

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO., LTD.,
SAMSUNG ELECTRONICS AMERICA, INC., AND
QUALCOMM INCORPORATED,
Petitioner

v.

DAEDALUS PRIME LLC,
Patent Owner.

DECLARATION OF TREVOR MUDGE
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF
U.S. PATENT NO. 10,705,588

I.	INTRODUCTION	1
A.	<u>Qualifications</u>	<u>1</u>
B.	<u>Materials Considered</u>	<u>3</u>
II.	<u>LEGAL STANDARDS</u>	<u>4</u>
A.	<u>Anticipation</u>	<u>5</u>
B.	<u>Obviousness</u>	<u>6</u>
C.	<u>Level of Ordinary Skill in the Art</u>	<u>9</u>
D.	<u>Claim Construction</u>	<u>10</u>
III.	<u>SUMMARY OF GROUNDS</u>	<u>11</u>
IV.	<u>THE '588 PATENT</u>	<u>12</u>
A.	<u>Overview of the '588 Patent</u>	<u>12</u>
B.	<u>Prosecution History of the '588 Patent</u>	<u>13</u>
V.	<u>TECHNOLOGY BACKGROUND</u>	<u>14</u>
A.	<u>Power, voltage, and frequency control approaches</u>	<u>14</u>
B.	<u>Multi-core processing approaches</u>	<u>18</u>
C.	<u>NorthBridge – SouthBridge Architecture</u>	<u>20</u>
D.	<u>Interconnects and interconnect approaches</u>	<u>21</u>
E.	<u>Memory and Cache</u>	<u>23</u>
VI.	<u>THE PRIOR ART IN THE APPLIED INVALIDITY GROUNDS</u>	<u>24</u>
A.	<u>Machnicki</u>	<u>24</u>
B.	<u>Cho</u>	<u>26</u>
C.	<u>Gold</u>	<u>30</u>
D.	<u>Aguilar</u>	<u>33</u>
E.	<u>Shakkarwar</u>	<u>34</u>
VII.	<u>CLAIM CONSTRUCTION</u>	<u>37</u>
VIII.	<u>DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS</u>	<u>37</u>
A.	<u>Ground 1: Claims 1-2, 4, 8-9, 11, 15-20 are rendered obvious by Machnicki</u>	<u>37</u>
1.	<u>Independent Claim 1</u>	<u>37</u>

a.	Element 1[pre]: A processor comprising:.....	<u>37</u>
b.	Element 1[a]: a plurality of cores including a first core operable at a first voltage and <u>at a first</u> frequency, and.....	<u>38</u>
c.	Element 1[b]: a second core operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency;	<u>40</u>
d.	Element 1[c]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and	<u>43</u>
e.	Element 1[d]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first, second, and third frequencies.	<u>46</u>
2.	Dependent Claim 2: The processor of claim 1, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.....	<u>50</u>
3.	Dependent Claim 4: The processor of claim 2, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency.	<u>52</u>
4.	Independent Claim 8	<u>53</u>
a.	Element 8[pre]: A system comprising:.....	<u>53</u>
b.	Element 8[a]: a system memory;.....	<u>53</u>
c.	Element 8[b]: a plurality of cores coupled to the system memory;.....	<u>54</u>

- d. Element 8[c]: a first core of the plurality of cores operable at a first voltage and at a first frequency, and a second core of the plurality of cores operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency; 55
- e. Element 8[d]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and 56
- f. Element 8[e]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first and second frequencies.....56
5. Dependent Claim 9: The system of claim 8, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.....56
6. Dependent Claim 11: The system of claim 9, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency. 56
7. Dependent Claim 15: The system of claim 8, further comprising: a chipset coupled to the plurality of cores, the chip set comprising a plurality of interfaces to couple the cores to a plurality of I/O devices.....57
8. Dependent Claim 16: The system of claim 15, wherein one of the I/O devices comprises an audio generation device.....59
9. Dependent Claim 17: The system of claim 16, wherein one of the I/O devices comprises a video rendering device. 60
10. Dependent Claim 18: The system of claim 17, wherein one of the I/O devices comprises a user input device. 60
11. Dependent Claim 19: The system of claim 18, wherein one

	of the I/O devices comprises a network communication device.	<u>60</u>
12.	Independent Claim 20	<u>61</u>
a.	Element 20[pre]: An apparatus comprising:.....	<u>61</u>
b.	Element 20[a]: means for operating a first core at a first voltage and at a first frequency while concurrently operating a second core at a second voltage independent of the first voltage and at a second frequency different from the first frequency;	<u>62</u>
c.	Element 20[b]: means for independently operating a graphics processor at a third frequency different from the first frequency and the second frequency	<u>62</u>
d.	Element 20[c]: interconnect means to couple the graphics processor to a memory device, the interconnect means operable at a fourth frequency different from the first, second, and third frequencies.	<u>63</u>
B.	Ground 2: Claims 3, 7, 10, and 14 are rendered obvious by Machnicki and Aguilar.	<u>64</u>
1.	A POSITA would have been motivated to combine Machnicki and Aguilar's teachings and would have had a reasonable expectation of success.....	<u>64</u>
2.	Dependent Claim 3: The processor of claim 1, further comprising: a non-core power domain comprising a cache shared by the plurality of cores.....	<u>65</u>
3.	Dependent Claim 7: The processor of claim 3, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively.....	<u>68</u>
4.	Dependent Claim 10: The system of claim 8, further comprising: a non-core power domain comprising a cache shared by the plurality of cores.....	<u>69</u>
5.	Dependent Claim 14: The system of claim 10, further comprising: power management circuitry and/or logic to	

	cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively.....	<u>69</u>
C.	Ground 3: Claims 5-6, and 12-13 are rendered obvious by Machnicki and Shakkarwar.....	<u>69</u>
1.	A POSITA would have been motivated to combine Machnicki and Shakkarwar’s teachings and would have had a reasonable expectation of success.....	<u>69</u>
2.	Dependent Claim 5: The processor of claim 4, wherein the second power management circuitry and/or logic includes a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.....	<u>71</u>
3.	Dependent Claim 6: The processor of claim 5, wherein the second power management circuitry and/or logic is to select the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.....	<u>75</u>
4.	Dependent Claim 12: The system of claim 11, wherein the second power management circuitry and/or logic includes a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.....	<u>77</u>
5.	Dependent Claim 13: The system of claim 12, wherein the second power management circuitry and/or logic is to select the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.....	<u>77</u>
D.	Ground 4: Claims 1-2, 4, 8-9, <u>11</u> and 20 are rendered obvious by Cho and Gold.....	<u>78</u>
1.	A POSITA would have been motivated to combine Cho and Gold’s teachings and would have had a reasonable expectation of success.....	<u>78</u>
2.	Independent Claim 1	<u>85</u>
a.	Element 1[pre]: <u>A processor comprising:</u>	<u>85</u>

- b. Element 1[a]: a plurality of cores including a first core operable at a first voltage and at a first frequency, and.....86
 - c. Element 1[b]: a second core operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency; 87
 - d. Element 1[c]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and 96
 - e. Element 1[d]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first, second, and third frequencies. 100
- 3. Dependent Claim 2: The processor of claim 1, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.....104
- 4. Dependent Claim 4: The processor of claim 2, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency. 107
- 5. Independent Claim 8 108
 - a. Element 8[pre]: A system comprising: 108
 - b. Element 8[a]: a system memory; 109
 - c. Element 8[b]: a plurality of cores coupled to the system memory; 110

- d. Element 8[c]: a first core of the plurality of cores operable at a first voltage and at a first frequency, and a second core of the plurality of cores operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency; 111
- e. Element 8[d]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and 111
- f. Element 8[e]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first and second frequencies.....111
- 6. Dependent Claim 9: The system of claim 8, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.....111
- 7. Dependent Claim 11: The system of claim 9, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency. 112
- 8. Independent Claim 20 112
 - a. Element 20[pre]: An apparatus comprising:.....112
 - b. Element 20[a]: means for operating 112
 - c. Element 20[b]: means for independently operating 113
 - d. Element 20[c]: interconnect means 114
- E. Ground 5: Claims 3, 7, 10, and 14-19 are rendered obvious by Cho, Gold, and Aguilar. 115
 - 1. A POSITA would have been motivated to combine Cho, Gold, and Aguilar's teachings and would have had a reasonable expectation of success.....115

2. Dependent Claim 3: The processor of claim 1, further comprising: a non-core power domain comprising a cache shared by the plurality of cores......119
3. Dependent Claim 7: The processor of claim 3, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively......123
4. Dependent Claim 10: The system of claim 8, further comprising: a non-core power domain comprising a cache shared by the plurality of cores......123
5. Dependent Claim 14: The system of claim 10, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively......123
6. Dependent Claim 15: The system of claim 8, further comprising: a chipset coupled to the plurality of cores, the chip set comprising a plurality of interfaces to couple the cores to a plurality of I/O devices......124
7. Dependent Claim 16: The system of claim 15, wherein one of the I/O devices comprises an audio generation device......126
8. Dependent Claim 17: The system of claim 16, wherein one of the I/O devices comprises a video rendering device......127
9. Dependent Claim 18: The system of claim 17, wherein one of the I/O devices comprises a user input device......128
10. Dependent Claim 19: The system of claim 18, wherein one of the I/O devices comprises a network communication device......128
- F. Ground 6: Claims 5-6, and 12-13 are rendered obvious by Cho, Gold, and Shakkarwar......129
 1. A POSITA would have been motivated to combine Cho, Gold, and Shakkarwar’s teachings and would have had a reasonable expectation of success.....129
 2. Dependent Claim 5: The processor of claim 4, wherein the second power management circuitry and/or logic includes

	<u>a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.</u>	<u>132</u>
3.	<u>Dependent Claim 6: The processor of claim 5, wherein the second power management circuitry and/or logic is to select the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.</u>	<u>136</u>
4.	<u>Dependent Claim 12.</u>	<u>138</u>
5.	<u>Dependent Claim 13.</u>	<u>138</u>
<u>IX.</u>	<u>CONCLUSION</u>	<u>138</u>

I. INTRODUCTION

1. I have been retained by Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc. (collectively, “Samsung” or “Petitioner”), as an independent expert in this proceeding before the Patent Trial and Appeal Board (“PTAB” or “Board”). I understand that Samsung is requesting that the Board institute an *inter partes* review (“IPR”) proceeding of U.S. Patent No. 10,705,588 (“the ’588 Patent”) (Ex-1001), currently assigned to Daedalus Prime LLC (“Patent Owner” or “PO”).

2. I am not and have never been an employee of Samsung. I am being compensated at my usual and customary rate of \$600 per hour. No part of my compensation depends on the outcome of this proceeding, and I have no other interest in this proceeding.

3. I have been asked to provide my independent analysis of the ’588 Patent in light of the prior art publications cited below. I have also been asked to consider the state of the art and prior art available as of October 27, 2011. Based on the prior art discussed in this declaration, it is my opinion that Claims 1-20 of the ’588 Patent is unpatentable for the reasons provided below.

A. Qualifications

4. I am currently a the Bredt Family Professor of Computer Science and Engineering at the University of Michigan, Ann Arbor. All of my opinions stated in

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

this declaration are based on my own personal knowledge and professional judgment. In forming my opinions, I have relied on my over 45 years of research, academic, industry, and consulting engineering experience in IC (integrated circuit) processing, semiconductor devices, and computer architecture with an emphasis on power and energy control.

5. I am over 18 years of age and, if I am called upon to do so, I would be competent to testify as to the matters set forth herein. I understand that a copy of my current curriculum vitae, which details my education and professional and academic experience, is being submitted by Petitioner as Exhibit 1003. The following provides an overview of some of my experience that is relevant to the matters set forth in this declaration.

6. I received the Ph.D. in Computer Science from the University of Illinois, Urbana. I am the Brecht Family Professor of Computer Science and Engineering at the University. I am the author of numerous papers on computer architecture, programming languages, VLSI design, and computer vision. I have chaired 57 PhD theses in these areas. In 2014 I received the ACM/IEEE CS Eckert-Mauchly Award for “pioneering contributions to low-power computer architecture and its interaction with technology.” This is known as the computer architecture community's most prestigious award. I also received the University of Illinois

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

Distinguished Alumni Award. I am a Life Fellow of the IEEE, a Fellow of the ACM,
and a member of the IET and the British Computer Society.

7. Based on my experience and education, I believe that I am qualified to
opine as to the knowledge and level of skill of one of ordinary skill in the art at the
time of the alleged invention of the '588 Patent, as well as the state of the art at that
time.

B. Materials Considered

8. In forming my opinions, I have reviewed the following documents:¹

<u>Ex-1001</u>	U.S. Patent No. 10,705,588
Ex-1002	Declaration of Dr. Trevor Mudge
Ex-1003	Curriculum Vitae of Dr. Trevor Mudge
Ex-1004	File History of U.S. Patent No. 10,705,588
Ex-1005	U.S. Published Patent Application No. 2010/0073068 (“Cho”)
Ex-1006	U.S. Patent No. 8,468,373 (“Machnicki”)
Ex-1007	U.S. Patent No. 7,552,346 (“Aguilar”)
Ex-1008	U.S. Patent No. 7,392,411 (“Shakkarwar”)
Ex-1009	U.S. Published Patent Application No. 2009/0235108 (“Gold”)
Ex-1010	INTENTIONALLY LEFT BLANK
Ex-1011	INTENTIONALLY LEFT BLANK
Ex-1012	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.

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

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
Ex-1020	Claim Mapping Table
<u>Ex-1021</u>	INTENTIONALLY LEFT BLANK
Ex-1022	Bharadwaj Amrutur, Nandish Mehta, Satyam Dwivedi, and Ajit Gupte, Adaptative Techniques to Reduce Power in Digital Circuits (July 2011)
Ex-1023	Churiwala, Sanjay, Tackling multiple clocks in SoCs, https://www.eetimes.com/tackling-multiple-clocks-in-socs/
Ex-1024	System Bus Fundamentals, http://www.edwardbosworth.com/My5155_Slides/Chapter12/SystemBusFundamentals.htm
Ex-1025	Intel® E8501 Chipset North Bridge (NB), Datasheet (May 2006)
Ex-1026	U.S. Patent No. 8,397,090 to Gunther (“Gunther”)
Ex-1027	U.S. Patent No. 9,086,883 to Thomson (“Thomson”)
Ex-1028	U.S. Patent No. 7,263,457 to White (“White”)
Ex-1029	U.S. Patent No. 8,438,416 to Kocev (“Kocev”)
Ex-1030	U.S. Patent Publication No. 2007/0250736 to Gaskins (“Gaskins”)

II. LEGAL STANDARDS

9. In forming my opinions and considering the subject matter of the ’588 Patent and its claim in light of the prior art, I am relying on certain legal principles

that counsel in this case explained to me. My understanding of these concepts is summarized below.

10. I understand that earlier publications and patents may act to render a patent unpatentable for one of two reasons: (1) anticipation, and (2) obviousness.

A. Anticipation

11. It is my understanding that the claims of a patent are anticipated by a prior art reference if each and every element of the claim is found either explicitly or inherently in the reference. I understand that inherency requires a showing that the missing descriptive matter in the claim is necessarily present in the allegedly anticipating reference, and that it would have been so recognized by a person of ordinary skill in the art (“POSITA”).

12. I understand that when a challenged claim covers several structures, either generically or as alternatives, the claim is deemed anticipated if any of the structures within the scope of the claim is found in the prior art reference.

13. Although anticipation typically involves the analysis of a single prior art reference, I understand that additional references may be used to show that the prior art reference has enabling disclosure (i.e., allows a POSITA to make the invention without undue experimentation), to explain the meaning of a term used in the prior art reference, and/or to show that a characteristic is inherent in the prior art reference.

