &

Petition for *Inter Partes* Review of U.S. Patent No. 9,075,605

SULLIVAN, LLP
865 S. Figueroa Street, 10th Floor
Los Angeles, CA 90017
(213) 443-3047
fredericklorig@quinnemanuel.com
daniellegilmore@quinnemanuel.com
pushkalmishra@quinnemanuel.com
nimahefazi@quinnemanuel.com

Mark Tung, Esquire
QUINN, EMANUEL, URQUHART &
SULLIVAN, LLP
555 Twin Dolphin Drive, 5th Floor
Redwood Shores, CA 94065
(650) 801-5016
marktung@quinnemanuel.com

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

Attorneys for Plaintiff

/s/ Kevin Bendix
Kevin Bendix