

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO., LTD.,
SAMSUNG ELECTRONICS AMERICA, INC., and QUALCOMM, INC.,
Petitioners

v.

DAEDALUS PRIME LLC,
Patent Owner.

Case No. IPR2023-00567
U.S. Patent No. 10,049,080

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 10,049,080**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8.....	1
III. GROUNDS FOR STANDING.....	4
IV. FEE AUTHORIZATION	4
V. PRECISE RELIEF REQUESTED	4
VI. THE CHALLENGED PATENT	5
VII. PROSECUTION HISTORY	7
VIII. THERE IS NO BASIS FOR DISCRETIONARY DENIAL	8
A. <i>Advanced Bionics</i> Part I	8
B. <i>Advanced Bionics</i> Part II	11
C. Mathieson and Sutardja Render Claims 1, 4, 7-9, 12, 15-17, 20, and 23-24 Unpatentable	12
IX. LEVEL OF ORDINARY SKILL IN THE ART	13
X. PRIORITY DATE	13
XI. CLAIM CONSTRUCTION	14
XII. BRIEF DESCRIPTION OF THE APPLIED PRIOR ART REFERENCES	14
A. Mathieson (Ex-1005).....	14
B. Carmack (Ex-1006)	15
C. Sutardja (Ex-1007, Ex-1008)	17
D. Rychlik (Ex-1009).....	17
XIII. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS	19
A. Ground 1: Claims 1-4, 7-12, 15-20, 23-24 Are Rendered Obvious By Sutardja (Ex-1007, Incorporating Ex-1008).....	19
1. Independent Claim 1	19
a. Element 1[preamble]: A multi-core processor comprising:	19

TABLE OF CONTENTS

(continued)

Page

b.	Element 1[a][i]: a first plurality of cores and a second plurality of cores that support a same instruction set,.....	20
c.	Element 1[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and.....	23
d.	Element 1[b][i]: power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores,.....	24
e.	Element 1[b][ii]: wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.....	31
2.	Dependent Claim 2: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.	33
3.	Dependent Claim 3: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that consume less power than corresponding logic gates of the first plurality of cores.	34
4.	Dependent Claim 4: The multi-core processor of claim 1, wherein the second plurality of cores each have a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.	35
5.	Dependent Claim 7: The multi-core processor of claim 1, wherein the first plurality of cores are at a maximum operating frequency in the state.	35

TABLE OF CONTENTS

(continued)

	Page
6. Dependent Claim 8: The multi-core processor of claim 1, wherein.....	36
a. Element 8[a]: the power management hardware is to enable all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores, ...	36
b. Element 8[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.	39
7. Independent Claims 9 and 17:.....	40
a. Element 9[preamble]: A method comprising:	40
b. Element 17[preamble]: A non-transitory machine readable medium containing program code that when processed by a machine causes a method to be performed, the method comprising:.....	40
c. Elements 9[a][i] and 17[a][i]: operating a multi-core processor such that a first plurality of cores and a second plurality of cores execute a same instruction set,.....	40
d. Elements 9[a][ii] and 17[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and.....	40
e. Elements 9[b][i] and 17[b][i]: disabling with power management hardware, from a state where the first plurality of cores and the second plurality of cores are enabled, all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores, ...	41

TABLE OF CONTENTS

(continued)

Page

f.	Element 9[b][ii] and 17[b][ii]: wherein an operating system executing on the multi-core processor monitors a demand for the multi-core processor and controls the power management hardware based on the demand.....	41
8.	Dependent Claims 10 and 18: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.	41
9.	Dependent Claims 11 and 19: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that consume less power than corresponding logic gates of the first plurality of cores.	41
10.	Dependent Claims 12 and 20: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the second plurality of cores at a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.	42
11.	Dependent Claims 15 and 23: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the first plurality of cores at a maximum operating frequency in the state.....	42
12.	Dependent Claims 16 and 24: The [method of claim 9/non-transitory machine readable medium of claim 17], further comprising.....	42
a.	Elements 16[a] and 24[a]: enabling, with the power management hardware, all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores, ...	42

TABLE OF CONTENTS

(continued)

	Page
b. Elements 16[b] and 24[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.....	42
B. Ground 2: Claims 5-6, 13-14, and 21-22 Are Rendered Obvious By Sutardja/Rychlik	43
1. Dependent Claims 5, 13, and 21:	43
a. Elements 5[a], 13[a], 21[a]: The [multi-core processor of claim 1/method of claim 9/non-transitory machine readable medium of claim 17], [wherein the power management hardware is to disable/further comprising disabling, with the power management hardware,] an additional core of the second plurality of cores for each continued drop in demand below a next lower threshold until one core of the second plurality of cores remains enabled, and	43
b. Elements 5[b], 13[b], 21[b]: [lower/lowering] an operating frequency or a supply voltage of the one core of the second plurality of cores as demand drops below a next lower threshold.....	46
2. Dependent Claims 6, 14, 22: [The multi-core processor of claim 5/method of claim 13/The non-transitory machine readable medium of claim 21], [wherein the power management hardware is to raise/further comprising raising, with the power management hardware,] a supply voltage or an operating frequency of said one core in response to higher demand.....	48
C. Ground 3: Claims 7, 15, and 23 Are Rendered Obvious By Sutardja/Carmack	48
1. Dependent Claims 7, 15, 23: The [multi-core processor of claim 1/method of claim 9/non-transitory machine readable medium of claim 17], wherein [the operating comprises operating] the first plurality of cores [are] at a	

TABLE OF CONTENTS

(continued)

	Page
maximum operating frequency in the state.....	48
D. Ground 4: Claims 1-4, 7-12, 15-20, and 23-24 Are Rendered Obvious By Mathieson/Sutardja	50
1. A POSITA would have been motivated to combine Mathieson with Sutardja	50
2. Independent Claim 1	51
a. Element 1[preamble]: A multi-core processor comprising:	51
b. Element 1[a][i]: a first plurality of cores and a second plurality of cores that support a same instruction set,.....	52
c. Element 1[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and.....	55
d. Element 1[b][i]: power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores,.....	56
e. Element 1[b][ii]: wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.....	62
3. Dependent Claim 2: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.	64
4. Dependent Claim 3: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that consume less power than corresponding logic gates of	

TABLE OF CONTENTS

(continued)

	Page
the first plurality of cores.	67
5. Dependent Claim 4: The multi-core processor of claim 1, wherein the second plurality of cores each have a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.	69
6. Dependent Claim 7: The multi-core processor of claim 1, wherein the first plurality of cores are at a maximum operating frequency in the state.	70
7. Dependent Claim 8: The multi-core processor of claim 1, wherein.....	72
a. Element 8[a]: the power management hardware is to enable all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores, ...	72
b. Element 8[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.	77
8. Independent Claims 9 and 17:.....	77
a. Element 9[preamble]: A method comprising:	77
b. Element 17[preamble]: A non-transitory machine readable medium containing program code that when processed by a machine causes a method to be performed, the method comprising:.....	77
c. Elements 9[a][i] and 17[a][i]: operating a multi-core processor such that a first plurality of cores and a second plurality of cores execute a same instruction set,.....	77
d. Elements 9[a][ii] and 17[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and.....	78

TABLE OF CONTENTS

(continued)

	Page
e. Elements 9[b][i] and 17[b][i]: disabling with power management hardware, from a state where the first plurality of cores and the second plurality of cores are enabled, all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores, ... f. Element 9[b][ii] and 17[b][ii]: wherein an operating system executing on the multi-core processor monitors a demand for the multi-core processor and controls the power management hardware based on the demand.....	78
9. Dependent Claims 10 and 18: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.	78
10. Dependent Claims 11 and 19: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that consume less power than corresponding logic gates of the first plurality of cores.	79
11. Dependent Claims 12 and 20: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the second plurality of cores at a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.	79
12. Dependent Claims 15 and 23: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the first plurality of cores at a maximum operating frequency in the state.....	79
13. Dependent Claims 16 and 24: The [method of claim 9/non-	

TABLE OF CONTENTS
(continued)

	Page
transitory machine readable medium of claim 17], further comprising.....	79
a. Elements 16[a] and 24[a]: enabling, with the power management hardware, all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores, ...	79
b. Elements 16[b] and 24[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.....	80
E. Ground 5: Claims 5-6, 13-14, and 21-22 Are Rendered Obvious By Mathieson/Sutardja/Rychlik	80
1. Dependent Claims 5, 13, and 21:	80
a. Elements 5[a], 13[a], 21[a]	80
b. Elements 5[b], 13[b], 21[b]	81
2. Dependent Claims 6, 14, 22	81
XIV. CONCLUSION.....	82

LIST OF EXHIBITS¹

Ex-1001	U.S. Patent No. 10,049,080 to George et al. (“the ’080 Patent”)
Ex-1002	Declaration of Dr. Trevor Mudge
Ex-1003	Curriculum Vitae of Dr. Trevor Mudge
Ex-1004	Prosecution History of the ’080 Patent (Application No. 15/431,527)
Ex-1005	U.S. Patent Pub. No. 2011/0213950 to Mathieson et al. (“Mathieson”)
Ex-1006	U.S. Patent Pub. No. 2009/0309243 to Carmack et al. (“Carmack”)
Ex-1007	U.S. Patent Pub. No. 2008/0288748 to Sutardja et al. (“Sutardja ’748”)
Ex-1008	U.S. Patent Pub. No. 2007/0083785 to Sutardja (“Sutardja ’785”)
Ex-1009	U.S. Patent Pub. No. 2011/0145615 to Rychlik et al. (“Rychlik”)
Ex-1010	Prosecution History of U.S. Patent No. 9,569,278 (“the ’278 Patent”)
Ex-1011	INTENTIONALLY LEFT BLANK
Ex-1012	INTENTIONALLY LEFT BLANK
Ex-1013	INTENTIONALLY LEFT BLANK
Ex-1014	INTENTIONALLY LEFT BLANK
Ex-1015	INTENTIONALLY LEFT BLANK
Ex-1016	INTENTIONALLY LEFT BLANK
Ex-1017	INTENTIONALLY LEFT BLANK
Ex-1018	INTENTIONALLY LEFT BLANK
Ex-1019	INTENTIONALLY LEFT BLANK

¹ Four-digit pin citations that begin with 0 are to the branded numbers added by Samsung in the bottom right corner of the exhibits. All other pin citations are to original page, column, paragraph, or line numbers.

Ex-1020	Claim Mapping Table
Ex-1021	INTENTIONALLY LEFT BLANK
Ex-1022	U.S. Patent Pub. No. 2006/0095807 to Grochowski (“Grochowski”)
Ex-1023	U.S. Patent Pub. No. 2012/0317568 to Aasheim (“Aasheim”)
Ex-1024	Jeffrey C. Mogul et al., <i>Operating Systems and Asymmetric Single-ISA CMPs: The Potential for Saving Energy</i> , Hewlett-Packard Development Company, L.P. (2007)
Ex-1025	Juan Carlos Saez et al., <i>Operating System Support for Mitigating Software Scalability Bottlenecks on Asymmetric Multicore Processors</i> , ACM 978-1-4503-004-5/10/05 (2010)
Ex-1026	U.S. Patent No. 7,093,147 to Farkas et al. (“Farkas”)
Ex-1027	Charles Lefurgy et al., <i>Energy Management for Commercial Servers</i> , Computer 39 (Dec. 2003).
Ex-1028	Yushi Shen et al., <i>Enabling the New Era of Cloud Computing: Data Security, Transfer, and Management</i> (Information Science Reference 2014).
Ex-1029	Stefanos Kaxiras and Margaret Martonosi, <i>Computer Architecture Techniques for Power-Efficiency</i> , in Synthesis Lectures on Computer Architecture #4 (Morgan & Claypool 2008).
Ex-1030	Vasanth Venkatachalam and Michael Franz, <i>Power Reduction Techniques For Microprocessor Systems</i> , 37 ACM Computing Surveys 195 (2005).
Ex-1031	Euiseong Seo et al., <i>Energy Efficient Scheduling of Real-Time Tasks on Multicore Processors</i> , 19 IEEE Transactions on Parallel and Distributed Systems 1540 (Nov. 2008).
Ex-1032	Rakesh Kumar et al., <i>Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction</i> , Proceedings of the 36 th International Symposium on Microarchitecture (MICRO-36 2003), IEEE Computer Society (2003).

Ex-1033	U.S. Patent No. 8,615,647 to Hum et al. (“Hum”)
---------	---

I. INTRODUCTION

Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., and Qualcomm, Inc. (collectively, “Petitioner”) request *inter partes* review (“IPR”) of Claims 1-24 of U.S. Patent No. 10,049,080 (“the ’080 Patent”) (Ex-1001), currently assigned to Daedalus Prime LLC (“PO”).