B. Obviousness

14. I understand that a claim is invalid as obvious if it would have been obvious to a person of ordinary skill in the art at the time the alleged invention was made. This means that even if all of the elements of the claim cannot be found in a single prior art reference that would anticipate the claim, a person of ordinary skill in the art who was aware of the prior art would have been able to come up with the claimed invention. This may be the case, for example, where the missing element represents only an insubstantial different over the prior art or a reconfiguration of a known system. I understand that in an obviousness determination, the person of ordinary skill in the art is presumed to have knowledge of all material prior art.

15. I understand that an obviousness analysis requires an understanding of the scope and content of the prior art, any differences between the alleged invention and the prior art, and the level of ordinary skill in evaluating the pertinent art.

16. I understand that when a product is available, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill in the art can implement a predictable variation, obviousness likely bars its patentability. For the same reason, if a technique has been used to improve one device and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique would have been obvious.

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

17. I understand that whether a prior art reference renders a patent claim unpatentable as obvious is determine from the perspective of a person of ordinary skill in the art at the time of the alleged invention. I have been told that there is not requirement that the prior art contain an express suggestion to combine known elements to achieve the claimed invention, but a suggestion to combine known elements to achieve the claimed invention may come from the prior art, as filtered through the knowledge of one skilled in the art. In addition, I have been told that the inferences and creative steps a person of ordinary skill in the art would employ are relevant to the determination of obviousness.

18. I understand that one may consider, e.g., whether (1) the change was merely the predictable result of using prior art elements according to their known functions, or whether it was the result of true inventiveness; (2) there is some teaching or suggestion in the prior art to make the modification or combination of elements claimed in the patent; (3) the claimed innovation applies a known technique that had been used to improve a similar device or method in a similar way; (4) the claimed invention would have been obvious to try, meaning that the claimed innovation was one of a relatively small number of possible approaches to the problem with a reasonable expectation of success by those skilled in the art; (5) the invention merely substituted one known element for another known element in order to obtain predictable results; (6) the invention merely applies a known technique to

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

a known device, method, or product to yield predictable results; or (7) known work in one field of endeavor may have prompted variations of it for use in either the same field or a different one based on design incentives or other market forces that would have been predictable to a person of ordinary skill in the art.

19. I further understand that certain factors may support or rebut the obviousness of a claim. I understand that such secondary considerations include, among other things, commercial success of the patented invention, skepticism of those having ordinary skill in the art at the time of the invention, unexpected results of the invention, any long-felt but unsolved need in the art that was satisfied by the alleged invention, the failure of others to make the alleged invention, praise of the alleged invention by those having ordinary skill in the art, and copying of the alleged invention by others in the field. I understand that there must be a nexus—that is, a connection—between any such secondary considerations and the alleged invention. I also understand that contemporaneous and independent invention by others is a secondary consideration tending to show obviousness.

20. I am not aware of any allegations by the named inventor of the '588 Patent or any assignee of the '588 Patent that any secondary considerations tend to rebut the obviousness of the '588 Patent.

21. Additionally, I understand that in considering obviousness, it is important not to use the benefit of hindsight derived from the patent under consideration.

C. Level of Ordinary Skill in the Art

22. When interpreting a patent, I understand that it is important to identify the relevant art pertaining to that patent, as well as the level of ordinary skill in that art at the time of the claimed invention. The “art” is the field of technology to which the patent is related.

23. I am informed and understand that the POSITA is a hypothetical person who is presumed to know the relevant prior art. I understand that the actual inventor’s skill is not determinative of the level of ordinary skill. I further understand that factors that may be considered in determining level of skill include: (i) the types of problems encountered in the art; (ii) prior art solutions to those problems; (iii) the rapidity with which innovations are made; (iv) the sophistication of the technology; and (v) the educational level of active workers in the field. I understand that not all such factors may be present in every case, and one or more of them may predominate.

24. I understand that a POSITA is one who is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. A POSITA would have had knowledge of circuit boards and

microcircuits, including components thereof, and related technologies as of October 27, 2010.

25. 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.

26. Based on my education and experience, I would have, at least, met the
criteria for a POSITA in October 27, 2011, and I still exceed it today.

D. Claim Construction

27. I understand that the United States Patent and Trademark Office
interprets claim terms in an *inter partes* review proceeding under the same claim
construction standard that is used in a United States federal court. I understand that
under this standard, the meaning of claim terms is considered from the viewpoint of
a POSITA at the time of the alleged invention.

28. I have been informed that claim terms are generally given their ordinary
and customary meaning as understood by one of ordinary skill in the art in light of
the specification and the prosecution history pertaining to the patent. I understand,

however, that claims terms are generally not limited by the embodiments described in the specification.

29. I understand that in addition to the claims, specification, and prosecution history, other evidence may be considered to ascertain the meaning of claim terms, including textbooks, encyclopedias, articles, and dictionaries. I have been informed that this other evidence is often less significant and less reliable than the claims, specification, and prosecution history.

III. SUMMARY OF GROUNDS

30. I understand that Petitioner challenges the validity of Claims 1-20 of the '588 Patent. In particular, I understand that Petitioner asserts the following grounds for challenging the claims of the '588 Patent:

Ground	Summary
1	Claims 1-2, 4, 8-9, 11, 15-20 are obvious in view of Machnicki (Ex-1006)
2	Claims 3, 7, 10, 14 are obvious over Machnicki (Ex-1006) in view of Aguilar (Ex-1007)
3	Claims 5-6, 12-13 are obvious over Machnicki (Ex-1007) in view of Shakkarwar (Ex-1008)
4	Claims 1-2, 4, 8-9, <u>11</u> are obvious over Cho (Ex-1005) in view of Gold (Ex-1009)

Ground	Summary
5	Claims 3, 7, 10, 14-20 are obvious over Cho (Ex-1005) in view of Gold (Ex-1009) and Aguilar (Ex-1007)
6	Claims 5-6, 12-13 are obvious over Cho (Ex-1005) in view of Gold (Ex-1009) and Shakkarwar (Ex-1008)

IV. THE '588 PATENT

A. Overview of the '588 Patent

31. The '588 Patent purports to describe a processor having “multiple domains,” including a “core domain” and a “non-core domain.” Ex-1001, Abstract.

Figure 3, annotated below, illustrates the '588 Patent's claimed processor, 300:




FIG. 3

Id., Fig. 3 (annotated).

32. As Figure 3 illustrates, the processor includes a “plurality of cores 310_a-310_n” (red and blue). Ex-1001, 5:36-40. Each core may “be of an independent power domain” and “can be configured to operate at an independent voltage and/or frequency” *Id.* at 5:51-44. The cores are connected to “interconnect 320” (green), which in turn connects the cores to a “shared cache 330” (maroon). *Id.* at 5:44-46.

33. The processor also includes “power control unit 355” (yellow). Ex-1001, 5:46-50. The power control unit 355 may include “frequency control logic 359,” which “perform[s] dynamic control of interconnect frequency responsive to a request for a given domain” in either the core or non-core domains. *Id.* at 5:51-55.

34. The ’588 Patent also discloses a “graphics domain” within the non-core domain that includes “one or more graphics engines.” Ex-1001, 2:19-21.

B. Prosecution History of the ’588 Patent

35. I understand the ’588 Patent was filed on January 16, 2019, as a continuation of application No. 16/249,103 (now U.S. Patent No. 10,428,181), claiming priority to application No. 13/282,892 (now U.S. Patent No. 8,832,478), filed on October 27, 2011.

36. On December 5, 2019, the applicant terminally disclaimed pending Claims 2, 9, and 21. Ex-1004, 0156. On March 2, 2020, the examiner allowed all pending claims based on the applicant’s terminal disclaimer. *Id.* at 0164.

37. I understand the prior art relied on in this Petition was not before the PTO during prosecution of the '588 Patent. Thus, I understand that the Petition presents substantially new arguments that were not considered during prosecution.

V. TECHNOLOGY BACKGROUND

38. To provide background for the element-by-element analysis of the claims to follow, below I will present an overview of the state of the art existing at the time of the alleged invention relating to processors generally, and as it relates to multi-core devices, and power, voltage and frequency management approaches, among other relevant technology. As I will describe below from a high-level functional perspective, in the next section based on specific prior art, and then in subsequent sections on a limitation-by-limitation basis, all of these technologies and techniques were well-known to those of ordinary skill in the art at the time of the alleged invention claimed in the '588 Patent, and a POSITA would have readily understood the combination of elements of Claim 1-20 to have been obvious.

A. Power, voltage, and frequency control approaches

39. There are two key relationships linking core frequency and voltage to core power consumption. They are: (1) power is proportional to V^2f – where V is the voltage to operate the core, and f is the frequency of operation of the core; and (2) the frequency of operation is directly proportional to the voltage, i.e., f is proportional to V. Higher frequency requires a great core voltage resulting in the

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

consumption of much more power. Observation (1) tells us that we can control power by simply changing f . For example, halving/doubling f , halves/doubles the production of results but also halves/doubles the power consumed. A more advantageous method for saving power is to control the voltage term, V . The fact that power depends on the square of the voltage (V^2) means that the savings can be significant: halving the voltage reduces the power consumption to one-fourth. Unfortunately, (2) tells us that we also have to reduce f , but only by half. Thus using two cores at half frequency to replace one fast core at full frequency cuts total power by half, but the same amount of computation is achieved, if we take frequency as a crude measure of performance. In the late 1990's power consumption was becoming a limiting factor to improvements in computer performance and multi-cores seemed to provide a solution. Although we have presented a highly simplified version of the underlying physics, the basic argument of using multi-cores to save power without sacrificing performance holds, and this led the industry to adopt to multi-cores. In 2005 Intel decide to adopt this strategy and referred to it as its "right-hand turn." This was after Intel had the first i486 processor that include an L1 cache in 1993. In fact, the single chip multi-core concept had been anticipated by researchers at Stanford University some years earlier with their Hydra Multiprocessor-on-a-Chip, in February 1998 (report CSL-TR-98-749, Stanford 1998). The control of power and speed settings on cores, the allocation of cores to tasks, and maximize power saving

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

when those resources are not needed is referred to as DVFS—dynamic voltage and frequency scaling (DVFS).

40. “Limiting power consumption presents a critical issue in computing, particularly in battery powered devices such as laptop computers and cell phones. Limiting power in other computer settings—such as server farms—warehouse-sized buildings filled with Internet service providers’ servers—is also important.”

41. I wrote this in 2000/1 in two review articles, at which time the problem of power consumption was well known and starting to become a design constraint for computer designers. The source of power consumption in integrated circuits is dictated by the voltage that they require and the frequency at which they operate. Power consumption increases in proportion to the frequency at which the core, an integrated circuit, operates but increases as the square of the voltage supply.


42. Additionally, increasing frequency improves performance but requires that the voltage also be increased. Thus, to improve performance by increasing frequency requires a quadratic increase in voltage supply. This in turn has translates to a significant increase in power consumption.

43. It is also a common approach for the “relationship between frequency and minimum supply voltage [to] be precisely captured in a table for a few frequency points” Ex-1022, 0005. This “guarantee[s] that circuit timing is met even for the worst case process and temperature conditions.” *Id.* For example, “[w]hen the

operating frequency is changed, the table is consulted for the appropriate supply voltage value and the voltage regulators are programmed to deliver this voltage.” *Id.*

44. There is also “the interplay between power consumption and performance in microprocessors, and particularly to the reduction of the former and the increase of the latter with respect to the operating temperature of the microprocessor.” Ex-1030 ¶7. U.S. Patent Publication 2007/0250736, filed June 11, 2007, discusses this technique and relates to a “method for increasing microprocessor performance when making a P-state transition by, rather than making a single frequency change, iteratively making multiple frequency changes while transitioning from the current P-State Voltage to the new P-state voltage.” *Id.* ¶15. It details incremental voltage increases and only increasing the operating frequency while monitoring the maximum operating temperature. Specifically, “the microprocessor 102 monitors its temperature while operating at a given frequency. That is, the temperature sensor 132 senses the current operating temperature and provides the temperature 134 to the voltage/frequency control 104.” *Id.* ¶72. Figure 9 is a “flowchart illustrating a method for dynamically operating the microprocessor 102 of FIG. 1 at or near optimum performance within a specified temperature range.” *Id.* ¶96.

Fig. 9



B. Multi-core processing approaches

45. Multi-core processing refers to the use of multiple cores or computer processing units (CPU) to speed up a computing task. This is in contrast to having one core that runs at a multiple of the speed of each core in the multi-core group. Speeding up a single core has limitations as we will see in the next section. The multi-core approach has its own limitations too, because many problems cannot be split into subtasks that can be run simultaneously on several cores. Dependences between subtasks—you need to crack an egg before you can fry it—pose a

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

limitation. One common approach that gets around this problem is to run separate
copies of the on each of the multiple cores and realize a throughput gain, instead of
trying to split the task into subtasks. This way more tasks are achieved per unit time
even though each copy still take the same amount of time it would have done on a
single core. For this to be effective we need to have many of the same task to be
done—many eggs to cook. In fact the multiple tasks need not even be the same.

46. Manufacturers of processor cores, before about the turn of the century,
had been able to increase performance by simply increasing frequency.
Unfortunately, the fact that power increases as square of the voltage finally caught
up with them and meant that increasing frequency quickly became impractical. The
solution that began to emerge was to employ multiple processors to do the work of
one. The reasoning was that two processors could run at half the frequency, but now
the quadratic dependency on supply voltage gave a positive effect, because the
voltage supply could be lowered at a quadratic rate. Thus the power consumption of
two processors could actually be less than that of one fast processor.

47. This observation led the industry to adopt the idea of multiple
processors operating at a lower frequency to replace one fast processor. In about
2005, Intel and its rivals began to perform a “right turn,” to quote their promotional
material, which meant they moved away from using a single fast processor and
replaced it with multiple slower cores on an integrated circuit to meet power

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

limitations, while at the same time maintaining performance. “Each processor may include on or more cores and each processor/core may perform operations independent of the other processors/cores.” Ex-1027, 8:46-48. The background of the ’588 Patent reflects this concept, which discusses a scenario with multiple cores, and was widely employed at the date of the patent. The multiple cores opened up further ways to control power consumption: when not needed some cores could be turned off.

48. Having multiple processors also require multiple clocks. This is because each processors having specialty functions and operate at different frequencies. Ex-1023. “Multiple clocks are inevitable in today's designs.” *Id.*

C. NorthBridge – SouthBridge Architecture

49. A NorthBridge is a common component of may processors. The NorthBridge may include “separate interfaces to communicate with and route communication from multiple logic cores of an integrated circuit to other sub-systems, components or peripherals external to the integrated circuit.” Ex-1028. U.S. Patent No. 8,438,416, filed October 21, 2010 utilizes a NorthBridge and “is directed to a method and apparatus for function-based dynamic power control of a bridge unit, such as a north bridge unit in a computer system.” Ex-1029, 2:44-46. The NorthBridge may also include “an interface to an I/O bridge (e.g., a South bridge), a memory controller, and a graphics processing unit.” *Id.* at 2:57-59.




Fig. 1

D. Interconnects and interconnect approaches

50. Interconnects and buses are commonly used in system components to interconnect the system to components such as memory elements, system components, resources and custom circuitry. Ex-1027. As shown in U.S. Published Patent No. 2010/0073068, published March 25, 2010, an interconnect (“NorthBridge 415”) couples the GPU to the memory and other components. Ex-1005, Fig. 4. Additionally, U.S. Patent No. 9,086,883, filed June 10, 2011, utilizes a bus/interconnect module “which may include an array of reconfigurable logic gates and/or implement a bus architecture.” Ex-1027, 8:67-9:5. “Communications may be provided by advanced interconnects, such as high performance networks on chip (NoCs).” Id. at 9:6-7.





FIG. 2

Ex-1027, Fig. 2.



Ex-1005, Fig.4.