II. MANDATORY NOTICES UNDER 37 C.F.R. §42.8

Real Parties-in-Interest: Petitioner identifies the following real parties-in-interest: Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., and Co-Petitioner Qualcomm Incorporated, and Qualcomm Technologies, Inc.

Petitioner Qualcomm Incorporated further identifies, without conceding any such party is a Real Party-in-Interest, Mazda Motor Corporation, Mazda North American Operations, Mazda Motor of America, Inc., Mercedes-Benz Group AG, Mercedes-Benz AG, Mercedes-Benz USA, LLC, and Visteon Corporation as parties that Patent Owner has accused of infringing the ’080 Patent based on incorporation of Qualcomm products in *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, 1:22-cv-01109 (D. Del.) and *Certain Semiconductors and Devices and Products Containing the Same, Including Printed Circuit Boards, Automotive Parts, and Automobiles*, ITC Inv. No. 337-TA-1332 (Violation).

Related Matters: Patent Owner has asserted the ’080 Patent against Petitioner in (1) *Daedalus Prime LLC v. Arrow Electronics, Inc. et al.*, 1:22-cv-

01107 (D. Del.); (2) *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, 1:22-cv-01109 (D. Del.); (3) *Daedalus Prime LLC v. Mazda Motor Corporation et al.*, 1:22-cv-01108 (D. Del.); (4) *Daedalus Prime LLC v. Samsung Electronics Co., Ltd. et al.*, 2:22-cv-00352 (E.D. Tex.); (5) *Certain Integrated Circuits, Mobile Devices Containing the Same, and Components Thereof*, U.S. International Trade Commission (“ITC”) Inv. No. 337-TA-1335 (Violation); and (6) *Certain Semiconductors and Devices and Products Containing the Same, Including Printed Circuit Boards, Automotive Parts, and Automobiles*, ITC Inv. No. 337-TA-1332 (Violation).

Lead and Back-Up Counsel:

- Lead Counsel:

William M. Fink (Reg. No. 72,332)
O’Melveny & Myers LLP
1625 Eye Street, NW
Washington, DC 20006
Telephone: (202) 383-5300
Fax: (202) 383-5414
E-Mail: tfink@omm.com

- Backup Counsel:

Ben Haber (Reg. No. 67,129)
Nicholas J. Whilt (Reg. No. 72,081)
Brian Cook (Reg. No. 59,356)
O’Melveny & Myers LLP
400 South Hope Street, 18th Floor
Los Angeles, CA 90071
Telephone: (213) 430-6000

Fax: (213) 430-6407

E-Mail: bhaber@omm.com

E-Mail: nwhilt@omm.com

E-Mail: bcook@omm.com

*Counsel for Petitioners Samsung Electronics Co., Ltd. and
Samsung Electronics America, Inc.*

Daniel Leventhal (Reg. No. 59,576)

Richard Zembek (Reg. No. 43,306)

Darren Smith (Reg. No. 64,261)

Norton Rose Fulbright US LLP

1301 McKinney, Suite 5100

Houston, TX 77010

Telephone: (713) 651-5151

Fax: (713) 651-5246

E-Mail: daniel.leventhal@nortonrosefulbright.com

E-Mail: richard.zembek@nortonrosefulbright.com

E-Mail: darren.smith@nortonrosefulbright.com

Counsel for Petitioner Qualcomm Incorporated

Eagle H. Robinson (Reg. No. 61,361)

Norton Rose Fulbright

98 San Jacinto Blvd., Ste. 1100

Austin, TX 78701

Telephone: (512) 474-5201

Fax: (512) 536-4598

E-Mail: eagle.robinson@nortonrosefulbright.com

Counsel for Petitioner Qualcomm Incorporated

Service Information: Petitioner consents to electronic service by email to

the following addresses:

- tfink@omm.com
- bhaber@omm.com
- nwhilt@omm.com
- bcook@omm.com
- daniel.leventhal@nortonrosefulbright.com
- richard.zembek@nortonrosefulbright.com

- darren.smith@nortonrosefulbright.com
- eagle.robinson@nortonrosefulbright.com

III. GROUNDS FOR STANDING

Petitioner certifies that the '080 Patent is available for review, and Petitioner is not barred or estopped from requesting review.

IV. FEE AUTHORIZATION

The PTO is authorized to charge any fees due during this proceeding to Deposit Account No. LA500639.

V. PRECISE RELIEF REQUESTED

Petitioner requests cancellation of Claims 1-24 under 35 U.S.C. §§103 on the following grounds, supported by a declaration from Dr. Trevor Mudge. Ex-1002.

Ground	Summary
1	Claims 1-4, 7-12, 15-20, and 23-24 are rendered obvious by Sutardja ² alone (Ex-1007, incorporating Ex-1008)
2	Claims 5-6, 13-14, and 21-22 are rendered obvious by Sutardja in view of Rychlik (“Sutardja/Rychlik”)

² For clarity, this ground relies on Sutardja '748 (Ex-1007) incorporating by reference Sutardja '785 (Ex-1008), as a single reference obviousness ground, referred to herein as the “combined Sutardja” or simply “Sutardja.” When reference is made to a particular patent, Ex-1007 refers to Sutardja '748, and Ex-1008 refers to the incorporated document Sutardja '785. To the extent that any relied upon teaching is argued to have not been fully incorporated from Ex-1008 into Ex-1007, it would have been obvious to combine such feature/teaching from Ex-1008 with Ex-1007, for the same reasons provided herein.

Ground	Summary
3	Claims 7, 15, and 23 are rendered obvious by Sutardja in view of Carmack (“Sutardja/Carmack”)
4	Claims 1-4, 7-12, 15-20, and 23-24 are rendered obvious by Mathieson ³ (Ex-1005, incorporating Ex-1006) in view of Sutardja (“Mathieson/Sutardja”)
5	Claims 5-6, 13-14, and 21-22 are rendered obvious by Mathieson/Sutardja in view of Rychlik (“Mathieson/Sutardja/Rychlik”)

VI. THE CHALLENGED PATENT

The '080 Patent is directed to power management in a multi-core processor. Ex-1001, 3:34-62. The '080 Patent explains that “a number of different power management schemes are incorporated into modern day computing systems.” Ex-1001, 2:20-22. Prior art power management schemes include “enabl[ing]/disabl[ing] entire cores and rais[ing]/lower[ing] their supply voltages and operating frequencies in response to system workload,” as illustrated in Figure 2.⁴ Ex-1001, 2:30-33.

³ For clarity, Ground 4 relies on Mathieson (Ex-1005) incorporating by reference Carmack (Ex-1006), referred to herein as simply “Mathieson.” To the extent that any relied upon teaching is argued to have not been fully incorporated from Ex-1006 into Ex-1005, it would have been obvious to combine such teaching from Ex-1006 with Ex-1005, for the same reasons provided herein.

⁴ All annotations and emphasis have been added, unless otherwise noted.

Prior Art Power Management




FIG. 2

The '080 Patent allegedly improves upon the prior art using "[s]ome basic concepts of electronic circuit power consumption," including that "the speed of operation of interconnected logic gates [] rises as the width of its driving transistors [] increase," as does the power consumed by those logic gates. Ex-1001, 2:43-52, 3:26-33. In particular, the '080 Patent implements a multi-core processor in which one or more cores are "designed to be lower performance and therefore consume less power than other cores," while nevertheless supporting the same instruction set, purportedly to achieve greater power savings. Ex-1001, 3:50-62, 4:20-29.

But multi-core processors with both high-power and low-power cores, each supporting the same instruction set, were well-known in the prior art. Ex-1002, ¶¶34, 54-57. Indeed, as explained below, the '080 Patent's claims would have been obvious.

VII. PROSECUTION HISTORY

The '080 Patent was filed February 13, 2017, as a continuation of Application No. 13/335,257 (now U.S. 9,569,278), filed December 21, 2011.

During prosecution, the Examiner rejected pending claims corresponding to issued claims 1, 4-9, 12-17, and 20-24 over U.S. Patent Pub. No. 2006/0095807 to Grochowski and pending claims corresponding to issued claims 2-3, 10-11, and 18-19 over Grochowski in view of U.S. Patent Pub. No. 2008/0263324 (Sutardja '324). Ex-1004, 0067-0072. Applicant argued that Grochowski failed to teach all of the limitations of the independent claims:

§§102 and 103 Rejections

Merely in the interest of compact prosecution, the Applicant notes that the cited portions of Grochowski do not appear to teach or suggest the Applicant's claims. For example, the cited portions of Grochowski do not appear to teach or suggest:

20. A multi-core processor comprising:

a first plurality of cores and a second plurality of cores that support a same instruction set, wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores, wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand. (Emphasis added.)

Ex-1004, 0056. Notably, Applicant did not dispute that Sutardja '324 teaches the transistor-related limitations of dependent Claims 2-3, 10-11, and 18-19. Ex-1004, 0054-0057.

As reflected in the file history, the purported novelty of the claims lay in the power management hardware. However, as explained in the Grounds below, this limitation (indeed all limitations of Claims 1-24) was well-known.

VIII. THERE IS NO BASIS FOR DISCRETIONARY DENIAL


The Board should not deny institution under 35 U.S.C. §314(a). The '080 Patent has not been challenged in any other post-grant proceeding or reexamination. The only proceedings relating to the '080 Patent with scheduled hearing dates are in the ITC (the concurrent District court litigations are stayed pending the ITC proceedings), however, concurrent ITC proceedings are not a basis to discretionarily deny IPR institution. *See* USPTO Memorandum by Katherine Vidal, June 21, 2022, *Interim Procedure For Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation*, at 5-7.

Nor should the Board deny institution under 35 U.S.C. §325(d), because both parts of the *Advanced Bionics, LLC v. Med-El Elektromedizinische Geräte GMBH* framework weigh against denial. IPR2019-01469, Paper 6, at 8 (PTAB Feb. 13, 2020) (precedential).

A. *Advanced Bionics* Part I

First, to the best of Petitioner's knowledge, none of Sutardja '785 (Ex-1008), Mathieson (Ex-1005), Carmack (Ex-1006), or Rychlik (Ex-1009) were presented to or considered by the Examiner during prosecution. Second, although Sutardja '748

(Ex-1007) and its parent application, Sutardja '324, are cited on the face of the '080 Patent, the Office never considered Sutardja '785 (Ex-1008), incorporated by reference into Sutardja '748 (Ex-1007) (i.e., the combined Sutardja applied in Ground 1). Sutardja '785 was not discussed and is not found in the history of the '080 Patent's nor its parent application. Sutardja '748 (Ex-1007) is a continuation of Sutardja '324, and it is also a continuation-in-part of Sutardja '785 (Ex-1008), which contains several important additional disclosures as discussed further in this petition. *See, e.g.*, Section XIII.A.1.b, XIII.A.1.d, XIII.A.6. The relationship between the different Sutardja references is shown below.



Thus, the combined Sutardja Ground (Ex-1007 incorporating Ex-1008), providing material and non-cumulative disclosures of a plurality of lower power cores, was not previously considered.⁵

Nor is the combined Sutardja Ground, or the separate and distinct Mathieson Ground, presented herein the same, substantially the same, or cumulative of Grochowski, relied upon by the Examiner in rejecting the independent claims of the '080 Patent. For instance, the Examiner relied on Grochowski's teaching of "parallelism in the program" and situations in which "parallelism increased" as teaching the claimed "demand." Ex-1004, 0069-0070. But Grochowski discloses **enabling** high-power (A) cores and **disabling** low-power (B) cores in response to decreasing "parallelism." Ex-1022, [0036], Table 1. In contrast, Sutardja discloses demand in the form of "system loading," and its multi-core processor exhibits the opposite behavior—high-power cores are **disabled** and low-power cores remain **enabled** in response to decreasing loading. Ex-1007, ¶¶259, 286; Ex-1008, ¶237, Figure 33; Section XII.A.1.d. Nor is Mathieson, in Ground 4, based on "parallelism" as Grochowski is. Mathieson instead includes express teachings to consider

⁵ The Office considered Sutardja '748 in rejecting the claims of the '080 Patent's parent application, but the applicant distinguished Sutardja '748 by arguing it lacked a plurality of lower power cores. Ex-1010, 0125, 0135. As explained below in detail, a plurality of lower power cores is expressly disclosed by Sutardja '785. Ex-1008 ¶¶215, 223, Fig. 2A.