E. Memory and Cache

51. It is common for integrated circuits to include a memory, such as a private or shared cache or a main or system memory. See Ex-1026. Also, multicore processors may include a multi-level cache that includes Level 1 (L1) caches and Level 2 (L2) caches. U.S. Patent No. 9,086,883 filed June 10, 2011, utilizes a L1 and L2 cache system. Ex-1027. “The L2 caches [] may be larger and slow than the L1 caches [] but smaller (and substantially faster) than the main memory unit. *Id.* at 9:31-34. Additionally, “[t]he L1 and L2 caches may be used to store data frequently accessed by the processing units, whereas the main memory 220 may be used to store larger files and data units being accessed by the processing cores.” *Id.* at 9:40-43.




FIG. 2

VI. THE PRIOR ART IN THE APPLIED INVALIDITY GROUNDS

A. Machnicki

52. Machnicki, titled “Modifying Performance Parameters in Multiple Circuits According to a Performance State Table Upon Receiving a Request to Change a Performance State,” issued on June 18, 2013, from an application filed on January 14, 2011.

53. I understand Machnicki is prior art under at least 35 U.S.C. § 102(e).

54. Machnicki discloses a system that “includes an integrated circuit (IC) 100 forming a system-on-a-chip (SoC) or application-specific integrated circuit

(ASIC) coupled to power supply 110.” Ex-1006, 4:9-12. The integrated circuit includes “a set of performance domains 120A-F,” which include several “component[s]” such as “processor 160A” (red), “graphics subsystem 160B” (purple), “memory control circuit 160E” (green), and a “second processor 160F” (blue). Id. at 4:9-25. Figure 1, annotated below, illustrates Machnicki’s integrated circuit 100:




FIG. 1

Id., Fig. 1 (annotated).

55. Machnicki also discloses that “both clock frequencies and voltages are configuration parameters that may be set for each of the performance domains (e.g.

domains 1 through m).” Ex-1006, 8:1-4. As shown in Table 1, each domain has a voltage, V , and a frequency, f :

TABLE I

	Domain 1	Domain 2	...	Domain m
state 1	$f_{1,1}; V_{1,1}$	$f_{2,1}; V_{2,1}$	...	$f_{m,1}; V_{m,1}$
state 2	$f_{1,2}; V_{1,2}$	$f_{2,2}; V_{2,2}$	...	$f_{m,2}; V_{m,2}$
state 3	$f_{1,3}; V_{1,3}$	$f_{2,3}; V_{2,3}$	...	$f_{m,3}; V_{m,3}$
...	...	...	...	...
state n	$f_{1,n}; V_{1,n}$	$f_{2,n}; V_{2,n}$	...	$f_{m,n}; V_{m,n}$

Id. at 8:8-17.

56. By using a performance state table “such as the one illustrated in Table I [], each performance domain may have its parameters changed independently.” Id. at 8:37-43. In the case of clock frequencies, Machnicki discloses “the frequencies in each frequency domain for any given state may be different from each other, and may be other than a multiple of each other.” Id. at 8:37-43


B. Cho

57. Cho, titled “Functional Block Level Thermal Control,” issued on March 25, 2010. Ex-1005, Cover.

58. I understand Cho is prior art under at least 35 U.S.C. § 102(b).

59. Cho discloses an “integrated circuit” that includes a “plurality of functional units” implemented on a die. Ex-1005, Abstract. Cho’s integrated circuit is a “multi-core processor configured for functional block level thermal control.”

Id. ¶38. Cho discloses that its processor “includes four primary functional units—two processor cores 405 and 410 [red and blue], a memory controller hub (NorthBridge 415) [green], and a graphics processing unit (GPU) 42 [purple].” *Id.* Figure 4, annotated below, illustrates Cho’s processor, 400:




Id., Fig. 4 (annotated).

60. Cho also discloses that “[e]ach of the functional units shown is coupled to receive corresponding clock signals” and is “further coupled to receive corresponding voltage control signals.” Ex-1005 ¶40. These “clk” and “vctrl”

signals are shown in the colored boxes corresponding to the same color functional units in Figure 4. Each functional unit's voltages “can be set independently of one another, even if all of the functional units initially receive power from the same source.” Id. ¶21. Moreover, a clock control unit 310, contained in Temperature Control Unit 102 (yellow), is also “configured to adjust the frequency of each of clock signals clk1-clk4, independently of one another, based on signals received from control logic 305.” Id.

61. Cho discloses that each functional unit has a “temperature threshold,” which may include “maximum,” “one or more intermediate,” or “minimum” thresholds at which the “TCU 102 may initiate a temperature control action.” Ex-1005 ¶¶20, 24. Cho further explains that “the temperature thresholds at which TCU 102 may initiate a temperature control action may be different for one or more of the functional units with respect to other ones of the functional units,” and “[t]his includes embodiments wherein the temperature thresholds are unique for each of functional units 105-120.” Id. ¶20. These temperature control actions include “adjusting the frequency of a clock signal received thereby, adjusting the voltage on a voltage plane of the functional unit, or a combination of both of these actions.” Id. ¶40. Specifically, “[c]hanging the frequency of individual clock signals may be performed for temperature control purposes.” Id. ¶36.

62. For example, in Figure 5, a flow chart illustrates three temperature control actions based on exceeding a threshold (light orange), still exceeding a threshold (orange), and dropping below a threshold (indigo). See Ex-1005 ¶¶43-44.



Id., Fig. 5 (annotated).

63. In the Figure 5 example, first temperature control action (515) may be to “reduce the clock frequency” and the second temperature control action (525) may be a subsequent reduction of the clock frequency.” Ex-1005 ¶44. Furthermore, if the temperature drops below a low threshold (530), in next temperature control action (545), “the clock frequency may be increased back to its previous value,” which may be the starting value or the first reduced value as shown in Figure 5. Id.

C. Gold

64. Gold, titled “Automatic Processor Overclocking,” was published on September 17, 2009, from an application filed on March 11, 2008. Ex-1009, Cover. I understand Gold is prior art under at least 35 U.S.C. § 102(b).

65. Gold discloses methods for “performing overclocking for a plurality of processing units,” including “a plurality of processing cores (each of which has a respective standard operating frequency.” Ex-1009 ¶5. Gold discloses a “clock generation unit” that is coupled to “each of the plurality of processing cores,” where the clock generation unit is configured to “generate a respective clock signal for each of the plurality of processing cores.” *Id.*

66. Gold assigns several “performance states” (Pstate) to each core, such as PMax, P0, P1, P2, which each includes a “corresponding input voltage and clock frequency.” *Id.* ¶39. Figure 5 illustrates Gold’s performance states with corresponding frequencies:

500


<i>P-State</i>	<i>Power Usage</i>	<i>Frequency</i>
<i>P Max</i>	<i>19W</i>	<i>2.90 GHz</i>
<i>P 0</i>	<i>15W</i>	<i>2.60 GHz</i>
<i>P 1</i>	<i>12W</i>	<i>2.28 GHz</i>
<i>P 2</i>	<i>9W</i>	<i>1.73 GHz</i>

FIG. 5A

Id., Fig. 5A. As Figure 5A illustrates, each core is operable at least four different frequencies, corresponding to four different power usages. Gold also teaches “a hardware throttling control system that aggressively reduces or even stops the clock of a processing core once thermal limitations are exceeded.” *Id.* ¶ 53.




FIG. 6


D. Aguilar

67. Aguilar, titled “Dynamically Adapting Software for Reducing a Thermal State of a Processor Core Based On Its Thermal Index,” issued on June 23, 2009, from an application filed on May 3, 2006. Ex-1007, Cover.

68. I understand Aguilar is prior art under at least 35 U.S.C. § 102(b).

69. Aguilar describes “processing system 200” that employs a “hub architecture,” including a “north bridge” “south bridge,” and “input/output (I/O) controller hub ICH (204).” Ex-1007, 3:44-52. Aguilar discloses that the NorthBridge contains a “cache.” Id. at 4:57-60. A plurality of input/output devices connect to Aguilar’s system through the I/O controller hub. For example, Aguilar teaches that “[a]udio adapter 216, keyboard and mouse adapter 220, modem 222, read only memory (ROM) 224, hard disk drive (HDD) 226, CD-ROM drive 230, universal serial bus (USB) ports and other communications ports 232, and PCI/PCIe devices 234 [orange] connect to south bridge [yellow] and I/O controller hub 204 [yellow] through bus 238 and bus 240” (blue). Id. at 3:55-60.

70. Figure 3, annotated below, illustrates Aguilar’s hub architecture, including the I/O controller hub 204 (yellow) and I/O devices (orange):




Ex-1007, Fig. 3 (annotated).

E. Shakkarwar

71. Shakkarwar, titled “Systems and Methods for Dynamic Voltage Scaling of Communication Bus to Provide Bandwidth Based on Whether an Application is Active,” issued on June 24, 2008. Ex-1008, Cover.

72. I understand Shakkarwar is prior art under at least 35 U.S.C. § 102(b).

73. Shakkarwar discloses “systems and methods for dynamic power management of electronic devices” Ex-1008, Abstract. Shakkarwar’s dynamic power management system includes several components, including “a data controller 302” (green) that is connected to a “graphics controller 305” (purple) through “data connection 314” (green), as illustrated in Figure 3:



Id., Fig. 4 (annotated).

74. Shakkarwar dynamically manages power by monitoring a component's "bandwidth ... for increasing or decreasing communication requirements." Ex-1008, 6:41-43. The system then detects "a change in processing condition[s]" for a component, and upon detecting a change, "determines bandwidth 403 for an associated data connection based on the active components." Id. at 6:63-7:7. To determine the appropriate bandwidth for a data connection, Shakkarwar utilizes "a look-up table [that] may include a list of clock speeds, and/or voltage levels required

to process information for specific devices and/or groups of devices.” Id. at 7:7-10.

Then, at step 405, “the data rate of a data connection is adjusted,” which occurs by “adjusting operating of a clock generation device.” Id. at 7:26-29.

75. Figure 4 illustrates this process:




FIG. 4

Ex-1008, Fig. 4.

76. Shakkarwar also discloses that “when an application that is known to use 3-D graphics is opened, system software can determine if the bus rate between

the data controller and the graphics controller is high enough to support the graphics requirements.” Ex-1008, 6:49-52.

VII. CLAIM CONSTRUCTION

77. I do not believe that any claim term requires explicit construction to resolve the issues presented in this Petition. I apply the plain and ordinary meaning to each claim term, as that plain and ordinary meaning would have been understood by a POSITA.

78. I reserve the right to offer opinions on any claim constructions proposed in this proceeding or to offer opinions on additional constructions in the district court.

VIII. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS

A. Ground 1: Claims 1-2, 4, 8-9, 11, 15-20 are rendered obvious by Machnicki.

1. Independent Claim 1

a. Element 1[pre]: A processor comprising:

79. Machnicki discloses Claim 1’s preamble to the extent it is limiting. Machnicki discloses its system “includes an integrated circuit (IC) 100 forming a system-on-a-chip (SoC) or application-specific integrated circuit (ASIC) coupled to power supply 110.” Ex-1006, 4:9-12. The integrated circuit includes “a set of performance domains 120A-F,” which include “component[s]” such as “processor 160A,” “graphics subsystem 160B,” “memory control circuit 160E,” and a “second

processor 160F.” *Id.* at 4:9-25. Figure 1, annotated below, illustrates Machnicki’s integrated circuit 100 (green):




FIG. 1

80. Ex-1006, Fig. 1 (annotated). Machnicki’s SoC 100 is thus a processor as described in the preamble.

b. Element 1[a]: a plurality of cores including a first core operable at a first voltage and at a first frequency, and

81. Machnicki discloses Element 1[a]. Machnicki’s SoC 100 includes “processor 160A” and “processor 160F.” Ex-1006, 4:16-25. Machnicki further discloses that “[i]ntegrated circuit 520 may include one or more instances of elements within integrated circuit 100 of FIG. 1.” *Id.* at 10:55-57. Machnicki’s

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

processors 160A and 160F disclose a “plurality of cores” because a POSITA would have known that each processor contains at least one core. See Section V.B. Indeed, the ’588 Patent explains that at the time, it was well known that processors could be “multicore” and include “a plurality of cores 310a-310 n.” Ex-1001, 5:37-40. See also id. at 1:20-26 (“Over time, processor designs have evolved from a collection of independent integrated circuits ... to multicore processors that include multiple processor cores within a single IC package); See Section V.B. A POSITA would thus have known processor 160A includes at least a “first core.” Figure 1, annotated below, illustrates Machnicki’s processors 160A (red) and 160F (blue), each containing at least one core:




FIG. 1

Ex-1006, Fig. 1 (annotated).

82. Machnicki's first core is also operable at a first voltage and first frequency. Machnicki discloses that "both clock frequencies and voltages are configuration parameters that may be set for each of the performance domains (e.g. domains 1 through m)." Ex-1006, 8:1-4. As shown in Table 1, each domain has n operating states (e.g., at least state 1, state 2, state 3, ... state n), and each operating state has a voltage, V, and a frequency, f:

TABLE I

	Domain 1	Domain 2	...	Domain m
state 1	$f_{1,1}; V_{1,1}$	$f_{2,1}; V_{2,1}$	...	$f_{m,1}; V_{m,1}$
state 2	$f_{1,2}; V_{1,2}$	$f_{2,2}; V_{2,2}$	...	$f_{m,2}; V_{m,2}$
state 3	$f_{1,3}; V_{1,3}$	$f_{2,3}; V_{2,3}$	...	$f_{m,3}; V_{m,3}$
...	...	...	...	...
state n	$f_{1,n}; V_{1,n}$	$f_{2,n}; V_{2,n}$	...	$f_{m,n}; V_{m,n}$

Id. at 8:8-17 (annotated).

83. Because each domain has n pairs of operable voltages and frequencies, a POSITA would have understood that processor 160A (that includes the "first core") within a first domain (performance domain 120A) is operable at a first voltage and a first frequency (e.g., as shown above, $f_{1,1}$; $V_{1,1}$).

- c. Element 1[b]: a second core operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency;**

84. Machnicki discloses Element 1[b].

85. Machnicki's SoC 100 includes "processor 160F." Ex-1006, 4:16-25. Processor 160F has at least one core, and that core is a "second core." Figure 4, annotated below, illustrates processor 160F (blue), including the second core:




FIG. 1

Id., Fig. 4 (annotated).

86. Machnicki's second core is also operable at a second voltage independent of the first voltage. Both "clock frequencies and voltages are configuration parameters that may be set for each of the performance domains (e.g. domains 1 through m)," including a "performance domain 120F" containing processor 160F (that includes the "second core" of processor 160F). Ex-1006, 8:1-4. As shown in Table 1, a second domain containing processor 160F has n operating states, and each operating state has a voltage, V, and a frequency, f:

TABLE I

	Domain 1	Domain 2	...	Domain m
state 1	$f_{1,1}; V_{1,1}$	$f_{2,1}; V_{2,1}$	...	$f_{m,1}; V_{m,1}$
state 2	$f_{1,2}; V_{1,2}$	$f_{2,2}; V_{2,2}$	...	$f_{m,2}; V_{m,2}$
state 3	$f_{1,3}; V_{1,3}$	$f_{2,3}; V_{2,3}$	...	$f_{m,3}; V_{m,3}$
...	...	...	...	...
state n	$f_{1,n}; V_{1,n}$	$f_{2,n}; V_{2,n}$	...	$f_{m,n}; V_{m,n}$

Id. at 8:8-17 (annotated).

87. Furthermore, Machnicki explains that “by using a performance state table such as the one illustrated in Table I [], each performance domain may have its parameters changed *independently*.” Ex-1006, 8:37-43. Thus, because voltage is a configuration parameter that may be set for a performance domain, the second core’s voltage can be set separately and is operable independent of the first voltage.

88. Moreover, Machnicki’s second core is operable at a second frequency different from the first frequency. As noted above, each performance domain, including frequency, “may have its parameters changed independently.” Ex-1006, 8:37-43. Machnicki discloses that “in the case of clock frequencies, *the frequencies in each frequency domain for any given state may be different from each other*, and may be other than a multiple of each other.” Id. at 8:37-43 (emphasis added). A POSITA would have understood that processor 160F (that includes the “second core”), which is within a second performance domain, is operable at a voltage

independent from the first and second frequency (e.g., as shown above, $f_{2,2}$; $V_{2,2}$)

different from that of the first core.