“workloads” against “threshold” values to determine the claimed “demand.” *See, e.g.,* Section XIII.D.2.d. None of these teachings relating to workloads, system loading parameters, and thresholds is cumulative of the Grochowski’s “parallelism” considerations as applied by the Examiner because they operate in substantially different ways than Grochowski.

Accordingly, *Becton, Dickinson* factors (a) (material differences), (b) (non-cumulative nature), and (d) (lack of overlap) weigh against discretionary denial under part one of *Advanced Bionics*. IPR2017-01586, Paper 8 at 17-18. Because the first part is not met, the Board need not proceed to the second part of the *Advanced Bionics* framework.

B. *Advanced Bionics* Part II

There is no evidence of Sutardja ’785 (Ex-1008) in the record. Nonetheless, if the Board were to consider the combined Sutardja under part-two of the *Advanced Bionics* framework, it would have been a material error for the examiner to have overlooked the teachings of Sutardja ’785 (Ex-1008) because the combined Sutardja, which includes Sutardja ’785, teaches all limitations of the independent claims as explained in Ground 1 below. *See, e.g., Commscope Tech. LLC v. Dali Wireless Inc.*, IPR2020-01473, Paper 18, at 16-17 (PTAB Mar. 12, 2021) (“Examiner erred by either overlooking, or failing to appreciate, the materiality of the disclosures” in prior art before it). Specifically, even if deemed presented, it was material error not to

apply Sutardja '324 or '748 incorporating Sutardja '785 (Ex-1008) in a rejection because fully-incorporated Sutardja discloses and teaches the claims pending during prosecution. The incorporated Sutardja '785 discloses what the applicant claimed was missing from the prior art: the plurality of low power cores. *See supra* n. 5. Thus under *Advanced Bionics* part 2, it was a material error to have overlooked the combined Sutardja disclosures invalidating these claims as explained in this Petition.

C. Mathieson and Sutardja Render Claims 1, 4, 7-9, 12, 15-17, 20, and 23-24 Unpatentable

Additionally, under the Board's *SAS* Guidance, the Board should also institute if only some challenges fall within the scope of Section 325(d) when the statutory purpose would not be undermined. *See SAS Q&As*, at D1 (June 5, 2018).⁶ Here, although Ground 4 is based on Mathieson and the combined Sutardja, Mathieson (which has never been previously considered, either alone or as a primary reference combined with Sutardja) teaches all of the limitations of independent claims 1, 8, and 20, except for "the same instruction set" feature, an indisputable teaching of Sutardja '324, '748, and '785 and which was never disputed during prosecution. Indeed, Mathieson as a primary reference combined with express teachings of Sutardja render all of the claims of Ground 4 unpatentable. *See Cradlepoint, Inc. v.*

⁶ Available at https://www.uspto.gov/sites/default/files/documents/sas_qas_20180605.pdf

3G Licensing S.A., IPR2021-01141, Paper 23 at 34-35 (PTAB Jan. 13, 2022) (“irrespective of whether § 325(d) is implicated for [one] Ground, it would be inappropriate to deny institution of the entire Petition under § 325(d)); *Berkeley Lights, Inc. v. Univ. of British Columbia*, IPR2021-01249, Paper 9 at 38 (PTAB Jan. 21, 2022) (declining to deny a petition where the primary reference in an alternative ground was never considered); *Grimco, Inc. v. Principal Lighting Group, LLC*, IPR2021-00968, Paper 13, at 20 (PTAB November 22, 2021) (declining to apply 325(d) to deny a petition where the primary prior art was new). Institution based on the Mathieson grounds thus would not undermine Section 325(d) purpose.

IX. LEVEL OF ORDINARY SKILL IN THE ART

A person of ordinary skill in the art at the relevant time (“POSITA”) would have had a bachelor’s degree in electrical engineering, computer science, computer engineering, material science, physics, applied physics, or a related field, and at least two years of experience in the research, design, development, or testing of electronic circuits or components or software for controlling electronic circuits or components, or the equivalent, with additional education substituting for experience and vice versa. Ex-1002, ¶25.

X. PRIORITY DATE

The earliest possible priority for the ’080 Patent is December 22, 2011. Petitioner takes no position on the proper priority date for each claim of the ’080

Patent. Assuming this earliest date, all prior art references asserted in this petition are prior art.

XI. CLAIM CONSTRUCTION

Petitioner interprets the claims of the '080 Patent according to 37 C.F.R. § 42.100(b). To resolve the particular grounds presented herein, Petitioner does not believe that any term requires explicit construction. Ex-1002, ¶¶77-78.⁷

XII. BRIEF DESCRIPTION OF THE APPLIED PRIOR ART REFERENCES

A. Mathieson (Ex-1005)

Mathieson was filed on May 25, 2010 and published September 1, 2011. Mathieson qualifies as prior art under at least 35 U.S.C. §§ 102(a) and (e).

Like the '080 Patent, Mathieson discloses a multi-core processor that includes a first set of cores and a second set of cores, which differ in their levels of power consumption because of differences in their constituent transistors. Ex-1005, ¶¶28-29. As a result, and as illustrated below in Figure 6, for a range of operating frequencies, Mathieson's first set of cores consume more power than its second set of cores, for the same operating frequency.

⁷ Claim construction procedures have not yet begun in the related district court or ITC matters. Petitioner reserves the right to revisit and revise these constructions in district court and the ITC as necessary. Additionally, Petitioner will request leave to submit any claim construction order as soon as it becomes available, so that it is timely made of record in the proceeding and can be considered by the Board. 37 C.F.R. § 42.100(b).




FIG. 6

Because Mathieson's cores consume different levels of power for the same operating parameters, Mathieson teaches that power can be conserved by switching between the sets of cores, based on the size of the workload. Ex-1005, ¶¶31-32. In particular, from a first mode of operation in which the first (and optionally, the second) set(s) of cores are enabled, Mathieson teaches switching to a second mode of operation in which only the second set of cores are enabled, for a decrease in workload below a "threshold value." Ex-1005, ¶¶30-32, 57, Fig. 4A.

B. Carmack (Ex-1006)

Carmack was filed on June 11, 2008 and published December 17, 2009. Carmack qualifies as prior art under at least 35 U.S.C. §§ 102(a) and (b). Carmack

belongs to the same patent family as Mathieson, all of Carmack's inventors are listed as inventors on Mathieson, and both references disclose multi-core processors that include cores with different power consumption properties. Ex-1005, Cover, ¶¶17-18; Ex-1006, Cover, ¶¶27-29; Ex-1002, ¶66. Indeed, Mathieson incorporates by reference Carmack's entire disclosure. Ex-1005, ¶1.

Like Mathieson, Carmack teaches dynamically selecting between running high-power cores, low-power cores, or both, depending on workload, as illustrated in Figure 3 of Carmack:




Figure 3

Carmack also discloses that, when both types of cores are used, the high-power cores may be operating at their maximum frequency. Ex-1006, ¶22.

C. Sutardja (Ex-1007, Ex-1008)


Sutardja '748 was filed June 30, 2008 and published November 20, 2008. Sutardja '748 qualifies as prior art under at least 35 U.S.C. § (b). Sutardja '748 incorporates by reference the disclosure of Sutardja '785 (Ex-1008) in its entirety. Ex-1007, ¶1. Both Sutardja '748 and Sutardja '785 (collectively “Sutardja”) disclose the same subject matter of a multi-core processor with both high-power and low-power cores. Ex-1007, ¶¶212-215; Ex-1008, ¶¶208, 215. And, like Mathieson and the '080 Patent, Sutardja teaches that the power consumption differences between cores are due to transistor differences. Ex-1007, ¶230; Ex-1008, ¶¶217, 225. In particular, Sutardja explains that “[g]ates of the transistors [of the HP core] may be larger than gates of the transistors [of the LP core].” Ex-1007, ¶230.

Also like Mathieson and the '080 Patent, Sutardja teaches operating the multi-core processor in a high-power mode (with the high-power cores and, optionally, the low-power cores enabled) and a low-power mode (with only the low-power cores enabled), based on system loading. Ex-1007, ¶¶223, 259.

D. Rychlik (Ex-1009)

Rychlik was filed on November 11, 2010 and published June 16, 2011. Rychlik qualifies as prior art under at least 35 U.S.C. §§ 102(a) and (e).

Rychlik discloses a power management technique for a multi-core processor in which cores are disabled, one-by-one, in response to falling levels of demand, as illustrated in Figure 12, below. Ex-1009, ¶¶68-71, 88, 96.



Rychlik also teaches the use of a “dynamic clock and voltage scaling (DCVS) algorithm” on each core to adjust the core’s frequency and/or voltage in response to changes in demand. Ex-1009, ¶¶34, 100.

XIII. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS

The '080 Patent contains 24 claims, of which 1, 9, and 17 are independent. Exhibit 1020 is a summary table showing correspondence between similar challenged claims and elements.

A. Ground 1: Claims 1-4, 7-12, 15-20, 23-24 Are Rendered Obvious By Sutardja (Ex-1007, Incorporating Ex-1008)

1. Independent Claim 1

a. Element 1[preamble]: A multi-core processor comprising:

Even if the preamble is limiting, Sutardja teaches a multi-core processor. Ex-1002, ¶¶79-80. Sutardja discloses “a multi-core processing system having a LP core and a HP core,” as illustrated below in Figure 3B. Ex-1007, ¶221; Ex-1002, ¶80.




FIG. 3B

b. Element 1[a][i]: a first plurality of cores and a second plurality of cores that support a same instruction set,

Sutardja teaches Element 1[a][i]. Ex-1002, ¶¶81-94.

Sutardja discloses “a high-speed, high-power (HP) core” and “a low-speed, low-power (LP) core,” both of which support the same instruction set: “The LP core 28 may use an [instruction set architecture or] ISA that is *the same as* or a subset of the ISA of the HP core 24.” Ex-1007, ¶225; *see also* Ex-1007, ¶213.

Sutardja also teaches a plurality of HP cores and a plurality of LP cores. Ex-1002, ¶¶83-94. Sutardja explains that “HP core 204 may comprise multiple identical cores,” and that “[w]hen applications demand still higher performance than that provided by one HP core, *multiple HP cores may be used.*” Ex-1007, ¶¶219, 248.

A POSITA would also have been motivated to include multiple LP cores. Ex-1002, ¶¶84-92. First, it was well-known to include multiple low-power cores together with multiple high-power cores in a multi-core processor. Ex-1002, ¶84; Ex-1023, ¶13; Ex-1024, 0003; Ex-1025, 0001. Second, Sutardja discloses that a set of “first applications ... that require relatively low processing speed during operation” may be executed on the LP core. Ex-1007, ¶226. However, because the processing capacity of the LP core is limited, “if a sufficient number of the first applications [] are being executed concurrently, the processing system [] may need to transition to and operate in the HP mode.” Ex-1007, ¶226. In such a scenario, one

or more of the first applications may be off-loaded onto the HP core, which may operate together with the LP core to meet the performance requirements of those concurrently executing first applications. Ex-1002, ¶¶85-87. However, because the HP core consumes more power for the same operating frequency and voltage as the LP core, spreading the execution of the first applications across an LP core and an HP core would consume more power than using a pair of LP cores. Ex-1002, ¶90. Thus, a POSITA would have recognized that routinely operating in the HP mode, while nevertheless executing applications that do not require the high speeds offered by the HP cores and which could be executed on an LP core, would frustrate Sutardja's purpose of "reduc[ing] power consumption." Ex-1007, ¶222; Ex-1002, ¶90; *see also* Ex-1026, 1:65-67 ("[U]sing a processor that ... exceeds the needs of the application software can lead to significant power waste and costs.").

Accordingly, a POSITA seeking to implement Sutardja's teachings in a system in which the concurrent execution of a large number of applications "that require relatively low processing speed during operation" is expected to routinely exceed the processing capacity of a single LP core would have been motivated to include multiple LP cores capable of handling such concurrent execution, without needing to rely on the additional processing capacity of the HP cores. Ex-1002, ¶91. This may be the case, for example, on a server, to which Sutardja teaches its disclosure may be applied. Ex-1007. ¶220. Server systems typically include both

multiple “simple” and multiple “high-performance” cores, because “[s]erver workloads can comprise multiple applications with varying computational and performance requirements.” Ex-1027, 0005; Ex-1002, ¶91.