- d. Element 1[c]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and**

89. Machnicki discloses Element 1[c].

90. Machnicki discloses the SoC's "performance domain 120B includes graphics subsystem 160B." Ex-1006, 4:9-25. The graphics subsystem 160B is "graphics processing circuitry" " because it is "involved in the display of images on a display device" and including "rendering hardware, refresh (of the display device) hardware, video encoders and/or decoders, video compression and decompression units, etc." Ex-1006, 7:1-7. Figure 1, annotated below, illustrates Machnicki's graphics subsystem 160B (purple):




FIG. 1

Ex-1006, Fig. 1 (annotated).

91. Machnicki also discloses that “[c]omponents within integrated circuit 100 may be *coupled to each other using any suitable bus* and/or interface mechanism.” Ex-1006, 4:31-33. Based on this disclosure, a POSITA would have known that the graphics subsystem 160B (the “graphics processing circuitry”) is coupled to the processors 160A and 160F (the “plurality of cores”), by e.g., using “any suitable bus.” See V.D. For example, graphics subsystem 160B (the “graphics processing circuitry”) would have been connected to processor 160A (the “first core”) and processor 160F (the “second core”) via a bus or other interface mechanism.

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

92. Moreover, the graphics processing circuitry performs graphic operations. Graphics unit or subsystem 160B may include “circuitry involved in the display of images on a display device,” and “[t]he produced images may be static images, or may be part of a video.” Ex-1006, 7:1-4. Graphics unit 160B may also include “rendering hardware, refresh (of the display device) hardware, video encoders and/or decoders, video compression and decompression units, etc.” *Id.* at 7:4-7.

93. Furthermore, Machnicki’s graphics processing circuitry is independently operable at a third frequency different from the first frequency and the second frequency. Machnicki discloses that “clock frequencies and voltages are configuration parameters that may be set for each of the performance domains (e.g. domains 1 through m),” including performance domain 120B containing graphics subsystem 160B. Ex-1006, 8:1-4. Furthermore, Machnicki explains that “by using a performance state table such as the one illustrated in Table I [], each performance domain may have its parameters changed independently.” *Id.* at 8:37-43. As shown in Table 1, a third domain containing graphics subsystem 160B has n operating states, and each operating state has a voltage, V, and a frequency, f:

TABLE I

	Domain 1	Domain 2	...	Domain m
state 1	$f_{1,1}; V_{1,1}$	$f_{2,1}; V_{2,1}$	...	$f_{m,1}; V_{m,1}$
state 2	$f_{1,2}; V_{1,2}$	$f_{2,2}; V_{2,2}$	...	$f_{m,2}; V_{m,2}$
state 3	$f_{1,3}; V_{1,3}$	$f_{2,3}; V_{2,3}$	...	$f_{m,3}; V_{m,3}$
...	...	...	...	...
state n	$f_{1,n}; V_{1,n}$	$f_{2,n}; V_{2,n}$	...	$f_{m,n}; V_{m,n}$

Id. at 8:8-17 (annotated).

94. As with the first and second cores, Machnicki also discloses that “in the case of clock frequencies, *the frequencies in each frequency domain for any given state may be different from each other*, and may be other than a multiple of each other.” Ex-1006, 8:37-43. A POSITA would have understood that graphics processor 160B is operable at a voltage and frequency (e.g., as shown above, $f_{3,3}$; $V_{3,3}$) different from that of the first core and second core.

- e. **Element 1[d]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first, second, and third frequencies.**

95. Machnicki discloses Element 1[d].

96. Machnicki discloses the SoC’s “performance domain 120E includes memory control circuit 160E.” Ex-1006, 4:9-25. Figure 1, annotated below, illustrates Machnicki’s memory control circuit 160E (green):




FIG. 1

Ex-1006, Fig. 1 (annotated).

97. Memory control circuit 160E acts to connect the other components within the integrated circuit to memory “*using any suitable bus* and/or interface mechanism.” Ex-1006, 4:31-33. The memory control circuit 160E comprising a bus is an “interconnect” because it is “configured to communicate on the memory interfaces, to queue memory requests from other components integrated circuit 100, and/or to communicate with the other components to fulfill memory read and write operations,” thus connecting the components. *Id.* at 7:30-34. As discussed in Section V.D, interconnects and buses are commonly used in system components to interconnect the system to components such as a graphic processing circuitry, memory elements, system components, resources and custom circuitry. Indeed, memory control circuit 160E sits between each component within the integrated

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

circuit and memory. For example, a POSITA would have understood that, in the configuration of Figure 1, a suitable bus would have linked each component to memory control circuit 160E. Ex-1006, 4:31-33. Memory control circuit 160E in turn links each component to memory. Id. at 7:30-34. By acting as the bridge between each component within the integrated circuit and memory, the memory control circuit 160E is an interconnect as described by the '588 Patent. See Section V.C-D.

98. The memory control circuit 160E also couples the graphics processing circuitry to a memory device. Machnicki discloses that “memory controller 160E may be configured to interface to one or more memory modules (e.g., single inline memory modules (SIMMs), dual inline memory modules (DIMMs), etc.) that include one or more of the memories.” Ex-1006, 7:21-30. Additionally, memory controller 160E “may be configured to communicate on the memory interfaces, to queue memory requests from other components in integrated circuit 100, and/or to communicate with the other components to fulfill memory read and write operations.” Id. at 7:30-34. As discussed in Section V.D, it was well known that interconnects provided communication, such as Machnicki’s memory controller 160E. A POSITA would have known that, to communicate with other components within the integrated circuit and to then fulfill read and write operations based on those communications, the memory control 160E would have been coupled to each

component, including the graphics processing circuitry, and configured to couple those components to memory.

99. Moreover, Machnicki's interconnect, is operable at a fourth frequency different than the first, second, and third frequencies. As shown in Table 1, a fourth domain containing memory control circuit 160E has n operating states, and each operating state has a voltage, V, and a frequency, f:

TABLE I

	Domain 1	Domain 2	...	Domain m
state 1	$f_{1,1}; V_{1,1}$	$f_{2,1}; V_{2,1}$	...	$f_{m,1}; V_{m,1}$
state 2	$f_{1,2}; V_{1,2}$	$f_{2,2}; V_{2,2}$	...	$f_{m,2}; V_{m,2}$
state 3	$f_{1,3}; V_{1,3}$	$f_{2,3}; V_{2,3}$	...	$f_{m,3}; V_{m,3}$
...	...	...	...	...
state n	$f_{1,n}; V_{1,n}$	$f_{2,n}; V_{2,n}$	...	$f_{m,n}; V_{m,n}$

Id. at 8:8-17 (annotated).

100. Machnicki discloses that “in the case of clock frequencies, the *frequencies in each frequency domain for any given state may be different from each other*, and may be other than a multiple of each other.” Ex-1006, 8:37-43. Because the memory control circuit 160E is in its own domain, a POSITA would have known the memory control circuit 160E is operable at a fourth frequency (e.g., as shown above, $f_{4,4}; V_{4,4}$) different than the cores and graphics processing circuitry's frequencies.

2. **Dependent Claim 2: The processor of claim 1, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.**

101. Machnicki discloses Claim 2.

102. Machnicki discloses that “[i]ntegrated circuit 100 also includes power management unit (PMU) 140, which in turn includes one or more performance configuration registers 150.” Ex-1006, 4:25-30. Figure 4, annotated below, illustrates power management unit (PMU) 140 (yellow):




FIG. 1

Id., Fig. 4 (annotated).

103. PMU 140 “may be a logic circuit or the like.” Ex-1006, 5:62-63. PMU 140 is coupled to “clock and/or voltage control unit (control unit) 130” and “may be configured to establish a corresponding performance state for each performance domain, in which “each performance domain may have its parameters changed *independently*.” Id. at 4:28-4:30, 4:56-60, 8:37-43. The PMU also “may be configured to control transitions between performance states in each performance domain” Id. at 4:56-60. Machnicki’s power management unit (PMU) 140 is “first power management circuitry and/or logic.” Specifically, it is power management “circuitry.”

104. Moreover, the power management unit (PMU) 140 causes the first core to operate at the first frequency and the second core to independently operate at the second frequency. For example, “[i]n the case where performance states or levels are reached by making voltage and/or clock frequency changes, PMU 140 may communicate new voltage and/or frequency settings to the control unit 130.” Ex-1006, 6:14-17. As discussed, the PMU is “configured to establish a corresponding performance state for each performance domain” in which “the frequencies in each frequency domain for any given state may be different from each other.” Id. at 4:56-60, 8:37-43. Because Machnicki states that the frequencies in each frequency domain for any given state may be different from each other, and PMU 140 is responsible for establishing each performance state within each domain, a POSITA

would have understood the PMU causes the first core to operate at the first frequency and the second core to independently operate at the second frequency.

3. Dependent Claim 4: The processor of claim 2, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency.

105. Machnicki discloses Claim 4.

106. Machnicki discloses that if its software communicates with the graphics subsystem 160B, it calls “graphic driver 420.” Ex-1006, 10:23-26. This driver communicates with “PMU driver 460,” which “may be used to interact with PMU 140.” Id. at 10:22-27. PMU driver 460 is “second power management circuitry and/or logic.” Specifically, these drivers are second power management “logic.”

107. Moreover, PMU driver 460 causes the graphics processing circuitry to operate at the third frequency. Machnicki discloses the PMU driver 460 may “monitor the activity in the components driven by the drivers 420-450 in a device activity table or the like,” and based on these monitored activities, the PMU driver 460 “may select the performance state for each performance domain to be used in the various performance states.” Id. at 10:27-37. Machnicki also discloses that “performance states or levels are reached by making voltage and/or clock frequency changes.” Id. at 6:14-17. Thus, a POSITA would have known that Machnicki uses PMU driver 460 to monitor activity in graphics subsystem 160B driven by graphics

driver 420, and based on the activities being monitored, PMU driver 460 may select the performance state (in terms of voltage/frequency changes) for graphics subsystem 160B. In other words, PMU driver 460 (second power management “circuitry and/or logic”) causes the graphics processing circuitry (160B) to operate at the third frequency.

4. Independent Claim 8

a. Element 8 ``` [pre]: A system comprising: ```

108. Machnicki discloses Claim 8’s preamble to the extent it is limiting.

109. Machnicki explains the “system includes an integrated circuit (IC) 100 forming a system-on-a-chip (SoC) or application-specific integrated circuit (ASIC) coupled to power supply 110.” Ex-1006, 4:9-25. The SoC includes performance domains, each with components, including “[P]erformance domain 120A includes two components processor 160A [] performance domain 120B includes graphics subsystem 160B, [] performance domain 120E includes memory control circuit 160E, and performance domain 120F includes optional second processor 160F.” *Id.*

b. Element 8[a]: a system memory;

110. Machnicki discloses Element 8[a].

111. Machnicki discloses that its system includes “memory 530,” which may be “any type of memory” such as “SRAM, nonvolatile RAM (NVRAM, such as ‘flash’ memory), dynamic RAM (DRAM) such as synchronous DRAM (SDRAM),

double data rate (DDR, DDR2, DDR3, etc.) SDRAM, RAMBUS DRAM, etc.,
 and/or any type of tangible memory.” Ex-1006, 11:20-29. Figure 5, annotated below,
 illustrates Machnicki’s memory 530 (maroon):




FIG. 5

Id., Fig. 5 (annotated).

c. Element 8[b]: a plurality of cores coupled to the system memory;

112. For the reasons described above for Element 1[a], Machnicki discloses a plurality of cores within an integrated circuit.

113. As shown in Figure 5, annotated below, the integrated circuit (green), which may be the integrated circuit 100 discussed above, and therefore the cores of integrated circuit 100, are coupled to the system memory (maroon):




FIG. 5

Id., Fig. 5 (annotated). Machnicki explains that “[i]ntegrated circuit 520 may include one or more instances of elements within integrated circuit 100.” *Id.* at 10:55-57. A POSITA would have understood that system 500 includes elements within integrated circuit 100, such as the cores.

- d. **Element 8[c]: a first core of the plurality of cores operable at a first voltage and at a first frequency, and a second core of the plurality of cores operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency;**

114. For the reasons described above for Elements 1[a] and 1[b], Machnicki discloses Element 8[c].

- e. **Element 8[d]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and**

115. For the reasons described above for Element 1[c], Machnicki discloses Element 8[d].

- f. **Element 8[e]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first and second frequencies.**

116. For the reasons described above for Element 1[d], Machnicki discloses Element 8[e].

- 5. **Dependent Claim 9: The system of claim 8, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.**

117. For the reasons described above for Claim 2, Machnicki discloses Claim 9.

- 6. **Dependent Claim 11: The system of claim 9, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency.**

118. For the reasons described above for Claim 4, Machnicki discloses Claim 11.

7. **Dependent Claim 15: The system of claim 8, further comprising: a chipset coupled to the plurality of cores, the chip set comprising a plurality of interfaces to couple the cores to a plurality of I/O devices.**

119. Machnicki discloses Claim 15.

120. Machnicki discloses that “performance domain 120D may include two or more networking or other peripherals 160D.” Ex-1006, 4:16-20. Networking peripherals and other peripherals 160D “may include a variety of circuitry,” including “a media access controller (MAC) unit for the supported network, as well as physical layer circuitry.” Id. at 7:13-15. Other peripherals may also include “any other desired peripherals, and/or peripheral interface controllers configured to control off-chip peripheral interfaces such as Peripheral Component Interconnect (PCI), PCI express (PCIe), firewire, Universal Serial Bus (USB), etc.” Id. at 7:13-20. A POSITA would have known that the peripherals are a chipset comprising a plurality of interfaces so that the input/output devices described above can connect to Machnicki’s SoC. Moreover, a POSITA would have understood that a “peripheral interface controller” would be implemented in a chip, such as a USB controller for USB or a PCI root hub for PCI. Figure 1, annotated below, illustrates, the peripherals are coupled to the cores because “[c]omponents within integrated circuit 100 may be coupled to each other using any suitable bus and/or interface mechanism”:




FIG. 1

Id., Fig. 1 (annotated), 4:31-33.

121. Machnicki also discloses “[i]ntegrated circuit 520 may include one or more instances of elements within integrated circuit 100,” and include “various additional circuitry such as ... video and/or audio processing circuitries, on-chip peripherals and/or peripheral interfaces to couple to off-chip peripherals, etc.” Ex-1006, 10:53-64. Figure 5, annotated below, illustrates Machnicki’s integrated circuit coupled to one or more “peripherals 540.”




FIG. 5

Id., Fig. 5 (annotated).

122. A POSITA would have understood that to couple the integrated circuit to multiple peripherals 540, the circuit would include a chipset with a plurality of interfaces. Moreover, a POSITA would have understood that the chipset would be coupled to the plurality of cores because peripherals 540 are connected to the integrated circuit 520 containing the cores. Finally, peripherals 540 include a plurality of I/O devices, including “a display screen, including touch display screens or multitouch display screens, keyboard or other input devices, microphones, speakers, etc.” *Id.* at 11:10-15.

8. Dependent Claim 16: The system of claim 15, wherein one of the I/O devices comprises an audio generation device.

123. Machnicki discloses Claim 16.

124. Machnicki discloses that the peripherals 540 include a plurality of I/O devices, including “speakers, etc.” Ex-1006, 11:10-15. A POSITA would have known that speakers are an “audio generation device.”

9. Dependent Claim 17: The system of claim 16, wherein one of the I/O devices comprises a video rendering device.

125. Machnicki discloses Claim 17.

126. Machnicki discloses that peripherals 540 include a plurality of I/O devices, including “a display screen, including touch display screens or multitouch display screens” Ex-1006, 11:10-15. A POSITA would have known that a display screen is a “a video rendering device.”