Implementing multiple LP cores would achieve Sutardja’s purpose of power-efficient processing by using a core of the minimum power consumption suitable to achieve the processing requirements for any given executing task. Ex-1002, ¶92; Ex-1026, 2:17-19 (for power conservation purposes, when choosing from among cores offering different performance levels, “the target processor ... sought for [a] job is the one that consumes a minimum of power and still yields acceptable performance”). Moreover, a POSITA would have had a reasonable expectation of success in doing so, since Sutardja already envisions using multiple HP cores, and there is nothing that precludes using multiple LP cores for the same reasons and in the same manner as multiple HP cores. Ex-1002, ¶92.

Furthermore, incorporated Sutardja ’785 discloses a multi-processor system that includes multiple primary processors, in the form of a high-power CPU and a high-power GPU (the claimed first plurality of cores), and multiple secondary processors, in the form of a low-power CPU and a low-power GPU (the claimed second plurality of cores). Ex-1008, ¶¶215, 223. These multiple primary and secondary processors are illustrated below, in annotated Figure 2A. Ex-1002, ¶¶93-94.




FIG. 2A

Sutardja '785 explains at least one benefit of using multiple primary and secondary processors: it “increase[s] the speed of graphics processing and performance.” Ex-1008, ¶5. A POSITA seeking to implement Sutardja would have applied both the teachings of Ex-1007 and incorporated Ex-1008, such that both the HP and LP cores could, and indeed should, be implemented with a plurality of cores to achieve the attendant benefits taught by each reference. Ex-1008, ¶5; Ex-1002, ¶94.

- c. **Element 1[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and**

Sutardja teaches Element 1[a][ii]. Ex-1002, ¶¶95-96.

The second cores (Sutardja's LP cores) consume less power, for a same applied operating frequency and supply voltage, than the first cores (Sutardja's HP cores). Notably, the '080 Patent explains that the reason "[the] lower power core will consume less power than [the] higher power core," "for a same applied clock or operating frequency, [is] because of its narrower drive transistor widths." Ex-1001, 4:13-16. Just like the '080 Patent, the transistors of Sutardja's LP core are narrower than those of its HP cores: "The transistors 106 of the HP core 24 tend to consume more power during operation in the active state than the transistors 110 of the LP core 28," because "[g]ates of the transistors 106 may be larger than gates of the transistors 110." Ex-1007, ¶230. Indeed, Sutardja contemplates that its HP and LP cores may consume different levels of power: "the LP and HP cores ... have different performance and capabilities (e.g., different processing speeds, *power consumption*, etc.)." Ex-1007, ¶213. Accordingly, the LP cores, with its narrower transistors, consume less power.

- d. **Element 1[b][i]: power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores,**

Sutardja teaches Element 1[b][i]. Ex-1002, ¶¶97-111.

First, Sutardja teaches a state where the first plurality of cores (the HP cores) and the second plurality of cores (the LP cores) are both enabled. Ex-1002, ¶¶98-99. As illustrated in Figure 4 below, at step 306 (purple), Sutardja's processing system may be in the HP mode, corresponding to a state in which the HP cores (the claimed first plurality of cores) are enabled. Ex-1007, ¶¶221, 223, 245.




FIG. 4

Sutardja explains that “the LP core 28 may be in the active state during all or part of the HP mode,” so a POSITA would have understood that step 306 can correspond to a state in which both the HP and LP cores are enabled. Ex-1007, ¶223; Ex-1002,

¶99. A POSITA would have been motivated to enable both Sutardja's HP and LP cores to operate together when workload is large, such that the processing resources of Sutardja's HP cores, alone, are not sufficient to achieve the desired performance level. Ex-1002, ¶99. As another example, a POSITA would have been motivated to enable both Sutardja's HP and LP cores, by actively operating the HP cores and maintaining the LP cores in "standby" mode. Ex-1007, ¶224; Ex-1002, ¶99. As Sutardja explains, doing so would "allow quicker transitions to the active state" for the LP cores, when transitioning to the LP mode. Ex-1007, ¶224; Ex-1002, ¶99.

Second, Sutardja teaches a demand in the form of "system load" and/or "power/performance demand." Ex-1007, ¶218 ("[C]ores are dynamically enabled (i.e., activated) or disabled (i.e., deactivated) based on the system load. ***The system load may be partly determined based on resource utilization and performance demanded*** by the applications."); ¶286 (disclosing "switching cores according to the change in the ***power/performance demand***").

Third, Sutardja teaches disabling all of the first plurality of cores (Sutardja's HP cores) for a drop in demand (system load) below a threshold, without disabling any of the second plurality of cores (Sutardja's LP cores). Ex-1002, ¶¶101-106. As illustrated in annotated Figure 4 (reproduced below), step 308 (green) involves determining whether to switch from the HP mode to the LP mode. Ex-1002, ¶101.




FIG. 4

Sutardja explains that this decision is based on system loading (the claimed demand): “After executing instructions using the HP core 204, system loading may eventually decrease. When the system loading *decreases to a predetermined point*, the core switching module 212 may transition operation from the HP core 204 back to the LP core 200.” Ex-1007, ¶259. Thus, if system loading (the claimed demand) “decreases to a predetermined point” (the claimed threshold), Sutardja’s system transitions to the LP mode (illustrated in steps 310-318 (blue) of Figure 4, below). Ex-1007, ¶245; *see also* Ex-1007, ¶286 (“The core profile module 412 may generate an event when the *power/performance requirement changes to a value* greater or *less than a*

threshold value. ... The event is a request for switching cores according to the change in the power/performance demand.”).

As Sutardja explains, “[i]n the LP mode, the LP core 28 operates in the active state and processes threads and the HP core 24 is in the inactive state.” Ex-1007, ¶224. Sutardja further explains that in “the inactive state,” “[t]he inactive core may be completely turned off (no power supplied).” Ex-1007, ¶224. A core that is “completely turned off” is disabled. Ex-1002, ¶104.




FIG. 4

A POSITA would have also understood that the transition from the HP mode (with both the HP and LP cores enabled) to the LP mode (with only the LP cores

enabled) would occur without disabling any of the LP cores (the claimed second plurality of cores). First, Sutardja discloses that during the transition, the LP cores receive “threads” and other “information such as registers, checkpoints and/or program counters” transferred from the HP cores. Ex-1008, ¶237. A POSITA would have understood that the LP cores would need to be provided with sufficient power (i.e., be enabled) in order to receive these threads and information. Ex-1002, ¶106. Additionally, a POSITA would have recognized that disabling the LP cores momentarily, only to re-enable them, would detrimentally affect performance and power consumption, potentially leading to lost threads during the transfer. Ex-1002, ¶106. Indeed, Sutardja explains that leaving the LP cores enabled, in a “standby” mode enables “quicker transitions to the active state” for the LP cores, when transitioning to the LP mode. Ex-1007, ¶224; Ex-1002, ¶106.

Finally, Sutardja teaches that the transition from the HP mode to the LP mode is performed using hardware (the claimed power management hardware). Ex-1002, ¶¶107-111. Sutardja explains that a “core switching module” is used to transition the system between the HP and LP modes. Ex-1007, ¶¶258, 261, 263; Ex-1002, ¶107. The hardware controlled by the core switching module corresponds to the claimed power management hardware. Ex-1002, ¶108.

Sutardja explains that the core switching module “may be integrated with [a] hypervisor module 213,” as illustrated in Figure 5. Ex-1007, ¶253.




FIG. 5

Sutardja teaches that the hypervisor may “operate[] above the kernel” or it “may operate below the kernel.” Ex-1007, ¶254. Because the kernel is the lowest level of the operating system, a POSITA would have understood that a hypervisor that “operate[s] below the kernel” is installed on underlying hardware. Ex-1007, ¶¶249, 254; Ex-1002, ¶110; Ex-1028, 0003-0004 (a hypervisor may take one of two forms—a “hypervisor [that] runs directly on [] hardware,” or “a hypervisor that runs inside the operating system”). A POSITA would have understood that the hardware on which the disclosed hypervisor is installed corresponds to the claimed power management hardware. Ex-1002, ¶111.

- e. **Element 1[b][ii]: wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.**

Sutardja teaches Element 1[b][ii]. Ex-1002, ¶¶112-117.

Sutardja discloses an operating system executing on the multi-core processor. Ex-1002, ¶113. Sutardja explains that “[t]he LP and HP cores ... may execute *the same operating system (OS)*,” and that “the OS may include Linux® or other suitable operating system.” Ex-1007, ¶249. Sutardja further explains that “[a] kernel module 220 [] is a component of the OS [and] manages system resources.” Ex-1007, ¶249. Sutardja’s kernel module is illustrated in Figure 5:




FIG. 5

Sutardja also discloses that the operating system monitors a demand for the multi-core processor. Ex-1002, ¶115. As explained for Element 1[b][i], Sutardja teaches demand in the form of “system loading.” Section XIII.A.1.d; Ex-1007, ¶259. Sutardja discloses that a core profile module included in a “power management system (PMS) module” (illustrated in Figure 5 above) monitors this demand: “[a] core profile module 240 **monitors system loading.**” Ex-1007, ¶¶251, 260.

Sutardja’s “FIG. 5 is a functional block diagram.” Ex-1007, ¶195. A POSITA would have understood that functional block diagrams like Figure 5 are used to depict how various software modules function together. Ex-1002, ¶116. Thus, a POSITA would have understood that the PMS Module is executed by the kernel module (part of the Linux OS) to control the functioning of the multi-core processor. *Id.* Sutardja thus discloses an operating system that includes kernel module and PMS module 242. *Id.*

Sutardja further discloses that the operating system (kernel module and PMS module) controls the power management hardware (core switching module) based on the demand: “the PMS module 242, and/or the kernel module 220 ... generate[s] a system call to a ... core switching module 212 to **trigger a transition.**” Ex-1007, ¶253; *see also id.*, ¶¶262-263; Ex-1002, ¶117.

2. **Dependent Claim 2: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.**

Sutardja teaches Claim 2, as the Examiner found during prosecution and the Applicant did not dispute. Ex-1004, 0054-0057, 0067-0072; Ex-1002, ¶¶118-121. Sutardja's LP and HP cores are comprised of logic gates, which are comprised of transistors. Ex-1007, ¶230. Sutardja's "HP and LP cores 24 and 28 include transistors 106 and 110, respectively," as illustrated in Figure 3B. Ex-1007, ¶230; Ex-1002, ¶120.




FIG. 3B

Sutardja explains that “[g]ates of the transistors 106 may be larger than gates of the transistors 110.” Ex-1007, ¶230. Because the gates of transistors 106 of the HP core are larger than the gates of transistors 110 of the LP core, Sutardja teaches that the second plurality of cores (Sutardja’s LP cores) comprise logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores (Sutardja’s HP cores). Ex-1002, ¶121.

3. Dependent Claim 3: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that consume less power than corresponding logic gates of the first plurality of cores.

Sutardja teaches Claim 3, as the Examiner found during prosecution and the Applicant did not dispute. Ex-1004, 0054-0057, 0067-0072; Ex-1002, ¶¶122-124. As previously explained, the cores in Sutardja are comprised of logic gates comprised of transistors. Ex-1007, ¶230; Ex-1002, ¶124. Sutardja explains that “[t]he transistors 106 of the HP core 24 tend to consume more power during operation ... than the transistors 110 of the LP core 28.” Ex-1007, ¶230. Accordingly, a POSITA would have understood that the logic gates of the LP core (comprising transistors 110) consume less power than the logic gates of the HP core (comprising transistors 106). Ex-1002, ¶124.

- 4. Dependent Claim 4: The multi-core processor of claim 1, wherein the second plurality of cores each have a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.**

Sutardja teaches Claim 4. Ex-1002, ¶¶125-127. Sutardja provides an example in which “the low-speed processor may operate at speeds that are less than or equal to 500 MegaHertz (MHz) while the high-speed processor may operate at speeds that are greater than or equal to 1 GigaHertz (GHz).” Ex-1007, ¶210; *see also id.*, ¶234. Thus, the second plurality of cores (LP cores) each have a maximum operating frequency (e.g., 500 MHz) that is less than a maximum operating frequency (e.g., greater than 1 GHz) of the first plurality of cores (HP cores).

- 5. Dependent Claim 7: The multi-core processor of claim 1, wherein the first plurality of cores are at a maximum operating frequency in the state.**

Sutardja teaches Claim 7. Ex-1002, ¶¶128-131. Sutardja discloses “chang[ing] the clock frequency of Core2 [(the HP core)] within a frequency range from low to high ... according to the performance level demanded by the application.” Ex-1007, ¶304; *see also id.*, ¶301. A POSITA would have been motivated to set the high end of this frequency range to the HP cores’ maximum frequency, to enable the HP cores to provide their highest levels of performance, when needed (i.e., when demand is high). Ex-1002, ¶130.

A POSITA also would have recognized that the HP cores would be operating at their maximum operating frequency in the state in which both the HP and LP cores

are enabled. Ex-1002, ¶131. In particular, a POSITA would have understood that operating both the HP and LP cores at the same time would consume more power than operating the HP cores alone. *Id.* Accordingly, a POSITA would have recognized that, to continue to meet the performance levels required by rising workload (demand), Sutardja's multi-core system would first increase the operating frequencies of the HP cores to their maximum values, before enabling the LP cores, because doing so would conserve power. *Id.*

6. Dependent Claim 8: The multi-core processor of claim 1, wherein

- a. Element 8[a]: the power management hardware is to enable all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores,**

Sutardja teaches Element 8[a]. Ex-1002, ¶¶132-139.

First, Sutardja teaches a state where the second plurality of cores (LP cores) are enabled, and the first plurality of cores (HP cores) are disabled: "In the LP mode, the LP core 28 operates in the active state and processes threads and the HP core 24 is in the inactive state." Ex-1007, ¶224. As illustrated in Figure 4 below, in step 306 (blue), Sutardja's system may be operating in the LP mode, with only the LP cores enabled. Ex-1007, ¶¶224, 246; Ex-1002, ¶133.




FIG. 4

Second, as explained for Element 1[b][i], Sutardja teaches a demand in the form of “system loading” and/or “power/performance demand.” *See* Section XIII.A.1.d.

Third, Sutardja teaches enabling all of the first plurality of cores (HP cores) for an increase in demand (system load) above the threshold, without disabling any of the second plurality of cores (LP cores). Ex-1002, ¶¶135-138. As illustrated in Figure 4, above, step 328 (green) involves determining whether to switch from the LP mode to the HP mode. Ex-1002, ¶135.

Sutardja explains that this decision is based on system loading (demand): “When the LP core 200 reaches its highest operating speed and additional system loading occurs, ... the core switching module 212 may transition (switch) the operation from the LP core 200 to the HP core 204.” Ex-1007, ¶253; *see also* Ex-1007, ¶260. Thus, if system loading (the claimed demand) increases above the capacity of the LP core (i.e., above the claimed threshold), Sutardja’s system transitions to the HP mode (steps 330-338 (purple) of Figure 4, above). Ex-1007, ¶253; *see also* Ex-1007, ¶286 (“The core profile module 412 may generate an event when the *power/performance requirement changes to a value greater or less than a threshold value*.... The event is a request for switching cores according to the change in the power/performance demand.”); Ex-1002, ¶137.

Furthermore, as previously explained, Sutardja discloses that “the LP core [] may be in the active state during all or part of the HP mode.” Ex-1007, ¶223; Ex-1008, ¶208. A POSITA would have been motivated to enable both Sutardja’s HP and LP cores to operate together, for example, when the workload to be processed by Sutardja’s system is large, such that the processing resources of Sutardja’s HP cores, alone, are not sufficient to achieve the desired performance. Ex-1002, ¶138; Ex-1024, 0003-0004 (“[I]n times of heavy load, low-power cores can offer additional parallelism without significant increases in power consumption”). As another example, a POSITA would have been motivated to enable both Sutardja’s

HP and LP cores together, by actively operating the HP cores and maintaining the LP cores in “standby” mode. Ex-1007, ¶224; Ex-1002, ¶138. As Sutardja explains, doing so would “allow quicker transitions to the active state” for the LP cores, when transitioning back to the LP mode. Ex-1007, ¶224; Ex-1002, ¶138. Accordingly, a POSITA would have understood that the transition to the HP mode may occur without disabling the LP cores (second plurality of cores). Ex-1002, ¶138. In particular, a POSITA would have recognized that disabling the LP cores momentarily, only to re-enable them (either in active or standby mode), would detrimentally affect performance and power consumption, and potentially lead to lost threads and data during the transition. Ex-1002, ¶138.

Finally, as explained for Element 1[b][i], Sutardja also teaches that power management hardware is used to performed the above-described transition. Section XIII.A.1.d; Ex-1002, ¶139.

- b. Element 8[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.**

Sutardja teaches Element 8[b] for the same reasons as explained above for Element 1[b][ii]. *See* Section XIII.A.1.e; Ex-1002, ¶140.

7. Independent Claims 9 and 17:

a. Element 9[preamble]: A method comprising:

Even if the preamble is limiting, Sutardja teaches a method. *See, e.g.*, Ex-1007, Cls. 15-28; Ex-1002, ¶141.

b. Element 17[preamble]: A non-transitory machine readable medium containing program code that when processed by a machine causes a method to be performed, the method comprising:

Even if the preamble is limiting, Sutardja teaches a non-transitory machine readable medium containing program code. Ex-1002, ¶142. For example, Sutardja discloses “[a] computer program stored on a computer readable medium and executed by a processor....” Ex-1007, Cl. 43.

c. Elements 9[a][i] and 17[a][i]: operating a multi-core processor such that a first plurality of cores and a second plurality of cores execute a same instruction set,

Sutardja teaches Elements 9[a][i] and 17[a][i]. *See* Section XIII.A.1.b; Ex-1002, ¶143.

d. Elements 9[a][ii] and 17[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and

Sutardja teaches Elements 9[a][ii] and 17[a][ii]. *See* Section XIII.A.1.c; Ex-1002, ¶144.

- e. **Elements 9[b][i] and 17[b][i]: disabling with power management hardware, from a state where the first plurality of cores and the second plurality of cores are enabled, all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores,**

Sutardja teaches Elements 9[b][i] and 17[b][i]. *See* Section XIII.A.1.d; Ex-1002, ¶145.

- f. **Element 9[b][ii] and 17[b][ii]: wherein an operating system executing on the multi-core processor monitors a demand for the multi-core processor and controls the power management hardware based on the demand.**

Sutardja teaches Elements 9[b][ii] and 17[b][ii]. *See* Section XIII.A.1.e; Ex-1002, ¶146.

- 8. **Dependent Claims 10 and 18: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.**

Sutardja teaches Claims 10 and 18. *See* Section XIII.A.2; Ex-1002, ¶147.

- 9. **Dependent Claims 11 and 19: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that consume less power than corresponding logic gates of the first plurality of cores.**

Sutardja teaches Claims 11 and 19. *See* Section XIII.A.3; Ex-1002, ¶148.

- 10. Dependent Claims 12 and 20: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the second plurality of cores at a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.**

Sutardja teaches Claims 12 and 20. *See* Section XIII.A.4; Ex-1002, ¶149.

- 11. Dependent Claims 15 and 23: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the first plurality of cores at a maximum operating frequency in the state.**

Sutardja teaches Claims 15 and 23. *See* Section XIII.A.5; Ex-1002, ¶150.

- 12. Dependent Claims 16 and 24: The [method of claim 9/non-transitory machine readable medium of claim 17], further comprising**

- a. Elements 16[a] and 24[a]: enabling, with the power management hardware, all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores,**

Sutardja teaches Elements 16[a] and 24[a]. *See* Section XIII.A.6.a; Ex-1002,

¶151.

- b. Elements 16[b] and 24[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.**

Sutardja teaches Element 16[b]. *See* Section XIII.A.6.b; Ex-1002, ¶¶152.

B. Ground 2: Claims 5-6, 13-14, and 21-22 Are Rendered Obvious By Sutardja/Rychlik

1. Dependent Claims 5, 13, and 21:


- a. **Elements 5[a], 13[a], 21[a]: The [multi-core processor of claim 1/method of claim 9/non-transitory machine readable medium of claim 17], [wherein the power management hardware is to disable/further comprising disabling, with the power management hardware,] an additional core of the second plurality of cores for each continued drop in demand below a next lower threshold until one core of the second plurality of cores remains enabled, and**

Sutardja/Rychlik teaches Elements 5[a], 13[a], and 21[a]. Ex-1002, ¶¶153-161.

Rychlik discloses “multicore CPU 402 [that] may include a zeroth core 410, a first core 412, and an Nth core 414.” Ex-1009, ¶34.

Rychlik teaches demand for the cores in the form of a “workload”: “[A] parallelism monitor 440 may monitor the workload on the cores 410, 412, 414.” Ex-1009, ¶37. The workload (demand) may take the form of “a total system load,” or “a running average of the degree of parallelism in the workload.” Ex-1009, ¶38.

Rychlik also discloses disabling an additional core of the plurality of cores for each continued drop in demand below a next lower threshold. *See, e.g.*, Ex-1009, ¶96; Ex-1002, ¶156. This process is illustrated in Figures 10 and 12 of Rychlik, annotated below. Note, for ease of explanation, the annotations and discussion to follow assumes Rychlik’s multi-core CPU includes 3 cores. Ex-1002, ¶156.



At step 1202 (green) Rychlik's zeroth core, first core, and Nth core are enabled. Ex-1009, ¶68. If, at step 1208 (dark blue), the degree of workload parallelism (demand) satisfies "[an] Nth sleep condition," corresponding to "the parallelism [being] less than or equal to [a] *threshold value*" associated with the Nth sleep condition, then in step 1216 (purple) one of Rychlik's cores is disabled. Ex-1009, ¶¶68-69. Assuming the process started with 3 enabled cores, this leaves Rychlik's zeroth core and first core enabled, as illustrated in step 1002 (red). Ex-1009, ¶71. If, at step 1008 (yellow), the degree of workload parallelism (the claimed demand) satisfies "a first sleep condition," corresponding to "the parallelism [being] less than or equal to [a]

threshold value” associated with the first sleep condition (the claimed next lower threshold), in step 1016 (pink) one of Rychlik’s two remaining cores is disabled, leaving only one core enabled. Ex-1009, ¶¶61-64.

Rychlik similarly teaches disabling an additional core each time system load (demand) falls below a next lower threshold load value: “[O]ne or more cores may be turned on or off based on the load value.... [A]s the load value decreases below **another threshold** indicating that having a third core operational is wasting power, the third core may be powered off and the workload may be spread across the two remaining cores.” Ex-1009, ¶88; Ex-1002, ¶160. A POSITA would have understood that, just like the process illustrated in Figures 10 and 12, this process would continue, as load continued to drop, until only one of Rychlik’s cores remained enabled. Ex-1002, ¶160.

A POSITA would have been motivated to apply Rychlik’s teachings to Sutardja’s LP cores, to disable an additional core of Sutardja’s LP cores (the claimed second plurality of cores), for each continued drop in workload (the claimed demand) below a next lower threshold, until only one LP core remained enabled. Ex-1002, ¶161. Both Rychlik and Sutardja relate to power management in multi-core processors. Ex-1007, ¶212; Ex-1009, ¶88. A POSITA implementing Sutardja’s teachings with multiple LP cores would understand that disabling one or more of those LP cores when demand is low would save power. Ex-1002, ¶161; *see also* Ex-

1031, Fig. 4. Accordingly, a POSITA would have been motivated to search for references, such as Rychlik, that disclose implementation details of processor scaling methods. Ex-1002, ¶161. A POSITA would have been motivated to apply Rychlik’s teachings to Sutardja’s LP cores so as to avoid “wasting power.” Ex-1009, ¶88; Ex-1002, ¶161. A POSITA would have had a reasonable expectation of success in implementing Rychlik’s power management technique in Sutardja’s multi-core system, because it amounts to using a known technique (disabling cores based on demand, as taught by Rychlik) to improve a similar system (Sutardja’s multi-core system with multiple LP cores) in the same way (by reducing power consumption during periods of low demand). Ex-1002, ¶161.

b. Elements 5[b], 13[b], 21[b]: [lower/lowering] an operating frequency or a supply voltage of the one core of the second plurality of cores as demand drops below a next lower threshold.

Sutardja/Rychlik teaches Elements 5[b], 13[b], and 21[b]. Ex-1002, ¶¶162-167.

Rychlik and Sutardja both teach lowering an operating frequency or a supply voltage of the one core of the second plurality of cores as demand drops. Ex-1002, ¶¶163-166. Rychlik discloses including a “dynamic clock and voltage scaling (DCVS) algorithm” on each core. Ex-1009, ¶34. Rychlik explains that, when only one core is enabled (one core of the second plurality of cores), the frequency of the

core may be raised to accommodate rising levels of demand: “the system may respond to instantaneous transient loads by ramping up the frequency of the [single] CPU core.” Ex-1009, ¶100. A POSITA would have understood that, similarly, Rychlik’s system would lower the frequency of the core in response to lower levels of demand (i.e., demand falling below a next lower threshold). Ex-1002, ¶164.