10. Dependent Claim 18: The system of claim 17, wherein one of the I/O devices comprises a user input device.

127. Machnicki discloses Claim 18.

128. Machnicki discloses peripherals 540 include a plurality of I/O devices, including a “keyboard or other input devices” Ex-1006, 11:10-15. A POSITA would have known that a keyboard is a “user input device.”

11. Dependent Claim 19: The system of claim 18, wherein one of the I/O devices comprises a network communication device.

129. Machnicki discloses Claim 19.

130. Machnicki discloses “networking peripherals” that may include “a media access controller (MAC) unit for the supported network, as well as physical layer circuitry” ” and “devices for various types of wireless communication, such as

Wi-fi, Bluetooth, cellular, global positioning system, etc.” Ex-1006, 7:11-20, 11:8-

10. A POSITA would have known that a network peripheral is a “network communication device.”

12. Independent Claim 20

131. Claim 20 is nearly identical to Claim 1, except for being written in “means-plus-function” form. A comparison between Claim 1 and Claim 20 is provided below for reference:

~~1. A processor~~20. An apparatus comprising:

~~a plurality of cores including~~means for operating a first core ~~operable~~ at a first voltage and at a first frequency, ~~and while concurrently operating~~ a second core ~~operable~~ at a second voltage independent of the first voltage and at a second frequency different from the first frequency;

~~graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry~~means for independently ~~operable~~operating a graphics processor at a third frequency different from the first frequency and the second frequency; and

~~an~~ interconnect means to couple the graphics ~~processing circuitry~~processor to a memory device, the interconnect means operable at a fourth frequency different from the first, second, and third frequencies.

132. The “functions” recited in Claim 20 are the same as what is recited in Claim 1, so the analysis below cross-references the discussion of the prior art analysis for Claim 1.

a. Element 20[pre]: An apparatus comprising:

133. Machnicki teaches the preamble of Claim 20 to the extent it is limiting. Machnicki teaches an “integrated circuit (IC) 100 forming a system-on-a-chip (SoC)

or application-specific integrated circuit (ASIC) coupled to power supply 110,
 which is an apparatus. Ex-1006, 4:9-12.

- b. Element 20[a]: means for operating a first core at a first voltage and at a first frequency while concurrently operating a second core at a second voltage independent of the first voltage and at a second frequency different from the first frequency;**

134. As described in Section VIII.A.1.b-d, above, the corresponding structure is a “power control unit 355” and equivalents, performing the function recited in the remainder of this claim element. Notably, these functions are the same as those discussed above with respect to Elements 1[a], 1[b], and 1[c]. Machnicki discloses Element 20[a] because it discloses equivalent structure performing the recited functions.

135. Specifically, Machnicki’s equivalent structure is its “power management unit (PMU) 140.” Ex-1006, 4:25-30. Machnicki’s PMU is discussed in Sections XII.A.1.b-d, above. Moreover, as set forth in those sections, PMU 140 provides the capability to operate the first core, second core, and GPU at independent and distinct voltages and frequencies, respectively.

- c. Element 20[b]: means for independently operating a graphics processor at a third frequency different from the first frequency and the second frequency**

136. As described in Section VIII.A.3, above, the corresponding structure is a “power control 355” and equivalents, performing the function recited in the

remainder of this claim element. Notably, this function is the same as those discussed above with respect to Claim 4. PMU driver 460 may be used to interact with PMU 140 and causes the graphics processing circuitry to operate at the third frequency. Ex-1006, 10:22-27. Machnicki discloses Element 20[b] because it discloses equivalent structure performing the recited functions.

- d. Element 20[c]: interconnect means to couple the graphics processor to a memory device, the interconnect means operable at a fourth frequency different from the first, second, and third frequencies.**

137. As described in Section VIII.A.1.e, above, the corresponding structure is a “interconnect 315” and equivalents, performing the function recited in the remainder of this claim element. Notably, this function is the same as those discussed above with respect to Element 1[d]. Machnicki discloses Element 20[c] because it discloses equivalent structure performing the recited functions.

138. Specifically, Machnicki’s equivalent structure is “memory control circuit 160E” that acts to connect the processor’s various components. Ex-1006, 7:30-34. Memory control circuit 160E is discussed in Sections XII.A.1.e, above. Moreover, as set forth in those sections, memory control circuit 160E performs the recited function.

B. Ground 2: Claims 3, 7, 10, and 14 are rendered obvious by Machnicki and Aguilar.**1. A POSITA would have been motivated to combine Machnicki and Aguilar's teachings and would have had a reasonable expectation of success.**

139. In addition to the reasons set forth below with respect to specific claim elements, a POSITA would have been motivated to combine the teachings of Machnicki and Aguilar and would have had a reasonable expectation of success in doing so because each relates to the same well-known technologies. See Section V.

140. Both relate to and describe methods for optimizing performance between multiple components within computer systems. Ex-1006 (Machnicki), Abstract; Ex-1007 (Aguilar), Abstract. For example, Machnicki discloses “[s]ystems and methods for coordinating performance parameters in multiple domains” Ex-1006, Abstract. The system manages performance by “causing a first parameter of a first circuit serving the first domain to be modified to a first modified parameter based on [a] request, and causing a second parameter of a second circuit serving the second domain to be modified to a second modified parameter based on [a] request.” *Id.* The parameter may include “a clock frequency,” and in other cases, “a parameter may include a voltage.” *Id.* Similarly, Aguilar discloses “[a] computer implemented method, data processing system, and computer usable code ... for optimizing thermal performance of a computer system” by requesting “a thermal index” for “*each of the set of system resources* to form a set of thermal indexes.” *Id.* (emphasis

added). Then, “[a]n action from a set of actions is identified to reduce resource utilization” for each resource. *Id.*

141. A POSITA would have been motivated to combine Machnicki and Aguilar because they would have recognized a benefit of incorporating Aguilar’s teachings of thermal indexes into a system like Machnicki to improve temperature control and resource utilization throughout a computer system. Aguilar discloses that “identification of a set of system resources” associated with a computer system “is performed and a thermal index is requested for *each [] set* of system resources.” Ex-1007, Abstract (emphasis added). A “set of actions is identified to reduce resource utilization,” and the actions are “implemented to reduce a thermal state of the computer system.” *Id.* Thus, a POSITA would have known that by incorporating Aguilar’s teachings into Machnicki, temperature control and overall utilization of the computer system could be improved.

142. A POSITA would also have had a reasonable expectation of success in combining the teachings of Machnicki and Aguilar. A POSITA would have understood that these references disclose interrelated teachings based on routine technologies that would have been amendable to various well-understood and predictable combinations, as explained in further detail below.

- 2. Dependent Claim 3: The processor of claim 1, further comprising: a non-core power domain comprising a cache shared by the plurality of cores.**

143. Machnicki and Aguilar teach Claim 3. Machnicki discloses non-core power components, which form a “non-core power domain,” including “graphics 160B” (purple), “memory controller 160E” (green), and “PMU 140” (yellow), as illustrated in Figure 1, annotated below:




FIG. 1

Ex-1006, Fig. 1 (annotated).

144. A POSITA would have understood that Machnicki’s memory control 160E contained a cache, specifically an on-chip or “Level 1” cache, because such “L1” caches were common components of high-performance processors at the time of the ’588 Patent. As discussed in Section V.E., it is common for integrated circuits such as Machnicki’s to include a shared cache. Machnicki teaches that memory

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

within its system may be a cache. For example, Machnicki discloses that “configuration parameters may be component specific,” and may depend on “cache sizes in various caches ...” Ex-1006, 5:18-20. Machnicki also teaches, “[r]educing cache size may reduce power consumption because a portion of the cache may not need to be accessed, and even more reduction may be realized if the unused portion may be turned off.” Id. at 5:43-46.

145. This cache is shared by the plurality of cores. And as discussed in Section V.E, it was common for multicore processors to include a multi-level cache that includes Level 1 (L1) caches and Level 2 (L2) caches. In fact, the i486 processor included an L1 cache in 1993. Machnicki’s memory controller 160E “may be configured to communicate on the memory interfaces, to queue memory requests from other components in integrated circuit 100, and/or to communicate with the other components to fulfill memory read and write operations.” Ex-1006, 7:30-34. Through the memory controller 160E, the components within the circuit, including the cores, communicate requests to the system memory and fulfill memory read and write operations, thus sharing the memory.

146. A POSITA would have understood that one design choice for memory controller 160E would have included a cache because caches were common components of nearly all memory systems at the time of the ’588 Patent. Indeed, Aguilar discloses such an arrangement in which “memory may be ... a cache such as

found in north bridge and memory controller 202 in Fig. 2.” Ex-1007, 4:57-60. A POSITA would have been motivated to implement Machnicki’s memory controller with a cache memory, for example as taught by Aguilar, because including a cache makes retrieving data from a computer’s memory more efficient. Data can be retrieved easily from cache memory, and data stored in caches is typically more readily available to processors and other system components than data stored in system’s main memory. Moreover, [c]omponents within integrated circuit 100 may be coupled to each other using any suitable bus and/or interface mechanism.” Ex-1006, 4:31-33.

- 3. Dependent Claim 7: The processor of claim 3, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively.**

147. For the reasons described above for Elements 1[c] and 1[d], Machnicki and Aguilar disclose Claim 7. A POSITA would have known Machnicki’s power management circuitry and/or logic causes the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores because the Graphics 160B and Memory Control 160E, which are part of the non-core power domain, operate at different frequencies than the cores as explained in Sections VIII.A.1.d-e.

- 4. Dependent Claim 10: The system of claim 8, further comprising: a non-core power domain comprising a cache shared by the plurality of cores.**

148. See discussion in Section VIII.A.4, regarding Claim 8. For the reasons described above for Claim 3, Machnicki and Aguilar disclose Claim 10.

- 5. Dependent Claim 14: The system of claim 10, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively.**

149. For the reasons described above for Claim 7, Machnicki and Aguilar disclose Claim 14.

C. Ground 3: Claims 5-6, and 12-13 are rendered obvious by Machnicki and Shakkarwar.

- 1. A POSITA would have been motivated to combine Machnicki and Shakkarwar's teachings and would have had a reasonable expectation of success.**

150. In addition to the reasons set forth below with respect to specific claim elements, a POSITA would have been motivated to combine the teachings of Machnicki and Shakkarwar and would have had a reasonable expectation of success in doing so because each relates to the same well-known technologies.

151. Both relate to and describe similar methods for optimizing performance between multiple components within a computer system. Ex 1006 (Machnicki), Abstract; Ex 1008 (Shakkarwar), Abstract. For example, Machnicki discloses “[s]ystems and methods for coordinating performance parameters in multiple

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

domains” Ex-1006, Abstract. The system manages performance by “causing a first parameter of a first circuit serving the first domain to be modified to a first modified parameter based on [a] request, and causing a second parameter of a second circuit serving the second domain to be modified to a second modified parameter based on [a] request.” *Id.* The parameter may include “a clock frequency,” and in other cases, “a parameter may include a voltage.” *Id.* Similarly, Shakkarwar relates to the “efficient use of electronic devices and, more particularly, to systems and methods for power management of electronic devices.” Ex-1008, 1:8-10. Shakkarwar discloses “[s]ystems and methods for dynamic power management of electronic devices” so that “power consumption of the system may be dynamically managed while providing sufficient bandwidth for the system.” Ex-1008, Abstract. Shakkarwar dynamically manages power by monitoring a component’s “bandwidth ... for increasing or decreasing communication requirements.” *Id.* at 6:41-43. The system then detects “a change in processing condition[s]” for a component, and upon detecting a change, “determines bandwidth 403 for an associated data connection based on the active components.” *Id.* at 6:63-7:7. The data rate “of a data connection is [then] adjusted,” including by adjusting a clock frequency or voltage, to reach optimal performance levels. *Id.* at 7:26-29. In this manner, “power consumption associated with communicating and processing information via data buses may be

optimized while satisfying bandwidth requirements for the computer system.” Id. at 3:17-20; *see also id.* at 2:12-15, Cl. 21.

152. A POSITA would have been motivated to combine the teachings of Machnicki and Shakkarwar because they would have recognized a benefit of implementing a look-up table for dynamically managing power, such as taught by Shakkarwar, to optimize performance of components within a processor as taught by Machnicki. As noted above, Shakkarwar teaches that an advantage to its system is that it dynamically “manag[es] power consumption of electronic devices while [also] satisfying bandwidth requirements for communicating and/or processing information.” Id. at 2:34-39.


153. A POSITA would also have had a reasonable expectation of success in combining the teachings of Machnicki and Shakkarwar. A POSITA would have understood that these references disclose interrelated teachings based on routine technologies that would have been amendable to various well-understood and predictable combinations, as explained in further detail below.

2. Dependent Claim 5: The processor of claim 4, wherein the second power management circuitry and/or logic includes a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.

154. Machnicki and Shakkarwar teach Claim 5. As explained above in Section VIII.A.3, Machnicki discloses a second power management circuitry and/or

logic. Machnicki in view of Shakkarwar teaches that this power management circuitry would have been implemented to include a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect. *See* section VIII.C.1 (discussing Shakkarwar's teachings).

155. Shakkarwar's dynamic power management system includes several components, including “a data controller 302” (green) connected to a “graphics controller 305” (purple) through “data connection 314” (green), as illustrated in Figure 3:




Id., Fig. 3 (annotated).

156. A POSITA would have understood that the data controller 302 and data connection 314 are an “interconnect,” and the “graphics controller 305” is “graphic processing circuitry.” For example, data controller 302 connects graphics controller 305 and other components within Shakkarwar to each other and to system memory 303. Ex-1008, 5:10-15, Fig. 4. Graphics controller 305 includes a “3-D engine 305e” and performs “3-D graphics operations” *Id.* at 6:7-16.

157. Shakkarwar also maps selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect. Shakkarwar discloses at step 401, its system monitors “bandwidth for a component,” such as the interconnect or graphics processing circuitry, for “increasing or decreasing communication requirements.” *Id.* at 6:41-43. For example, “when an application that is known to use 3-D graphics is opened, system software can determine if the bus rate between the data controller [the ‘interconnect’] and the graphics controller [the ‘graphics processing circuitry’] is high enough to support the graphics requirements.” *Id.* at 6:49-52.

158. At step 402, the system detects “a change in processing condition[s]” for a specific component, and upon detecting a change, it “determines bandwidth 403 for an associated data connection ***based on the active components.***” Ex-1008, 6:63-7:7. To determine the appropriate bandwidth for a data connection, Shakkarwar

utilizes “*a look-up table [that] may include a list of clock speeds, and/or voltage levels required to process information for specific devices and/or groups of devices.*” *Id.* at 7:7-10. Figure 4 illustrates this process:



159. *Id.*, Fig. 4. Accordingly, Shakkarwar teaches determining the appropriate bandwidth for a component (such Graphics controller 305) utilizing a look-up table that includes a list of appropriate clock speeds. It also considers other active components, such as the data controller 302 (the “interconnect”) and its clock speeds. As such a POSITA would have understood that Shakkarwar teaches to map selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect, using look-up tables.

160. A POSITA would have been motivated to implement these teachings of Shakkarwar in Machnicki for the reasons described above in Section VIII.C.1, including because Shakkarwar teaches, for example, that an advantage to its system is that it dynamically “manag[es] power consumption of electronic devices while satisfying bandwidth requirements for communicating and/or processing information.” Ex-1008, 2:34-39. A POSITA would have been motivated to include this advantage in Machnicki to improve responsiveness of the system and increase the agility of the system to respond to dynamically changing conditions, including the changing bandwidth requirements of Machnicki’s components during various activity states.

- 3. Dependent Claim 6: The processor of claim 5, wherein the second power management circuitry and/or logic is to select the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.**

161. Machnicki and Shakkarwar disclose Claim 6. As explained above for Claim 5, Machnicki and Shakkarwar disclose second power management circuitry that determines bandwidth for an associated data connection based on the other active components by using a look-up table of clock speeds required to process information for specific devices. *See* Section VIII.C.1 (discussing Shakkarwar’s teachings).

162. Shakkarwar further discloses that “[a]t step 405, the data rate of a data connection is adjusted. In one embodiment, the data rate can be adjusted by adjusting operating of a clock generation device.” Ex-1008, 7:26-29. Figure 4 illustrates adjusting the data rate of a data connection:




FIG. 4

Id., Fig. 4.

163. Shakkarwar also specifically discloses selecting “the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry

at the third frequency.” In particular, Shakkarwar discloses that “when an application that is known to use 3-D graphics is opened, system software can determine *if the bus rate between the data controller [the “interconnect”] and the graphics controller [the “graphics processing circuitry”] is high enough to support the graphics requirements.*” Ex-1008, 6:49-52. A POSITA would have understood that by applying these teachings from Shakkarwar to Machnicki, the fourth frequency for the interconnect would be selected based on detecting the frequency of the graphics processing circuitry.

4. **Dependent Claim 12: The system of claim 11, wherein the second power management circuitry and/or logic includes a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.**

164. See discussion in Sections VIII.A.4 and VIII.A.6, regarding Claims 8 and 11. For the reasons described above for Claim 5, Machnicki and Shakkarwar disclose Claim 12.

5. **Dependent Claim 13: The system of claim 12, wherein the second power management circuitry and/or logic is to select the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.**

165. See discussion in Sections VIII.A.4 and VIII.C.4, regarding Claims 8 and 12. For the reasons described above for Claim 6, Machnicki and Shakkarwar disclose Claim 13.


D. Ground 4: Claims 1-2, 4, 8-9, 11 and 20 are rendered obvious by Cho and Gold.

1. A POSITA would have been motivated to combine Cho and Gold's teachings and would have had a reasonable expectation of success.

166. In addition to the reasons set forth below with respect to specific claim elements, a POSITA would have been motivated to combine the teachings of Cho and Gold (Cho-Gold) and would have had a reasonable expectation of success in doing so because each addresses similar issues corresponding to the same well-known technologies.

167. Cho and Gold share a common inventor, Hanwoo Cho, and were developed contemporaneously for the same assignee, Global Foundries. Both relate to and describe similar methods for optimizing the performance of components within a computer system by controlling frequency. Ex-1005 (Cho), Abstract; Ex-1009 (Gold), Abstract. Additionally, each reference specifically teaches the same direct relationship between temperature and frequency; namely a lower operating frequency will result in a lower operating temperature. Ex-1005 ¶¶24, 33, 36; Ex-1009, Fig. 7.

168. Cho teaches an SOC device 400 that has at least four different functional units: A first core, a second core, a GPU, and a Northbridge. As shown in Figure 4, each of these four functional units are driven by a distinct clock frequency, clk1 through clk4, annotated in the figure in red, blue, green and purple:



169. Regarding these four clock frequencies, Cho teaches that they are changed “independently of one another, based on signals received from control logic 305.” Ex-1005 ¶35. And, “a change in the clock frequency for a clock signal provided to one functional unit need not affect the frequencies of the respective clock signals provided to other ones of the functional units.” *Id.* ¶22. Indeed, Cho’s “clock control unit” that is coupled to the four functional units “provide[s] *separate* clock signals to each of the functional units.” *Id.* ¶35.

170. The reason that Cho does this is to individually and separately control the temperature of each of the 4 functional units. For example, if a particular functional unit is operating at a maximum temperature, its operating frequency may be reduced to lower the operating temperature of that unit. Ex-1005 ¶¶30, 34, 39-40.

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

Cho further teaches that a functional unit may have four different monitored temperature states: a maximum temperature, a minimum temperature, a first intermediate temperature, and a second intermediate temperature. Id. ¶24. Cho teaches the direct relationship between temperature and operating frequency (lower operating frequency will result in lower operating temperature), but does not specifically identify four explicit frequencies that correspond to each of a maximum temperature, a minimum temperature, a first intermediate temperature, and a second intermediate temperature. One of skill in the art, however, would have understood that, due to the direct relationship between operating temperature and operating frequency, the four corresponding frequencies would have been different from each other.

171. This understanding is further made clear by Gold. Gold teaches “a hardware throttling control system that aggressively reduces or even stops the clock of a processing core once thermal limitations are exceeded.” Ex-1005 ¶53. Gold further specifically discloses four specific frequencies and corresponding power dissipations:


500


<i>P-State</i>	<i>Power Usage</i>	<i>Frequency</i>
<i>P Max</i>	19W	2.90 GHz
<i>P 0</i>	15W	2.60 GHz
<i>P 1</i>	12W	2.28 GHz
<i>P 2</i>	9W	1.73 GHz

FIG. 5A

172. As Gold teaches, the four specific frequencies may be reduced, stepping through PMax, P0, P1, and P2. This will result in corresponding reductions in power usage, causing commensurate reduction in operating temperature. As the teachings of Cho and Gold make clear, a system that is operable at four different temperatures (as taught by Cho), would have been implemented by operability at four different frequencies (as taught by Gold).

173. Additionally, like Cho, Gold discloses multiple components (e.g., “Core 0” and “Core 1”) that are each operable at different frequencies:


**FIG. 5B**

174. Ex-1009, Fig. 5B (annotated), ¶5 (“In one embodiment, an apparatus includes a plurality of processing cores (each of which has a respective standard operating Frequency”). A POSITA would have thus recognized Cho and Gold’s systems disclose interrelated teachings that would be especially compatible in the combined Gold-Cho system.

175. Moreover, a POSITA implementing the Cho-Gold teachings would have understood that each of Cho’s four independently controlled clocks, clk1- clk4, would have been operable at distinct frequencies, as taught by Gold. A POSITA

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

would have been motivated to implement Cho's teachings with four different frequencies to more particularly control the operating temperature of each component. A POSITA would have been motivated to assign specific and different frequency values (as taught by both Cho and Gold) to Cho's maximum and minimum temperature thresholds, and one or more intermediate temperature thresholds because doing so would have allowed for more predictable and controllable temperature actions using "thermal information corresponding to various utilization, temperature, and power entry/exit criteria" (Ex-1009 ¶6) as taught by Gold. As shown in Gold's Figure 7, this predictable relationship allows for improved power and temperature budgeting and controllability of each different component. Ex-1009 ¶66.

**FIG. 7**

176. Indeed, Gold specifically teaches to monitor cores to prevent and control overheating. Ex-1005 ¶57 (“one or more temperatures or thermal characteristics are monitored to ensure that overclocked cores are not overheating.”).


177. A POSITA would have had a reasonable expectation of success in implementing Gold’s different frequencies with Cho’s temperature reduction actions because Gold specifically explains its temperature and frequency control by including flow charts, such as shown in Figures 3, 4, and 6.

2. Independent Claim 1

a. Element 1[pre]: A processor comprising:

178. Cho teaches the preamble of Claim 1 to the extent it is limiting.

179. Cho teaches a “multi-core processor” that is configured for “functional block level thermal control.” Ex-1005 ¶38. Cho’s processor includes “four primary functional units—two processor cores 405 and 410, a memory controller hub (NorthBridge 415), and a graphics processing unit (GPU) 420.” *Id.* Figure 4, annotated below, illustrates Cho’s processor, 400 (green):




Id., Fig. 4 (annotated).

180. Accordingly, Cho discloses a processor as described in the preamble.

b. Element 1[a]: a plurality of cores including a first core operable at a first voltage and at a first frequency, and

181. Cho discloses Element 1[a].

182. Cho discloses that “processor 400 includes four primary functional units,” including “two processor cores 405 [red] and 410” (blue). Ex-1005 ¶38. Cho’s processor core 405 (red), illustrated in annotated Figure 4, is a “first core.”



Id., Fig. 4 (annotated).

183. Cho discloses that “each of the functional units ... each includes a voltage plane through which power is supplied independent of power supplied to the other functional units.” Ex-1005 ¶21. A “[p]ower control unit 320” provides a “corresponding voltage control signal to each of the functional units,” and each

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge


“voltage control signals vctrl1-vctrl4 corresponds to a unique one of the functional units.” Id. ¶37. Cho’s first core is thus operable at a first voltage, set by the “vctrl1” signal.

184. Cho also discloses that a “clock control unit is coupled ... to provide separate clock signals to each of the functional units (e.g., clk1 to a first functional unit, clk2 to a second functional unit, and so forth.” Ex-1005 ¶35. The clock control unit “adjust[s] the frequency of each of clock signals clk1-clk4, independently of one another, based on signals received from control logic 305.” Id. Cho’s first core is thus operable at a first frequency, set by “clk1” (see, id., Fig. 4 (red box around clk1 and vctrl1)).

c. Element 1[b]: a second core operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency;

185. Cho and Gold discloses Element 1[b].

186. Cho discloses that “processor 400” includes “two processor cores 405 and 410” Ex-1005 ¶38. Cho’s processor core 410 (blue), shown in annotated Figure 4, is a “second core.”



Id., Fig. 4 (annotated).

187. The second core is operable at a “second voltage independent of the first voltage.” Each functional unit, including the second core, has “a voltage plane through which power is supplied independent of power supplied to the other functional units.” Ex-1005 ¶21. Cho also discloses that “*the voltages on the voltage planes for each of the functional units can be set independently of one another.*”

Id. (emphasis added). For example, “each of the functional units ... is coupled to receive a voltage control signal that can change the voltage on a corresponding voltage plane *without affecting the voltages on the voltage planes* of any of the other ones of the functional units.” Id. ¶40. Thus, the second core is operable at a second voltage, “vctrl2,” independent of the first voltage.

188. Additionally, the “second core is operable at a ... second frequency different from the first frequency.” First, as noted above, Cho discloses a “clock control unit” that is coupled “to provide *separate* clock signals to each of the functional units.” Ex-1005 ¶35. And, “a change in the clock frequency for a clock signal provided to one functional unit need not affect the frequencies of the respective clock signals provided to other ones of the functional units.” *Id.* The clock signals operate in frequencies “independently of one another, based on signals received from control logic 305.” *Id.* The second frequency of core 2 set by “clk2,” which is a different signal than “clk1” that sets the frequency for core 1.

189. Additionally, as discussed above in Section XII.D.1 (regarding motivations to combine Cho and Gold), Cho also discloses each functional unit’s operating temperature is proportional to a frequency at which the functional unit is operating. Specifically, “TCU 102 is further configured to control the frequency of clock signals clk1-clk4 based on temperature data that is reported by functional units 105-120.” Ex-1005 ¶22. For example, Cho explains that “TCU 102 may lower the frequency of the clock signal” of a functional unit if the operating temperature is above a desired setting. *Id.* As explained in Section VIII.D.1, Cho teaches at least four different temperature settings: a maximum value, a minimum value, a first


at all.” *Id.* ¶203. Thus, Conroy explicitly contemplates redistributing substantially all of its variable power budget to the CPU when the GPU has “almost no load” or is “hardly used at all,” and vice versa.

192. Thus, it is my opinion that Conroy teaches that the power sharing logic is to dynamically allocate substantially all of the variable power budget to the first domain for a first workload, and to dynamically allocate substantially all of the variable power budget to the second domain for a second workload executed after the first workload, notwithstanding that the lower-bound of “substantially all” may not be reasonably ascertainable in other contexts.

- c. **Claim 17: The system of claim 14, wherein the power sharing logic is to further allocate the first portion of the variable power budget according to a first minimum reservation value for the first domain and allocate the second portion of the variable power budget according to a second minimum reservation value for the second domain.**

193. In my opinion, Conroy teaches Claim 17.

194. Conroy discloses a “power usage range” between a “minimum possible power consumption” and a “maximum power consumption.” Ex-1006 ¶53, Fig. 1. Conroy also explains that even when a subsystem has the lowest power allocation utilization of 0, Conroy’s power manager still allocates power to that subsystem to maintain the subsystem in a “low power mode.” *Id.* ¶198. A POSITA would have understood that the reason to have a low power mode is so that a component can be

**FIG. 5B**

191. Ex-1009, Fig. 5B (annotated), ¶5 (“In one embodiment, an apparatus includes a plurality of processing cores (each of which has a respective standard operating frequency”). Indeed, Gold specifically teaches four different operating frequencies. Gold teaches that its processors may be operable at distinct “performance states” (Pstate), such as PMax, P0, P1, P2, which each includes a “corresponding input voltage and clock frequency.” Ex-1009 ¶39. Figure 5 illustrates Gold’s performance states with four different frequency states and

corresponding power usages: 2.9 GHz and 19 Watts, 2.6 GHz and 15 Watts, 2.28 GHz and 12 Watts, and 1.73 GHz and 9 Watts.


500

<i>P-State</i>	<i>Power Usage</i>	<i>Frequency</i>
<i>P Max</i>	19W	2.90 GHz
<i>P 0</i>	15W	2.60 GHz
<i>P 1</i>	12W	2.28 GHz
<i>P 2</i>	9W	1.73 GHz

FIG. 5A

Id., Fig. 5A. Figure 5A thus teaches operability at four different frequencies.

192. As detailed in Section VIII.D.1, a POSITA would have been motivated to implement Gold's teachings of different specific frequencies and performance states. A POSITA would have been motivated to assign specific and different frequency values (as taught by Gold) to Cho's maximum temperature and first intermediate temperature because doing so would have allowed for more predictable and controllable temperature actions using "thermal information corresponding to various utilization, temperature, and power entry/exit criteria" as taught by Gold. Ex-1009 ¶6. As shown in Gold's Figure 7, this predictable relationship allows for improved power and temperature budgeting and controllability of each different component. *Id.* ¶66.

**FIG. 7**

193. Additionally, a POSITA would have had a reasonable expectation of success in implementing Gold's different frequencies with Cho's temperature reduction actions because Gold specifically includes a flow chart teaching how to make such controlled frequency reductions (step 650) based on temperature (step 610):





FIG. 6

194. The disclosed temperature control actions in Figure 5 the accompanying description at paragraphs 43-45, can be mapped to series of frequency

values, “ F_n ” for each functional unit, as shown in Figure 5’ further annotated to include a table.




Id., Fig. 5 (annotated).

195. As shown, Figure 5’ discloses that each function unit is operable at least three possible clock frequencies: F_{begin} , corresponding to the starting frequency or a reversal from a previous reduction of frequency; F_{lower} , corresponding to a first reduced frequency; and F_{min} , corresponding to a subsequently reduced frequency.

196. Finally, Cho explains that in addition to the three temperature control actions discussed in Figure 5 (and described in Ex-1005 ¶¶43-45), “a method based on multiple temperature thresholds *with one or more intermediate thresholds is possible.*” Ex-1005 ¶46. “In such an embodiment, each time the temperature reported by the functional unit exceeds one of the threshold values, its frequency is reduced

to a value commensurate with a given temperature range.” *Id.* With even one intermediate threshold, T_{inter} , each functional unit is operable at yet a fourth frequency, F_{inter} , as shown in further annotated Figure 5”.




Id., Fig. 5” (annotated with intermediate threshold).

197. Accordingly, the second core having a second frequency, set by clk_2 , is operable at a frequency (e.g., F_{inter}) different than first core frequency, set by clk_1 (e.g., F_{begin}), where $F_{inter} < F_{begin}$.

- d. **Element 1[c]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and**

198. Cho discloses Element 1[c].

199. Cho's multi-core processor includes "four primary functional units," one of which is a "graphics processing unit (GPU) 420." Ex-1005 ¶38, Fig. 4. The graphics processing unit 420 is "graphics processing circuitry." Figure 4, annotated below, illustrates Cho's graphics processing unit (GPU) 420 (purple):




200. Ex-1005, Fig. 4 (annotated). As Figure 4 also demonstrates, the graphics processing unit is coupled to Cho's first and second cores through NorthBridge 415 and Temperature Control Unit 102 (green lines). *Id.* A POSITA would have recognized the slash through the green lines as indicating multibit paths over which the cores may communicate with GPU 420 to perform graphic operations because "diagonal slash" notation denotes that lines with similar functions are grouped together. See Ex-1024.