Sutardja also teaches the use of voltage and frequency scaling on its LP core (one core of the second plurality of cores). Ex-1007, ¶¶301-302; Ex-1002, ¶¶165-166. Sutardja explains that for low workloads (demand), the frequency and/or voltage of the LP core may be reduced to save power: “The [system] may monitor CPU utilization and determine a lowest possible clock frequency and/or supply voltage for operating the ... LP cores ... [to] minimize impact on performance based on workload while saving power.” Ex-1007, ¶237.

As explained above for Elements 5[a], 13[a], and 21[a], a POSITA would have been motivated to implement Rychlik’s power management technique in Sutardja’s multi-core processor. Section XIII.B.1.a; Ex-1002, ¶167. A POSITA would also have been motivated to apply the voltage/frequency teachings of Rychlik and Sutardja to the combined Sutardja-Rychlik system, in order to conserve power when only one of Sutardja’s LP cores remains enabled, and demand continues to fall (i.e., demand drops below the claimed next lower threshold). Ex-1002, ¶167.

2. **Dependent Claims 6, 14, 22: [The multi-core processor of claim 5/method of claim 13/The non-transitory machine readable medium of claim 21], [wherein the power management hardware is to raise/further comprising raising, with the power management hardware,] a supply voltage or an operating frequency of said one core in response to higher demand.**

Sutardja/Rychlik teaches these claims. Ex-1002, ¶¶168-169. Rychlik explains that, when only one core is enabled (one core of the second plurality of cores), the frequency of the core may be raised to accommodate rising levels of demand: “[T]he system may respond to instantaneous transient loads by ramping up the frequency of the [single] CPU core.” Ex-1009, ¶100. As explained for Claims 5, 13, and 21, above, a POSITA would have been motivated to implement Rychlik’s power management techniques in Sutardja’s multi-core processor. *See* Section XIII.B.1; Ex-1002, ¶169.

C. Ground 3: Claims 7, 15, and 23 Are Rendered Obvious By Sutardja/Carmack

1. **Dependent Claims 7, 15, 23: The [multi-core processor of claim 1/method of claim 9/non-transitory machine readable medium of claim 17], wherein [the operating comprises operating] the first plurality of cores [are] at a maximum operating frequency in the state.**

Sutardja/Carmack teaches these claims. Ex-1002, ¶¶170-173. Carmack teaches that when first and second cores 110, 120, with differing power consumption properties, are operating together, the second cores 120 (first plurality of cores) may be operating at their peak (i.e., maximum) frequency: “For situations where the

overhead of coordinating asymmetric cores becomes too high, the system can lower the peak frequency of the faster core 120 to match the maximum frequency of the slower core 110....” Ex-1006, ¶22. Conversely, for situations where the overhead of coordinating is *not* too high, a POSITA would have understood that Carmack’s faster core 120 would be operating at its maximum frequency. Ex-1002, ¶171.

A POSITA would have been motivated to combine the teachings of Sutardja and Carmack, and would have had a reasonable expectation of success in doing so because each relates to power reduction techniques for multi-core processors that focus on the inclusion of cores with different power consumption properties within the same processing system. Ex-1006, ¶¶17-18; Ex-1007, ¶212; Ex-1002, ¶172.

Sutardja discloses that the operating frequency of its HP cores may be varied from a low value to a high value, “according to the performance level demanded.” Ex-1007, ¶304. A POSITA would have been motivated to set the high value to the maximum frequency of Sutardja’s HP cores, as taught by Carmack, to maintain desired performance levels during periods of high demand. Ex-1002, ¶173. A POSITA would have had a reasonable expectation of success in doing so, because it amounts to a simple substitution of one known element (Sutardja’s “high” operating frequency value) for another (Carmack’s maximum operating frequency) with predictable results (Sutardja’s HP cores operating at their maximum frequency). Ex-1002, ¶173.

D. Ground 4: Claims 1-4, 7-12, 15-20, and 23-24 Are Rendered Obvious By Mathieson/Sutardja


1. A POSITA would have been motivated to combine Mathieson with Sutardja

In addition to the reasons provided below for specific claim elements, a POSITA would have been motivated to combine the teachings of Mathieson and Sutardja, and would have had a reasonable expectation of success in doing so because each relates to the same well-known subject matter of power management in multi-core processors, both teaching techniques for reducing power consumption in response to decreased demand/workload. Ex-1005, Abstract, ¶¶31-32; Ex-1007, ¶259; Ex-1002, ¶175. And the techniques of both references rely on the inclusion of cores with different power consumption properties within the same processing system. Ex-1005, ¶20; Ex-1007, ¶212. Furthermore, both references contemplate that power consumption differences between cores may be achieved through differences in constituent transistors. Ex-1005, ¶28; Ex-1007, ¶230. At least because of these similarities, a POSITA seeking to implement Mathieson's teachings would have looked to Sutardja for additional implementation details. Ex-1002, ¶175.

2. Independent Claim 1

a. Element 1[preamble]: A multi-core processor comprising:

Even if the preamble is limiting, Mathieson discloses a multi-core processor in the form of “a central processing unit (CPU) 102,” illustrated below in Figure 1. Ex-1002, ¶¶176-177.



Mathieson’s CPU 102 is a multi-core processor because it includes multiple cores: “The CPU 102 includes one or more ‘fast’ cores 130 and one or more ‘shadow’ or slow cores 140.” Ex-1005, ¶20.

b. Element 1[a][i]: a first plurality of cores and a second plurality of cores that support same instruction set,

Mathieson and Sutardja teach Element 1[a][i]. Ex-1002, ¶¶178-187.

Mathieson discloses a first plurality of cores in the form of “a first set of cores 210,” and a second plurality of cores in the form of “a second set of cores 220,” as illustrated in Figure 2, below. Ex-1005, ¶25.




FIG. 2

Mathieson discloses that each set of cores 210 and 220 may include any number of cores: “[T]he first set of cores 210 includes one or more cores 212 ... and the second set of cores 220 includes one or more cores 222.” Ex-1005, ¶25; *see also* Ex-1005, ¶28.

Mathieson/Sutardja teaches that the first plurality of cores and the second plurality of cores support a same instruction set. Ex-1002, ¶¶181-187.

First, Mathieson explains that “the cores included in the first set of cores 210 may implement substantially *the same functionality* as the cores included in the second set of cores 220.” Ex-1005, ¶26. Because an instruction set architecture (ISA) specifies what a processor is capable of doing and how it gets done, a POSITA would have understood that a straightforward way to enable both of Mathieson’s sets of cores 210 and 220 to implement the same functionality, as taught be Mathieson, would be to implement the same ISA in each core. Ex-1002, ¶182. Cores operating the same instruction set will be able to perform the same basic functionality. *Id.*

Second Mathieson explains that “[w]hen execution of [] processing operations switches from the first set of cores to the second set of cores, ... the processor state [is transferred] from the first set of cores to the second set of cores.” Ex-1005, ¶43 (“[T]he processing state may be *directly transferred* to the other set of cores via a dedicated bus.”). A POSITA would have understood that a straightforward way ensure the ability to directly transfer the processing state is for both sets of cores to support the same ISA. Ex-1002, ¶184 (explaining that a POSITA would have understood that additional state translations would be required to transfer between different instructions sets, and Mathieson does not require such additional translations).

Based on these teachings, a POSITA seeking to produce Mathieson's multi-core processor would have been motivated to search for implementation details in similar references that teach cores with different power consumption properties that implement the same functionality. Ex-1002, ¶185. Sutardja is one such reference, which provides express teachings to use the "same instructions set" for different cores. Ex-1002, ¶185. *See* Section XIII.A.1.b.

As explained for Ground 1, the differences in speed and power consumption between Sutardja's LP cores and HP cores may be due to differences in transistor gate size between the cores. Sections XIII.A.1.b, XIII.A.2-3; Ex-1007, ¶¶212, 230. "[T]o achieve dynamic compatibility" between the two types of cores, "the instruction set of the LP core may be *the same as* or a subset of the instruction set of the HP core." Ex-1007, ¶270. A POSITA would have understood that "dynamic compatibility" between Sutardja's LP and HP cores corresponds to the ability to dynamically switch between the two types of cores when executing operations. Ex-1002, ¶186. This would ensure that both types cores support execution of those operations, and therefore, the functionality associated with those operations. *Id.*

Accordingly, to achieve Mathieson's teaching that both sets of cores implement the same functionality, a POSITA would have been motivated to design both types of cores to support the same instruction set, as taught by Sutardja. Ex-1002, ¶187.

- c. **Element 1[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and**

Mathieson teaches Element 1[a][ii]. Ex-1002, ¶¶188-193.

As illustrated in Figure 6 of Mathieson, for a large frequency range (up to 610), the power consumed by Mathieson's second set of cores (blue) is less than the power consumed by Mathieson's first set of cores (red) for *the same operating frequency* (e.g., " F_0 "). Ex-1002, ¶189.




FIG. 6

Indeed, Mathieson explains that for “particular operating characteristics,” including “operating frequency [and] voltage,” either an operating mode in which the first set

of cores is enabled or an operating mode in which the second set of cores is enabled may be “more optimal,” with “the CPU 102 [] us[ing] less power to operate in [that] mode.” Ex-1005, ¶¶30-31, 37.

Just like the '080 Patent, Mathieson explains that the difference in power consumption between cores 210 and 220, for the same operating characteristics (e.g., frequency and voltage), is due to transistor differences between the cores: “[T]he cores 212 [included in the first set of cores 210] may be manufactured using *faster transistors* that have *significant static leakage*.... [A]t slower clock speeds, the static leakage of the fast transistors can dominate the overall power consumption.” Ex-1005, ¶28; Ex-1001, 4:13-16 (“[B]ecause of its narrower drive transistor widths, for example, a lower power core will consume less power than a higher power core” “for a same applied clock or operating frequency”); Ex-1002, ¶¶191-193.

- d. **Element 1[b][i]: power management hardware to, from a state where the first plurality of cores and the second plurality of cores are enabled, disable all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores,**

Mathieson teaches Element 1[b][i]. Ex-1002, ¶¶194-205.

First, Mathieson teaches a state where the first plurality of cores (cores 210) and the second plurality of cores (cores 220) are both enabled. Ex-1002, ¶¶195-196. As illustrated in Figure 4A below, at step 402 (purple), CPU 102 is in the first mode

of operation, corresponding to a state in which the first set of cores 210 are enabled.

Ex-1005, ¶55.




FIG. 4A

Mathieson teaches that “the CPU 102 may operate in both the first mode and the second mode [(in which the second set of cores 220 is enabled)] simultaneously,” so a POSITA would have understood that step 402 can correspond to a state in which both sets of cores 210 and 220 are enabled. Ex-1005, ¶¶30-31; Ex-1002, ¶196. Indeed, a POSITA would have been motivated to enable both sets of cores to operate together at step 402, for example, when the workload to be processed by CPU 102

is large, such that the processing resources of the first set of cores 210, alone, are not sufficient to achieve the desired performance level. Ex-1002, ¶196.

Second, Mathieson teaches demand in the form of “workload.” Ex-1005, ¶31 (in the first mode of operation, the workload is evaluated to determine whether “*the workload* has [] decreased to a point where the CPU 102 would use less power to operate in the second mode”).

Third, Mathieson teaches disabling all of the first plurality of cores (cores 210) for a drop in demand (Mathieson’s workload) below a threshold without disabling any of the second plurality of cores (cores 220). Ex-1002, ¶¶198-200. In Mathieson’s Figure 4A (reproduced below), step 406 (green) involves “*evaluating the workload* [by] determining whether a processing parameter associated with processing the one or more operations is greater than or less than *a threshold value*.” Ex-1005, ¶¶32, 56, Fig. 4A.




FIG. 4A

If the workload (demand), as determined by the processing parameter, is less than the threshold value, CPU 102 switches to the second mode of operation, corresponding to a state in which the cores 210 (the claimed first plurality of cores) are disabled and the cores 220 (the claimed second plurality of cores) are enabled (step 408 (blue) above). Ex-1005, ¶¶30, 57.

A POSITA would have understood that the second set of cores 220 would remain enabled throughout the switch from the first mode of operation (with both the first and second sets of cores 210, 220 enabled) to the second mode of operation (with only the second set of cores 220 enabled), because disabling cores 220 only to

re-enable them would detrimentally affect performance and power consumption, and potentially lead to lost data during the transition. Ex-1002, ¶200.