201. Cho's graphics processing circuitry is also independently operable at a third frequency different from the first frequency and the second frequency. As explained in Section VIII.D.2.b, a "clock control unit" provides "separate clock signals to each of the functional units" Ex-1005 ¶35. The clock signals operate in frequencies "*independently of one another*, based on signals received from control logic 305." *Id.* Specifically, the third frequency of the GPU is set by "clk4," which is a different signal than both "clk1" and "clk2" that set the frequencies for core 1 and core 2. And, as discussed above in Section VIII.D.1 and VIII.D.2.c, "TCU 102 is further configured to control the frequency of clock signals clk1-clk4 based on temperature data that is reported by functional units 105-120." Ex-1005 ¶22. Thus, based on Cho's teachings, a POSITA would have understood that core 1 is operable at a maximum temperature, settable by a first frequency clk1, core 2 is operable at a first intermediate temperature, settable by a different second frequency clk2, and the GPU is operable at a second intermediate temperature, settable by a different third frequency clk4. While Cho does not state the specific operating frequencies of clk1, clk2, or clk4, a POSITA would have understood these frequencies to be operable at different values to correspond to the respective different operating temperatures.

202. Additionally, to the extent that a POSITA desired additional information about four different operating frequencies and corresponding power

dissipations, these teachings are provided by Gold, as explained above in sections VIII.D.1 and VIII.D.2.c. And a POSITA would have been motivated to, and had a reasonable expectation of success in applying these teachings to Cho's unique and independently operable clock frequency, clk_4 , for the same reasons.

203. GPU is operable at a third frequency different from the first frequency and the second frequency for the reasons explained above with respect the annotated Figures 5'', for example, the frequency of the graphics processing circuitry, the third frequency, could be the reduced frequency, F_{lower} , after the maximum threshold has been exceeded (light orange, below):




Ex-1005, Fig. 5'' (annotated with intermediate threshold).

204. Accordingly, graphics processing circuitry is operable at a frequency, set by clk4 (e.g., F_{lower}) different from first and second frequencies, set by clk1 (e.g., F_{begin}) and clk2 (e.g., F_{inter}), where $F_{lower} < F_{inter} < F_{begin}$.

e. Element 1[d]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first, second, and third frequencies.

205. Cho discloses Element 1[d].

206. Cho discloses that one of the four primary functional units is “a memory controller hub (NorthBridge 415).” Ex-1005 ¶38. Cho’s memory controller hub (NorthBridge 415) is an “interconnect” because it connects the processor’s various components. As discussed in Section V.D, interconnects and buses are commonly used in system components to interconnect the system to components such as memory elements, system components, resources and custom circuitry. A POSITA would have recognized the slash through the green lines as indicating multibit paths interconnected at NorthBridge 415, and would have further recognized the term “northbridge” as referring to a commonly known interconnect at the time of the ’588 Patent. See Ex-1024; Section V.D. NorthBridge is a term coined is logic to connect the CPU to the memory and other high speed devices via the frontside bus. See Ex-1025. Figure 4, annotated below, illustrates memory controller hub (NorthBridge 415) (green):



Id., Fig. 4 (annotated).

207. As Figure 4 illustrates, memory controller hub (NorthBridge 415) (the “interconnect”) connects graphics processing unit (GPU) 420 (the “graphics processing circuitry”) to off-processor memory (orange) via a multibit path (green).

Id. As shown in Figure 4, “clk3” is the clock signal for NorthBridge 415. Id.

208. NorthBridge 415 is operable at a fourth frequency different from the first frequency and the second frequency. As explained in Section VIII.D.2.b, a “clock control unit” provides “separate clock signals to each of the functional units” Ex-1005 ¶35. The clock signals operate in frequencies “*independently of one another*,” based on signals received from control logic 305.” Id. Specifically, the fourth frequency of the NorthBridge 415 is set by “clk3,” which is a different signal


U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

than both “clk1” and “clk2” that set the frequency for core 1 and core 2. And, as discussed above in Section VIII.D.1 and VIII.D.2.c, “TCU 102 is further configured to control the frequency of clock signals clk1-clk4 based on temperature data that is reported by functional units 105-120.” Id. ¶22. Thus, based on Cho’s teachings, a POSITA would have understood that core 1 is operable at a maximum temperature, settable by a first frequency clk1, core 2 is operable at a first intermediate temperature, settable by a different second frequency clk2, and the NorthBridge is operable at a minimum temperature, settable by a different fourth frequency clk3. While Cho does not state the specific operating frequencies of clk1, clk2, or clk3, a POSITA would have understood these frequencies to be operable at different values to correspond to the respective different operating temperatures.

209. Additionally, to the extent that a POSITA desired additional information about four different operating frequencies and corresponding power dissipations, these teachings are provided by Gold, as explained above in sections VIII.D.1 and VIII.D.2.c. And a POSITA would have been motivated to, and had a reasonable expectation of success in applying these teachings to Cho’s unique and independently operable clock frequency, clk3, for the same reasons.

210. NorthBridge 415 is operable at a fourth frequency different from the first frequency and the second frequency for the reasoned explained above with respect the annotated Figures 5”, for example, frequency of the interconnect, fourth

frequency, could be the further reduced frequency, $F_{\min}$, after the maximum threshold has been exceeded and remains exceeded (orange, below), in which the “2nd Temp Control Action 525” is taken:



Ex-1005, Fig. 5'' (annotated with intermediate threshold). Accordingly, the interconnect is operable at a frequency, set by clk_3 (e.g., $F_{\min}$) different from first, second, and third frequencies, clk_1 (e.g., F_{begin}), clk_2 (e.g., F_{inter}), clk_3 (e.g., F_{lower}), where $F_{\min} < F_{\text{lower}} < F_{\text{inter}} < F_{\text{begin}}$.

211. In summary, each of the four functional units of processor 400 is operable at a frequency different from the others, as taught by Cho. A “clock control unit” provides “separate clock signals to each of the functional units,” clk_1-clk_4 , using independently controllable phase lock loops, PLL. Ex-1005 ¶35. Moreover,

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge


there are expressly disclosed a minimum of at least four frequency values that can be assigned independently to teach of the four functional units, with the express teaching to add one or more intermediate thresholds and corresponding frequencies.

Id. ¶¶44-46, Fig. 5.

3. **Dependent Claim 2: The processor of claim 1, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.**

212. Cho and Gold discloses Claim 2.


213. Cho discloses that a “temperature control unit 102 includes control logic 305” and “clock control logic 310.” Ex-1005 ¶32. Figure 4, annotated below, illustrates the temperature control unit 102 (yellow) containing control logic 305 and clock control logic 310:



Id., Fig. 4 (annotated).

214. The clock control unit 310 “provide[s] separate clock signals to each of the functional units (e.g., *clk1* to a first functional unit, *clk2* to a second functional unit, and so forth).” Ex-1005 ¶35. The clock control unit 310 includes “a plurality of independently controllable phase locked loops (PLLs), each coupled ... to provide a corresponding one of the clock signals to a functional unit.” *Id.* The PLLs of Cho’s clock control unit that provide clock signals to the first and second cores are a “first power management circuitry and/or logic.” See Section V.A.

215. Figure 3, annotated below, illustrates Cho's clock control unit 310 (yellow):



Id., Fig. 3 (annotated).

216. Cho also discloses that “[c]lock control unit 310 is configured to adjust the frequency of each of clock signals clk1-clk4, independently of one another, based on signals received from control logic 305.” Ex-1005 ¶35. For example, the clock control unit may cause the first and second cores to operate independently at first and second frequencies “for temperature control purposes.” Id. ¶36. Indeed, “if a given functional unit is reporting a temperature over a predetermined threshold, control logic 305 may direct clock control unit 310 to reduce the frequency of its

corresponding clock signal.” *Id.* A POSITA would have known that the PLLs of clock control unit 310 thus cause the first and second cores to operate independently at first and second frequencies.

4. Dependent Claim 4: The processor of claim 2, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency.

217. Cho and Gold discloses Claim 4.

218. The clock control unit 310 includes “a plurality of independently controllable phase locked loops (PLLs), each coupled ... to provide a corresponding one of the clock signals to a functional unit.” Ex-1005 ¶35. The PLL that provides clock signal clk4 to graphics processing unit 420 (orange) is a “second power management circuitry,” as illustrated in Figure 3, annotated below.




Fig. 3

Id., Fig. 3 (annotated).

219. As explained, “[c]lock control unit 310 is configured to adjust the frequency of each of clock signals clk1-clk4, independently of one another, based on signals received from control logic 305.” Ex-1005 ¶35. For example, the clock control unit may cause the graphics processing unit 420 to operate independently at a third frequency “for temperature control purposes.” *Id.* ¶36.

5. Independent Claim 8

a. Element 8 ``` : A system comprising: ```


220. Cho discloses Claim 8’s preamble to the extent it is limiting.

221. Cho discloses that its “integrated circuit” is “a processor, such as a microprocessor suitable for use in a computer system.” Ex-1005 ¶9.

b. Element 8[a]: a system memory;

222. Cho discloses Element 8[a].

223. Cho discloses that the processor includes “a memory controller hub (NorthBridge 415).” Ex-1005 ¶38. As Figure 4, annotated below, illustrates, the memory controller hub (NorthBridge 415) connects to system memory (maroon).




Id., Fig. 4 (annotated).

c. Element 8[b]: a plurality of cores coupled to the system memory;

224. For the reasons described above for Element 1[a], Cho discloses a plurality of cores.

225. As illustrated in Figure 4, annotated below, the first and second cores (red and blue) are connected to memory (maroon) through NorthBridge 415 (green):



Id., Fig. 4 (annotated).

- d. **Element 8[c]: a first core of the plurality of cores operable at a first voltage and at a first frequency, and a second core of the plurality of cores operable at a second voltage independent of the first voltage and at a second frequency different from the first frequency;**

226. For the reasons described above for Elements 1[a] and 1[b], Cho and Gold disclose Element 8[c].

- e. **Element 8[d]: graphics processing circuitry coupled to the plurality of cores to perform graphics operations, the graphics processing circuitry independently operable at a third frequency different from the first frequency and the second frequency; and**

227. For the reasons described above for Element 1[c], Cho and Gold disclose Element 8[d].

- f. **Element 8[e]: an interconnect to couple the graphics processing circuitry to a memory device, the interconnect operable at a fourth frequency different from the first and second frequencies.**

228. For the reasons described above for Element 1[d], Cho and Gold disclose Element 8[e].

6. **Dependent Claim 9: The system of claim 8, further comprising: first power management circuitry and/or logic to cause the first core to operate at the first frequency and to cause the second core to independently operate at the second frequency.**

229. For the reasons described above for Claim 2, Cho and Gold disclose Claim 9.

7. **Dependent Claim 11: The system of claim 9, further comprising: second power management circuitry and/or logic to cause the graphics processing circuitry to operate at the third frequency.**

230. For the reasons described above for Claim 4, Cho and Gold disclose Claim 11.

8. **Independent Claim 20**

a. **Element 20[pre]: An apparatus comprising:**

231. Cho teaches the preamble of Claim 20 to the extent it is limiting.

232. Cho teaches a “multi-core processor,” which is an apparatus, configured for “functional block level thermal control.” Ex-1005 ¶38.

b. **Element 20[a]: means for operating ...**

233. As described in Section VIII.A.1.b-d, above, the corresponding structure is a “power control unit” and equivalents, performing the function recited in the remainder of this claim element. Notably, these functions are the same as those discussed above with respect to Elements 1[a], 1[b], and 1[c]. Cho and Gold disclose Element 20[a] because they disclose equivalent structure performing the recited functions.

234. Specifically, Cho’s equivalent structure is its TCU 102, that “perform[s] temperature control actions.” Cho’s TCU is discussed in Sections XII.A.1.b-d, above. Cho discloses TCU 102 controls each functional unit (e.g. “core” or “GPU”) based on “a number of different [temperature] thresholds,”

including maximum, one or more intermediate, or minimum thresholds. Ex-1005 ¶24. For example, if a functional unit's temperature exceeds a maximum allowable threshold, TCU 102 may “reduc[e] its respective clock frequency or voltage in an attempt to reduce the temperature of the functional unit.” *Id.* Moreover, as set forth in those sections, TCU 102 provides the capability to operate the first core, second core, and GPU at independent and distinct voltages and frequencies, respectively. See Sections XII.A.1.b-d, above.

c. Element 20[b]: means for independently operating ...

235. As described in Section VIII.A.D.4, above, the equivalent structure is the “plurality of independently controllable phase locked loops (PLLs)” in clock control unit 310. Ex-1005 ¶35. The PLL that provides clock signal clk4 to graphics processing unit 420 is a “second power management circuitry,” as shown in Figure 3. Cho discloses Element 20[b] because it discloses equivalent structure performing the recited functions.




Fig. 3

d. Element 20[c]: interconnect means

236. As described in Section VIII.A.1.e, above, the corresponding structure is a “interconnect 315” and equivalents, performing the function recited in the remainder of this claim element. Notably, this function is the same as those discussed above with respect to Element 1[d]. Cho discloses Element 20[c] because it discloses equivalent structure performing the recited functions.

237. Specifically, Cho’s equivalent structure is “memory controller hub (NorthBridge 415)” that acts to connect the processor’s various components. Ex-1005 ¶38. Cho’s NorthBridge 415 is discussed in Sections VIII.A.1.e, above.

Moreover, as set forth in those sections, Cho's NorthBridge 415 performs the recited function.

E. Ground 5: Claims 3, 7, 10, and 14-19 are rendered obvious by Cho, Gold, and Aguilar.

1. A POSITA would have been motivated to combine Cho, Gold, and Aguilar's teachings and would have had a reasonable expectation of success.


238. In addition to the reasons set forth below with respect to specific claim elements, a POSITA would have been motivated to combine the teachings of Cho, Gold, (Cho-Gold) and Aguilar and would have had a reasonable expectation of success in doing so because each relates to the same well-known technologies.

239. Cho-Gold, and Aguilar relate to and describe similar methods for controlling and optimizing temperature within computer systems. Ex-1005 (Cho), Abstract; Ex-1007 (Aguilar), Abstract. For example, Cho discloses “a temperature control unit ... wherein the temperature control unit is configured to monitor a temperature of each of the plurality of functional units based on temperature information provided from the temperature sensors.” Ex-1005, Abstract. The temperature control unit is then configured to, if the temperature exceeds a first threshold value, “perform a first temperature control action ...” *Id.* Likewise, Gold discloses its processor “may be coupled to cooling subsystem 120” to “control the temperature(s) of processing subsystem 110” and tracks “temperature criteria” to determine whether to apply overclocking. Ex-1009 ¶¶19, 41. Similarly, Aguilar


U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

discloses “[a] computer implemented method, data processing system, and computer usable code ... for optimizing thermal performance of a computer system.” Ex-1007, Abstract. Aguilar requests “a thermal index is ... for each of the set of system resources to form a set of thermal indexes.” *Id.* Then, “[a]n action from a set of actions is identified to reduce resource utilization” and “reduce a thermal state of the computer system.” *Id.*

240. Moreover, Cho and Aguilar include similar NorthBridge-SouthBridge architecture. As illustrated in Figure 4, Cho relies on a connection between NorthBridge 415 (green) and SouthBridge (yellow) to carry out its computer functions:



241. Ex-1005, Fig. 4. Likewise, as illustrated in Figure 2, Aguilar also relies on a NorthBridge (green)-SouthBridge (yellow) structure:



242. Ex-1007, Fig. 2. Thus, given Cho and Aguilar’s similar computer architecture, a POSITA would have understood that these systems are particularly compatible. See Section V.C.