Carmack, incorporated by reference into Mathieson, similarly explains that, from a state in which the first plurality of cores (Carmack's second cores 120) and the second plurality of cores (Carmack's first cores 110) are enabled, all of the first plurality of cores (Carmack's second cores 120) are disabled for a drop in demand (workload), without disabling any of the second plurality of cores (first cores 110). Ex-1002, ¶201. As illustrated in Figure 3 of Carmack (reproduced below), Carmack teaches using only the second plurality of cores (first cores 110) when workload is within a first range, using only the first plurality of cores (second cores 120) when workload is within a second range, and using both the first and second pluralities of cores when workload is within a third range, with the workload being determined "periodically." Ex-1006, ¶¶20-21, 23, Fig. 3.




Figure 3

Thus, if workload (the claimed demand) is within the third range at a first measurement time, and then at the next measurement time the workload has fallen below a threshold value such that it lands within the first range, Carmack's system will transition from a state in which both the first and second cores 110 and 120 are enabled (step 340, purple) to a state in which only the first cores 110 (the claimed second plurality of cores) are enabled (step 320, blue). Ex-1002, ¶204. And, as explained above, a POSITA would have been motivated to keep the first cores 110 (the claimed second plurality of cores) enabled throughout the transition, based on power and performance considerations. *Id.*; Section XIII.A.1.d.

Finally, Mathieson also explains that disabling/enabling cores is done using hardware (the claimed power management hardware): “In one embodiment, the controller saves the processor state to the shared resource 230, *triggers a hardware mechanism that stops and powers off the first set of cores 210*, and boots the second set of cores 220.” Ex-1005, ¶43; Ex-1002, ¶205.

- e. **Element 1[b][iii]: wherein an operating system to execute on the multi-core processor is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.**

Mathieson teaches Element 1[b][iii]. Ex-1002, ¶¶206-211. As explained for Element 1[b][i], Mathieson teaches a demand for the multi-core processor in the form of a “workload.” Section XIII.D.2.d; Ex-1005, ¶¶30-33. Mathieson discloses that a controller within CPU 102 (Figure 2 below) monitors the demand: “[A] controller 240 included within the CPU 102 is configured to *evaluate at least a workload* associated with one or more operations to be executed by the CPU 102.” Ex-1005, ¶30.




FIG. 2

As Mathieson explains, “the controller is implemented in software and is executed by the CPU 102.” Ex-1005, ¶30. Mathieson’s controller 240 both manages the “processing operations” executing on CPU 102, and selects the underlying hardware (e.g., cores 210 and/or 220) on which those operations are executed. Ex-1005, ¶¶35, 62. Because the functions of an operating system include process and device management, a POSITA would recognize that the controller performs the functions of an operating system, and thus would view controller 240 as being part of an operating system executing on CPU 102. Ex-1002, ¶209.

Mathieson also discloses that controller 240 (the claimed operating system) controls the power management hardware based on the demand. Ex-1002, ¶¶210-211. Mathieson explains that the controller causes CPU 102 to switch between

operating modes based on workload (the claimed demand): “If the controller 240 [] detects that the workload has [] decreased to a point where the CPU 102 would use less power to operate in the second mode, then ***the controller 240 causes*** the CPU 102 to operate in the second mode.” Ex-1005, ¶31; *see also* Ex-1005, ¶30. Mathieson also explains that the controller 124 causes the CPU 102 to change operating modes by controlling power management hardware: “When execution of the processing operations switches from the first set of cores to the second set of cores, ... ***the controller ... triggers a hardware mechanism*** that stops and powers off the first set of cores 210, and boots the second set of cores 220.” Ex-1005, ¶43.

3. Dependent Claim 2: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.

Mathieson/Sutardja teaches Claim 2. Ex-1002, ¶¶212-220.

Mathieson’s first and second cores in CPU 120 are comprised of logic gates, which are comprised of transistors. Ex-1005, ¶¶27 (“[T]he power consumption associated with the CPU 102 is derived from ‘dynamic’ switching power and ‘static’ leakage power. The switching power loss is based on the charging and discharging of the each transistor and its associated capacitance, and increases with operating frequency and number of gates.”), 36. Moreover, this is fundamental—a POSITA would have understood that different kinds of logic gates, e.g., NOR, NAND, are

the functional building blocks of a processing core, and are formed using transistors. Ex-1002, ¶213.

Mathieson discloses that the transistors of the first set of cores 210 (the claimed first plurality of cores) differ from those of the second set of cores 220 (the claimed second plurality of cores): “[T]he cores 212 included in the first set of cores 210 comprise ‘fast’ cores ... [that] may be manufactured using *faster transistors that have significant static leakage* [as compared to the cores of the second set of cores 220].” Ex-1005, ¶28.

Mathieson teaches that the second plurality of cores (Mathieson’s second set of cores 220) comprise logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores (Mathieson’s first set of cores 210). Ex-1002, ¶215. A POSITA would have understood that a small number of design choices affect transistor speed and/or static leakage. Ex-1002, ¶215. Among these is the width of the transistor. Ex-1002, ¶215; *see also, e.g.*, Ex-1029, 0015 (“Leakage grows with the aspect ratio [(width/length)] of a transistor.”); Ex-1030, 0007 (“reducing the width [of a transistor] increases the transistor’s delay”). A POSITA would have understood that increasing the width of a transistor would increase both its speed and static leakage of the transistor. *Id.* Thus, a POSITA would have been motivated to make predictable use of transistors with greater widths for Mathieson’s first set of cores 210 (first plurality of cores), as compared with

Mathieson's second set of cores 220 (second plurality of cores), to achieve the properties disclosed by Mathieson—namely, that the transistors in the first set of cores 210 are faster than the transistors in the second set of cores 220, particularly because transistor width is one of a finite number of identified factors that affect transistor speed and static leakage. Ex-1002, ¶¶215-216.

To the extent that a POSITA would have sought additional implementation details about transistor widths, this features is also taught by Sutardja. As explained for Ground 1, Sutardja teaches Claim 2. *See* Section XIII.A.2.

A POSITA seeking to follow Mathieson's teaching that a first set of cores 210 (the claimed first plurality of cores) include "faster transistors that have significant static leakage," as compared to a second set of cores 220 (the claimed second plurality of cores), would have been motivated to search for additional references that disclose transistor design details that would achieve such differences. Ex-1002, ¶218. Sutardja is one such reference. Ex-1002, ¶219. Sutardja discloses transistors with the same properties as those of Mathieson's transistors: "The HP core 24 may be fabricated using a *high leakage/high speed* process. The LP core 28 may be fabricated using a *low leakage/low speed* process." Ex-1007, ¶230 ("The transistors 106 [of the HP core] may have higher leakage current than the transistors 110."); Ex-1002, ¶219. Sutardja further discloses that these differences may result from

differences in transistor widths: “Gates of the transistors 106 may be larger than gates of the transistors 110.” Ex-1007, ¶230; Section XIII.A.2.

Accordingly, a POSITA would have been motivated to implement the transistors in Mathieson’s second set of cores (second plurality of cores) with narrower gates than the transistors in Mathieson’s first set of cores (first plurality of cores), as taught by Sutardja, to achieve the properties that Mathieson teaches are exhibited by those transistors. Ex-1002, ¶220. Moreover, as explained in Section XII.D.2.b, this feature would have been implemented by applying the combined Mathieson/Sutardja teachings to use the same instructions set and logic design in the two different sets of cores. A POSITA would have had a reasonable expectation of success in doing so, because Sutardja’s transistors exhibit the same speed and leakage properties as Mathieson’s transistors. Ex-1002, ¶220.

4. Dependent Claim 3: The multi-core processor of claim 1, wherein the second plurality of cores comprise logic gates that consume less power than corresponding logic gates of the first plurality of cores.

Mathieson and Sutardja teach Claim 3. Ex-1002, ¶¶221-225.

As explained for Claim 2, the cores in Mathieson are comprised of logic gates, which are comprised of transistors. Ex-1005, ¶¶27, 36. Moreover, a POSITA would have understood that logic gates in a processing core are formed from transistors. Ex-1002, ¶222.

Mathieson teaches Claim 3. Ex-1002, ¶¶223-224. Mathieson explains that power consumed by CPU 102 is consumed by its constituent transistors: “the power consumption associated with the CPU 102 ... is based on the charging and discharging of the each transistor and its associated capacitance ... [and] leakage in each transistor.” Ex-1005, ¶27. Mathieson further explains that “at slower clock speeds, the static leakage of the fast transistors [of the first set of cores] can dominate the overall power consumption.” Ex-1005, ¶28; *see also id.*, ¶73 (“[T]he lower total power associated with operating the processing complex at lower frequencies using the slow cores is based on the lower leakage power associated with the slow cores.”). Thus, a POSITA would have understood that at lower operating frequencies, the transistors (and the logic gates formed therefrom) of the second set of cores 220 (second plurality of cores) consume less power than the transistors/logic gates of the first set of cores 210 (first plurality of cores). Ex-1002, ¶224.

To the extent that a POSITA would have sought additional implementation details about logic gate power, this features is taught by Sutardja As explained for Ground 1, Sutardja teaches Claim 3. *See* Section XIII.A.3. And, as explained for Claim 2, a POSITA would have been motivated to use Sutardja’s transistors in Mathieson’s cores. *See* Section XIII.D.2.

5. **Dependent Claim 4: The multi-core processor of claim 1, wherein the second plurality of cores each have a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.**

Mathieson teaches Claim 4. Ex-1002, ¶¶226-229. As illustrated in Figure 6 (reproduced with annotations below), Mathieson's second set of cores (the claimed second plurality of cores) have a maximum frequency (labelled F_2) that is lower than the maximum frequency (labelled F_1) of Mathieson's first set of cores (the claimed first plurality of cores). Ex-1005, Fig. 6; Ex-1002, ¶227.




FIG. 6

Although Mathieson is silent on the maximum frequency of the first or second set of cores, it would have been readily apparent to a POSITA that the power usage of the second set of cores 608 becomes exponential after the frequency at line 610

and that, therefore, its maximum frequency is necessarily less than that of the first set of cores for which power usage is only gradually increasing in that range. Ex-1002, ¶228.

Carmack, incorporated by reference into Mathieson, similarly discloses that the maximum frequency of the second plurality of cores (Carmack's first cores 110) is less than the maximum frequency of the first plurality of cores (Carmack's second cores 120): "[T]he system can lower the peak frequency of the faster core 120 to match the maximum frequency of the slower core 110." Ex-1006, ¶22; Ex-1002, ¶229.

6. Dependent Claim 7: The multi-core processor of claim 1, wherein the first plurality of cores are at a maximum operating frequency in the state.

Mathieson teaches Claim 7. Ex-1002, ¶¶230-.

First, Mathieson teaches that "operating in both the first and second modes simultaneously may result in lower overall power efficiency." Ex-1005, ¶31. Mathieson also teaches that the operating frequency of the first set of cores (first plurality of cores) may be increased to accommodate larger workloads. Ex-1005, ¶31. Based on these two teachings, a POSITA would have recognized that, when only Mathieson's first set of cores are enabled, it would be more power-efficient to accommodate increasing workloads by increasing the operating frequency of the first set of cores, rather than by enabling Mathieson's second set of cores and

operating both sets of cores together. Ex-1002, ¶231. Thus, a POSITA seeking to implement Mathieson's CPU 102 in a power-efficient manner, would have designed the CPU such that, from a state in which only the first set of cores are enabled, when faced with increasing workloads, the operating frequency of the first set of cores would first be increased to its maximum value, and then, only after and the workload had continued to increase beyond the capacity of the first set of cores operating at their maximum frequency, the second set of cores would be enabled for operation together with the first set of cores (the claimed state). Ex-1002, ¶232.

Indeed, Carmack, incorporated by reference into Mathieson, Ex-1005, ¶1, discloses that when both the first and second cores 110, 120 are operating together, the second cores 120 (first plurality of cores) may be operating at their peak (i.e., maximum) frequency: "For situations where the overhead of coordinating asymmetric cores becomes too high, the system can lower the peak frequency of the faster core 120 to match the maximum frequency of the slower core 110...." Ex-1006, ¶22. Conversely, for situations where the overhead of coordinating is *not* too high, a POSITA would have understood that Carmack's faster core 120 would be operating at its maximum frequency. Ex-1002, ¶233.

7. **Dependent Claim 8: The multi-core processor of claim 1, wherein**
- a. **Element 8[a]: the power management hardware is to enable all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores,**

Mathieson teaches Element 8[a]. Ex-1002, ¶¶234-242.

First, Mathieson discloses a state where the second plurality of cores (Mathieson's cores 220) are enabled, and the first plurality of cores (Mathieson's cores 210) are disabled. *See* Section XIII.D.2.d. As illustrated in Figure 4A below, in step 408 (blue), CPU 102 is in the second mode of operation, corresponding to a state in which the cores 220 are enabled and cores 210 are disabled. Ex-1005, ¶30, Fig. 4A; Ex-1002, ¶235.




FIG. 4A

Second, as explained for Element 1[b][i], Mathieson teaches (1) the claimed demand in the form of a “workload” that is associated with “a processing parameter,” and (2) identifying the claimed increase in demand above the threshold by “determining whether [the] processing parameter ... is greater than ... a threshold value.” *See* Section XIII.D.2.d; Ex-1005, ¶¶31-31.

Third, Mathieson teaches enabling all of the first plurality of cores (Mathieson’s cores 210) for the increase in demand (workload) above the threshold (Mathieson’s threshold value), without disabling any of the second plurality of cores

(Mathieson's cores 220). Ex-1002, ¶237. Notably, while Figure 4A (reproduced above) illustrates method 400A terminating at step 408, Mathieson discloses that step 406 (green), during which "the controller determines whether a value of the processing parameter is above a threshold value," occurs "dynamically at regular time intervals." Ex-1005, ¶56. Accordingly, from a state in which only the second set of cores 220 (second plurality of cores) are enabled, a POSITA would understand that method 400A would return to steps 404 and 406 (green) where the workload (demand), as determined by the processing parameter, would again be evaluated and compared against the threshold value (the claimed threshold). Ex-1005, Fig. 4A; Ex-1002, ¶237. If the workload (demand) is greater than the threshold value, CPU 102 switches to the first mode of operation, corresponding to a state in which the first set of cores 210 (first plurality of cores) are enabled (step 402 (purple) of Figure 4A, above). Ex-1005, ¶56, Fig. 4A; Ex-1002, ¶237. As explained for Element 1[b][i], based on Mathieson's disclosure that "the CPU 102 may operate in both the first mode and the second mode simultaneously," a POSITA would have understood that both sets of cores 210 and 220 may be enabled and operating together at step 402 (purple), or at least would have been motivated to do so. Ex-1005, ¶31; Section XIII.D.2.d; Ex-1002, ¶238.

Furthermore, a POSITA would have understood that the second set of cores 220 (the second plurality of cores) would remain enabled throughout the transition

from the second mode of operation (with only the second set of cores 220 enabled), to the first mode of operation (with both the first and second sets of cores 210, 220 enabled), because disabling the second set of cores 220 only to re-enable them would detrimentally affect performance and power consumption. Ex-1002, ¶239.

Carmack, incorporated by reference into Mathieson, Ex-1005, ¶1, similarly illustrates that, from a state in which the second plurality of cores (Carmack's first cores 110) are operating, the first plurality of cores (Carmack's second cores 120) are enabled, for an increase in demand (workload), without disabling any of the second plurality of cores (Carmack's first cores 110). Ex-1002, ¶240. In particular, as illustrated in Figure 3 of Carmack (annotated below), if workload (demand) is within Carmack's first range at a first measurement time, and then at the next measurement time the workload has increased above a threshold value such that it lands within the third range, Carmack's system will transition from a state in which only the first cores 110 (second plurality of cores) are enabled (step 320, blue) to a state in which both the first and second cores 110 and 120 are enabled (step 340, purple). Ex-1006, ¶¶13, 21, Fig. 3; Ex-1002, ¶240.




Figure 3

And, as explained above, a POSITA would have been motivated to keep the first cores 110 (second plurality of cores) enabled through this transition, based on power and performance considerations. Ex-1002, ¶241.

Finally, as explained for Element 1[b][i], Mathieson also discloses that disabling/enabling cores is done using hardware (the claimed power management hardware). Section XIII.D.2.d; Ex-1005, ¶43; Ex-1002, ¶242.

- b. Element 8[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.**

For the reasons described above for Element 1[b][ii], Mathieson teaches Element 8[b]. Section XIII.D.2.e; Ex-1002, ¶243.

8. Independent Claims 9 and 17:

- a. Element 9[preamble]: A method comprising:**

Even if the preamble is limiting, Mathieson teaches a method. Ex-1002, ¶244. For example, Figure 4A of Mathieson presents “method steps for switching between modes of operation of a processing complex.” Ex-1005, ¶54, Cls. 1-11.

- b. Element 17[preamble]: A non-transitory machine readable medium containing program code that when processed by a machine causes a method to be performed, the method comprising:**

Even if the preamble is limiting Mathieson teaches a non-transitory machine readable medium containing program code. Ex-1002, ¶245. Mathieson discloses “[a] computer-readable medium including instructions that, when executed, cause a processing complex to perform the steps of” Ex-1005, Cl. 12.

- c. Elements 9[a][i] and 17[a][i]: operating a multi-core processor such that a first plurality of cores and a second plurality of cores execute a same instruction set,**

Mathieson teaches Elements 9[a][i] and 17[a][i]. *See* Section XIII.D.2.b; Ex-1002, ¶¶246.

- d. Elements 9[a][ii] and 17[a][ii]: wherein the second plurality of cores consume less power, for a same applied operating frequency and supply voltage, than the first plurality of cores; and**

Mathieson teaches Elements 9[a][ii] and 17[a][ii]. *See* Section XIII.D.2.c; Ex-1002, ¶¶247.

- e. Elements 9[b][i] and 17[b][i]: disabling with power management hardware, from a state where the first plurality of cores and the second plurality of cores are enabled, all of the first plurality of cores for a drop in demand below a threshold without disabling any of the second plurality of cores,**

Mathieson teaches Elements 9[b][i] and 17[b][i]. *See* Section XIII.D.2.d; Ex-1002, ¶¶248.

- f. Element 9[b][ii] and 17[b][ii]: wherein an operating system executing on the multi-core processor monitors a demand for the multi-core processor and controls the power management hardware based on the demand.**

Mathieson teaches Elements 9[b][ii] and 17[b][ii]. *See* Section XIII.D.2.e; Ex-1002, ¶¶249.

- 9. Dependent Claims 10 and 18: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that have narrower logic gate driver transistors than corresponding logic gates of the first plurality of cores.**

Mathieson teaches Claims 10 and 18. *See* Section XIII.D.3; Ex-1002, ¶¶250.

- 10. Dependent Claims 11 and 19: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating of the second plurality of cores comprises driving logic gates that consume less power than corresponding logic gates of the first plurality of cores.**

Mathieson teaches Claims 11 and 19. *See* Section XIII.D.4; Ex-1002, ¶¶251.

- 11. Dependent Claims 12 and 20: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the second plurality of cores at a maximum operating frequency that is less than a maximum operating frequency of the first plurality of cores.**

Mathieson teaches Claims 12 and 20. *See* Section XIII.D.5; Ex-1002, ¶¶252.

- 12. Dependent Claims 15 and 23: The [method of claim 9/non-transitory machine readable medium of claim 17], wherein the operating comprises operating the first plurality of cores at a maximum operating frequency in the state.**

Mathieson teaches Claims 15 and 23. *See* Section XIII.D.6; Ex-1002, ¶¶253.

- 13. Dependent Claims 16 and 24: The [method of claim 9/non-transitory machine readable medium of claim 17], further comprising**

- a. Elements 16[a] and 24[a]: enabling, with the power management hardware, all of the first plurality of cores for an increase in demand above the threshold without disabling any of the second plurality of cores,**

Mathieson teaches Elements 16[a] and 24[a]. *See* Section XIII.D.7.a; Ex-1002, ¶¶254.

- b. Elements 16[b] and 24[b]: wherein an operating system is to monitor a demand for the multi-core processor and control the power management hardware based on the demand.**

Mathieson teaches Element 16[b]. *See* Section XIII.D.7.b; Ex-1002, ¶¶255.

E. Ground 5: Claims 5-6, 13-14, and 21-22 Are Rendered Obvious By Mathieson/Sutardja/Rychlik

1. Dependent Claims 5, 13, and 21:

a. Elements 5[a], 13[a], 21[a]

Mathieson/Sutardja/Rychlik teaches Elements 5[a], 13[a], and 21[a]. Ex-1002, ¶¶256-260. Mathieson discloses that “the number of active cores [may be] increased or decreased” as part of Mathieson’s power management techniques. Ex-1005, ¶76. Based on this teaching, a POSITA would have been motivated to search for references, such as Rychlik, that disclose implementation details relating to enabling/disabling individual cores for power management purposes. Ex-1002, ¶257.

Just like Mathieson, Rychlik discloses power management techniques in multi-core systems. Ex-1009, Abstract. Ex-1002, ¶258. As explained for Ground 2, Rychlik teaches Elements 5[a], 13[a], and 21[a]. *See* Section XIII.B.1.a.

A POSITA would have been motivated to disable an additional core in Mathieson’s second set of cores, whenever the workload to be processed by Mathieson’s CPU 102 dropped below a next lower threshold, until only one core

remained enabled, as taught by Rychlik, because doing so would prevent Mathieson's second set of cores from "wasting power." Ex-1009, ¶88; Ex-1002, ¶259. A POSITA would have had a reasonable expectation of success in implementing Rychlik's power management technique in Mathieson's multi-core system, because it amounts to use of a known technique (disabling cores based on demand, as taught by Rychlik) to improve a similar system (Mathieson's multi-core system) in the same way (by reducing power consumption during periods of low demand). Ex-1002, ¶260.

b. Elements 5[b], 13[b], 21[b]

Mathieson/Sutardja/Rychlik teaches Elements 5[b], 13[b], and 21[b]. Ex-1002, ¶¶261-262. As explained for Ground 2, Rychlik teaches Elements 5[b], 13[b], and 21[b]. *See* Section XIII.B.1.b. And, as explained above for Elements 5[a], 13[a], and 21[a], a POSITA would have been motivated to implement Rychlik's power management technique in Mathieson's multi-core processor. Section XIII.E.1; Ex-1002, ¶262.

2. Dependent Claims 6, 14, 22

Mathieson/Sutardja/Rychlik teaches Claims 6, 14, and 22. Ex-1002, ¶263. As explained for Ground 2, Rychlik teaches the added limitations of these claims. *See* Section XIII.B.2. And, as explained above for Elements 5[a], 13[a], and 21[a], a

POSITA would have been motivated to implement Rychlik's power management technique in Sutardja's multi-core processor. Section XIII.E.1; Ex-1002, ¶263.

XIV. CONCLUSION

Accordingly Claims 1-24 of the '080 Patent should be cancelled.

Respectfully Submitted,

/s/ Ben Haber

Ben Haber (Reg. No. 67,129)

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner certifies that this petition includes 13,992 words, as measured by Microsoft Word, exclusive of the table of contents, mandatory notices under § 42.8, certificates of service, word count, and exhibits.

CERTIFICATE OF SERVICE (37 C.F.R. § 42.6(e)(1))

The undersigned hereby certifies that the above document was served on February 14, 2023, by filing this document through the Patent Trial and Appeal Board P-TACTS System, as well as delivering a copy via express mail upon the following attorneys of record for the Patent Owner and Petitioner:

NDWE LLP / Intel
99 Almaden Boulevard, Suite 575
San Jose, CA 95113

A courtesy copy was sent to the below counsel via electronic mail:

Andrea Leigh Fair
Charles Everingham, IV
Claire Abernathy Henry
1507 Bill Owens Parkway
Longview, TX 75604
Phone: (903) 757-6400
Fax: (903) 757-2323
Email: andrea@wsfirm.com
Email: ce@wsfirm.com
Email: claire@wsfirm.com

Michael T. Renaud
Mint Levin Cohn Ferris Glovasky & Popeo PC
One Financial Center, Suite 100
Boston, MA 02111
Phone: (617) 542-6000
Fax: (617) 542-2241
Email: mtrenaud@mintz.com

Brian E. Farnan
Michael J. Farnan
Farnan LLP
919 North Market Street, 12th Floor

Wilmington, DE 19801
Phone: (302) 777-0300
Fax: (302) 777-0301
Email: bfarnan@farnanlaw.com
Email: mfarnan@farnanlaw.com

Deanna Tanner Okun
AMS Trade LLP
1133 Connecticut Avenue NW, Suite 1200
Washington, DC 20036
Phone: (202) 467-6300
Fax: (202) 466-2006
Email: okun@adduci.com

Dated: February 14, 2023

Respectfully submitted,

/s/ Ben Haber

Ben Haber (Reg. No. 67,129)
O'Melveny & Myers LLP
400 South Hope Street, 18th Floor
Los Angeles, CA 90071
Telephone: (213) 430-6000
Fax: (213) 430-6407
E-Mail: bhaber@omm.com
Attorney for Petitioner