243. A POSITA would have been motivated to combine the teachings of Cho-Gold, and Aguilar because they would have recognized the benefit of implementing an I/O controller hub, such as taught by Aguilar’s element 204, and thermal indexes in the combined Cho-Gold system to improve temperature control and resource utilization throughout a complete computer system. Aguilar teaches that “identification of a set of system resources” associated with a computer system “is performed and a thermal index is requested for each set of system resources.” Ex-1007, Abstract (emphasis added). Aguilar teaches that a “set of actions is


identified to reduce resource utilization,” and the actions are “implemented to reduce a thermal state of the computer system.” *Id.* Thus, a POSITA would be motivated to implement Aguilar’s teachings with Cho-Gold to allow for temperature control of a wide range of system resources, including I/O devices, as taught by Aguilar.

244. A POSITA would also have had a reasonable expectation of success in combining the teachings of Cho, Gold, and Aguilar. A POSITA would have understood that these references disclose interrelated teachings, including well-known temperature control methods and NorthBridge/SouthBridge architectures, based on routine technologies that would have been amendable to various well-understood and predictable combinations, as explained in further detail below.

2. Dependent Claim 3: The processor of claim 1, further comprising: a non-core power domain comprising a cache shared by the plurality of cores.

245. Cho-Gold and Aguilar disclose Claim 3.

246. The non-core components in Cho’s processor, including the GPU 420 (purple), memory controller hub (NorthBridge 415) (green), and Temperature Control Unit 102 (yellow) form a “non-core power domain.”




Ex-1005 ¶38, Fig. 4 (annotated).

247. A POSITA would have understood that Cho's memory controller hub NorthBridge 415 contained a cache, specifically an on-chip or "Level 1" cache, because such "L1" caches were common components of high-performance processors at the time of the '588 Patent. And as discussed in Section V.E, it was common for multicore processors to include a multi-level cache that includes Level 1 (L1) caches and Level 2 (L2) caches. In fact, the i486 processor included an L1 cache in 1993.. Cache memory is typically SRAM memory, which is substantially faster than main memory comprised of DRAM, and including a fast, on-chip L1

cache therefore makes retrieving data from a computer's memory faster and more efficient. *Id.*


248. Moreover, as Figure 4 illustrates below, both Core 1 405 and Core 2 410 are connected to and share memory through controller hub (NorthBridge 415).



Ex-1005, Fig. 4 (annotated).

249. To the extent that a POSITA would have required any additional teaching that Cho's NorthBridge would have been implemented with a cache, Aguilar provides such teaching. Aguilar discloses that "memory may be ... a cache such as found in *north bridge* and memory controller hub 202 in Fig. 2." Ex-1007,

4:57-60. A POSITA would have been motivated to implement Cho's NorthBridge with a cache memory, for example as taught by Aguilar, as shown in annotated Figure 4 below, because, as noted above, including a cache makes retrieving data from a computer's memory more efficient, data can be retrieved easily from cache memory, and data stored in caches is typically more readily available to system components than data stored in main memory. Id., See Section V.E.



Ex-1005, Fig. 4 (annotated).

3. **Dependent Claim 7: The processor of claim 3, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively.**

250. For the reasons described above for Elements 1[c] and 1[d], Cho-Gold and Aguilar disclose Claim 7. A POSITA would have known Cho's power management circuitry and/or logic causes the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores because the GPU 420 and controller hub (NorthBridge 415), which are part of the non-core power domain, operate at different frequencies than the cores as explained in Sections VIII.D.2.d-e.

4. **Dependent Claim 10: The system of claim 8, further comprising: a non-core power domain comprising a cache shared by the plurality of cores.**


251. See discussion in Section VIII.D.5, regarding Claim 8. For the reasons described above for Claim 3, Cho-Gold and Aguilar disclose Claim 10.

5. **Dependent Claim 14: The system of claim 10, further comprising: power management circuitry and/or logic to cause the non-core power domain to operate at a different frequency from the first and second frequencies of the first and second cores, respectively.**

252. For the reasons described above for Claim 7, Cho, Gold, and Aguilar disclose Claim 14.

6. **Dependent Claim 15:** The system of claim 8, further comprising: a chipset coupled to the plurality of cores, the chip set comprising a plurality of interfaces to couple the cores to a plurality of I/O devices.


253. Cho-Gold and Aguilar disclose Claim 15. Cho discloses that NorthBridge 415 (green) is connected to a SouthBridge (yellow), as shown in annotated Figure 4:



Ex-1005, Fig. 4 (annotated).

254. A POSITA would have known that a SouthBridge is “a chipset.” See Section V.C. Moreover, as shown in Figure 4, the SouthBridge is connected to Core 1 405 and Core 2 410 through NorthBridge 415. *Id.*; Ex-1005, Fig. 4.

255. Additionally, Aguilar teaches a “south bridge and input/output (I/O) controller hub (ICH) 204.” Ex-1007, 3:44-52. Figure 2, annotated below, illustrates Aguilar’s SouthBridge (yellow) and I/O controller hub (yellow):



Id., Fig. 2 (annotated).


256. The SouthBridge and I/O controller hub, connect to several I/O devices (orange) via buses 240 and 238 (blue), which a POSITA would have known are “interfaces.” Ex-1007, 3:55-4:4. Aguilar discloses, “Audio adapter 216, keyboard and mouse adapter 220, modem 222, read only memory (ROM) 224, hard disk drive (HDD) 226, CD-ROM drive 230, universal serial bus (USB) ports and other

communications ports 232, and PCI/PCIe devices 234 connect to south bridge and I/O controller hub 204 through bus 238 and bus 240.” Id. at 3:55:60.

257. As explained above in Section VIII.E.1, a POSITA would have been motivated to combine the teachings of Cho-Gold and Aguilar. To the extent that a POSITA would have desired additional teachings about implementing I/O devices in Cho, those specific teachings are provided by Aguilar as described above. And a POSITA would have had a reasonable expectation of success in implementing in the combined Cho-Gold system a SouthBridge with I/O devices as described by Aguilar, because Cho already teaches using a SouthBridge. *Id.*; Ex-1005, Fig. 4.

7. Dependent Claim 16: The system of claim 15, wherein one of the I/O devices comprises an audio generation device.

258. Cho-Gold and Aguilar disclose Claim 16. Aguilar discloses that an “audio adapter 216” (orange) connects to a “south bridge and I/O controller hub 204 [yellow] through bus 238 and bus 240” (blue).” Ex-1007, 3:53-60.



Id., Fig. 2 (annotated).

259. Audio adapter 216 is an “audio generation device.”


8. Dependent Claim 17: The system of claim 16, wherein one of the I/O devices comprises a video rendering device.

260. Cho-Gold and Aguilar disclose Claim 17.

261. Aguilar discloses that “[i]nput/output or I/O devices (including but not limited to keyboards, *displays*, pointing devices, etc.) can be coupled to the system either directly or through intervening I/O controllers.” Ex-1007, 25:40-43. A display is a “video rendering device.”

9. Dependent Claim 18: The system of claim 17, wherein one of the I/O devices comprises a user input device.

262. Cho-Gold and Aguilar disclose Claim 18. Aguilar discloses that a “keyboard and mouse adapter 220” (orange) connect to a “south bridge and I/O controller hub 204 [yellow] through bus 238 and bus 240” (blue). Ex-1007, 3:53-60.




Id., Fig. 2 (annotated).

263. Keyboard and mouse adapter 220 are an “user input device.”

10. Dependent Claim 19: The system of claim 18, wherein one of the I/O devices comprises a network communication device.

264. Cho-Gold and Aguilar disclose Claim 19. Aguilar discloses that a “local area network (LAN) adapter 212” (orange) connect to a “south bridge and I/O controller hub 204 [yellow] through bus 238 and bus 240” (blue). Ex-1007, 3:53-60.



Id., Fig. 2 (annotated).

265. Local area network (LAN) adapter 212 is a “network communication device.”

F. Ground 6: Claims 5-6, and 12-13 are rendered obvious by Cho, Gold, and Shakkarwar.

1. A POSITA would have been motivated to combine Cho, Gold, and Shakkarwar’s teachings and would have had a reasonable expectation of success.

266. In addition to the reasons set forth below with respect to specific claim elements, a POSITA would have been motivated to combine the teachings of Cho-Gold and Shakkarwar and would have had a reasonable expectation of success in doing so because each relates to the same well-known technologies. See Section V.

267. Cho-Gold and Shakkarwar relate to and describe similar methods for managing power and optimizing performance within a computer system. Ex 1005 (Cho), Abstract; Ex 1008 (Shakkarwar), Abstract. For example, Cho discloses “a temperature control unit ... wherein the temperature control unit is configured to monitor a temperature of each of the plurality of functional units based on temperature information provided from the temperature sensors.” Ex-1005, Abstract. The temperature control unit is then configured to, if the temperature exceeds a first threshold value, “perform a first temperature control action” Id. The temperature control action may “result in a corresponding reduction of power consumption, and thereby an eventual decrease in temperature.” Id. ¶21. Cho explains that such actions “may be taken for the purposes of increasing performance” and “for conservation of power in times of low activity” Id. ¶37. Cho teaches similar approaches for reducing frequencies based on temperature. Id. ¶53.

268. Shakkarwar also relates to the “efficient use of electronic devices and, more particularly, to systems and methods for power management of electronic devices.” Ex-1008, 1:8-10. Shakkarwar discloses “[s]ystems and methods for dynamic power management of electronic devices” so that “power consumption may be dynamically managed while providing sufficient bandwidth for the system.” Ex-1008, Abstract. Shakkarwar dynamically manages power by monitoring a component’s “bandwidth ... for increasing or decreasing communication

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

requirements.” Id. at 6:41-43. The system then detects “a change in processing condition[s]” for a component, and upon detecting a change, “determines bandwidth 403 for an associated data connection based on the active components.” Id. at 6:63-7:7. The data rate “of a data connection is [then] adjusted” to reach optimal performance levels. Id. at 7:26-29. In this manner, “power consumption associated with communicating and processing information via data buses may be optimized while satisfying bandwidth requirements for the computer system.” Id. at 3:17-20; *see also id.* at 2:12-15, Cl. 21.


269. A POSITA would have been motivated to combine the teachings of Cho-Gold and Shakkarwar because they would have recognized a benefit of implementing a look-up table for dynamically managing power, such as taught by Shakkarwar, to optimize performance of components within a processor as taught by Cho. As noted above, Shakkarwar teaches that an advantage to its system is that it dynamically “manag[es] power consumption of electronic devices while [also] satisfying bandwidth requirements for communicating and/or processing information.” Id. at 2:34-39. Indeed, implementing Shakkarwar’s teachings in the combined Cho-Gold system would have been an obvious implementation detail to achieve the power and bandwidth optimization of Shakkarwar in Cho. Such a benefit would have been achieved in Cho-Gold by, for example, ensuring that Cho’s GPU and interconnect are being driving at compatible frequencies, so as not to generate

unnecessary clock cycles and corresponding heat, thus controlling temperature as suggested by Cho. A POSITA would have been motivated to implement Shakkarwar's teachings in Cho-Gold to achieve these benefits.

2. Dependent Claim 5: The processor of claim 4, wherein the second power management circuitry and/or logic includes a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.

270. Cho-Gold and Shakkarwar teach Claim 5. As explained above in Section VIII.D.4., Cho discloses a second power management circuitry and/or logic. Cho in view of Shakkarwar teaches that this power management circuitry would have been implemented to include a mapping of selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect.

271. Shakkarwar's dynamic power management system includes several components, including "a data controller 302" (green) connected to a "graphics controller 305" (purple) through "data connection 314" (green), as illustrated in Figure 3:




Id., Fig. 3 (annotated).

272. A POSITA would have understood that the data controller 302 and data connection 314 are an “interconnect,” and the “graphics controller 305” is “graphic processing circuitry.” For example, data controller 302 connects graphics controller 305 and other components within Shakkarwar to each other and to system memory 303. Ex-1008, 5:10-15, Fig. 4. Graphics controller 305 includes a “3-D engine 305e” and performs “3-D graphics operations” *Id.* at 6:7-16.

273. Shakkwarwar also maps selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect. Shakkwarwar discloses at step 401, its system monitors “bandwidth for a component,” such as the interconnect or graphics processing circuitry, for “increasing or decreasing communication requirements.” Ex-1008, 6:41-43. For example, “when an application that is known to use 3-D graphics is opened, system software can determine if the bus rate between the data controller [the ‘interconnect’] and the graphics controller [the ‘graphics processing circuitry’] is high enough to support the graphics requirements.” *Id.* at 6:49-52.

274. At step 402, the system detects “a change in processing condition[s]” for a specific component, and upon detecting a change, it “determines bandwidth 403 for an associated data connection *based on the active components.*” Ex-1008, 6:63-7:7. To determine the appropriate bandwidth for a data connection, Shakkwarwar utilizes “*a look-up table* [that] may include *a list of clock speeds, and/or voltage levels required to process information for specific devices and/or groups of devices.*” *Id.* at 7:7-10. Figure 4 illustrates this process:



Id., Fig. 4. Accordingly, Shakkarwar teaches determining the appropriate bandwidth for a component (such Graphics controller 305) utilizing a look-up table that includes a list of appropriate clock speeds. It also considers other active components, such as the data controller 302 (the “interconnect”) and its clock speeds. As such a POSITA would have understood that Shakkarwar teaches to map selectable frequencies of the graphics processing circuitry to selectable frequencies of the interconnect, using look-up tables.


275. A POSITA would have been motivated to implement these teachings of Shakkarwar in Cho-Gold for the reasons described above in Section VIII.F.1, including because Shakkarwar teaches, for example, that an advantage to its system

is that it dynamically “manag[es] power consumption of electronic devices while satisfying bandwidth requirements for communicating and/or processing information.” Id. at 2:34-39. A POSITA would have been motivated to include this advantage in Cho-Gold to improve responsiveness of the system and increase the agility of the system to respond to dynamically changing conditions, including the changing bandwidth requirements of Cho’s components during various activity states.

3. Dependent Claim 6: The processor of claim 5, wherein the second power management circuitry and/or logic is to select the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.

276. Cho-Gold and Shakkarwar disclose Claim 6. As explained above for Claim 5, Cho and Shakkarwar disclose second power management circuitry that determines bandwidth for an associated data connection based on the other active components by using a look-up table of clock speeds required to process information for specific devices.

277. Shakkarwar further discloses that “[a]t step 405, the data rate of a data connection is adjusted. In one embodiment, the data rate can be adjusted by adjusting operating of a clock generation device.” Ex-1008, 7:26-29. Figure 4 illustrates adjusting the data rate of a data connection:

**FIG. 4**

278. Shakkarwar also specifically discloses selecting “the fourth frequency for the interconnect based on detecting operation of the graphics processing circuitry at the third frequency.” In particular, Shakkarwar discloses that “when an application that is known to use 3-D graphics is opened, system software can determine if *the bus rate between the data controller [the “interconnect”] and the graphics controller [the “graphics processing circuitry”] is high enough to support the graphics requirements.*” Ex-1008, 6:49-52. A POSITA would have understood that

by applying these teachings from Shakkarwar to Cho-Gold, the fourth frequency for the interconnect would be selected based on detecting the frequency of the graphics processing circuitry.

4. Dependent Claim 12.

279. See discussion in Sections VIII.D.5 and VIII.D.7, regarding Claims 8 and 11. For the reasons described above for Claim 5, Cho-Gold and Shakkarwar disclose Claim 12.

5. Dependent Claim 13.

280. See discussion in Sections VIII.D.5 and VIII.F.4, regarding Claims 8 and 12. For the reasons described above for Claim 6, Cho-Gold and Shakkarwar disclose Claim 13.

IX. CONCLUSION

281. For the reasons set forth above, I believe that Claims 1-20 of the '588 Patent is unpatentable in view of the prior art.

282. In signing this declaration, I recognize that it will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place in the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

U.S. Patent No. 10,705,588
Declaration of Trevor Mudge

283. I reserve the right to supplement my opinions in the future to respond
to any arguments that the Patent Owner raises and to take into account new
information as it becomes available to me.

284. I declare under penalty of perjury under the laws of the United States
of America that the foregoing is true and correct.

Respectfully Submitted

/Trevor Mudge/

Trevor Mudge

Date: February 1, 2023