

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner,

v.

SINGULAR COMPUTING LLC,
Patent Owner.

IPR2023-00397
Patent 11,169,775 B2

Before JAMESON LEE, KEN B. BARRETT, and
JASON M. REPKO, *Administrative Patent Judges*.

BARRETT, *Administrative Patent Judge*.

DECISION
Granting Institution of *Inter Partes* Review
35 U.S.C. § 314

I. INTRODUCTION

A. *Background and Summary*

Google LLC (“Petitioner”)¹ filed a Petition requesting *inter partes* review of U.S. Patent No. 11,169,775 B2 (“the ’775 patent,” Ex. 1001). Paper 2 (“Pet.”). The Petition challenges the patentability of claims 1–5, 7–14, and 16–23 of the ’775 patent. Singular Computing LLC (“Patent Owner”)² filed a Preliminary Response to the Petition. Paper 7 (“Prelim. Resp.”).

An *inter partes* review may not be instituted “unless . . . the information presented in the petition . . . shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” 35 U.S.C. § 314(a) (2018). Having considered the arguments and evidence presented by Petitioner and Patent Owner, we determine that Petitioner has demonstrated a reasonable likelihood of prevailing on at least one of the challenged claims of the ’775 patent. Accordingly, we institute an *inter partes* review as to all the challenged claims of the ’775 patent on the sole ground of unpatentability set forth in the Petition.

B. *Related Proceedings*

Each party identifies the following district court litigation involving the ’775 patent as a related matter: *Singular Computing LLC v. Google LLC*, No. 1:21-cv-12110-FDS (D. Mass.). Pet. xviii; Paper 5, 1.

¹ Petitioner identifies Google LLC as the real party-in-interest. Pet. xvii.

² Patent Owner identifies Singular Computing LLC as the real party-in-interest. Paper 5, 1.

Petitioner filed another petition in IPR2023-00396 challenging claims of the '775 patent, and filed petitions in IPR2023-00394 and IPR2023-00395 challenging claims in related U.S. Patent No. 10,754,616. *See* Pet. xviii; Paper 5, 1.

Patent Owner also identifies *Singular Computing LLC v. Google LLC*, No. 1:19-cv-12551-FDS (D. Mass.) and *Google LLC v. Singular Computing LLC*, Appeal Nos. 2022-1866, 2022-1867, and 2022-1868 (Fed. Cir.). Paper 5, 1³.

C. The '775 Patent

The '775 patent states that “[l]ow precision computers can be efficient at finding possible answers to search problems.” Ex. 1001, code (57). The '775 patent relates to a computer system that augments low precision computing with a small amount of high precision computing, to improve search quality with little additional computing. *Id.* According to the '775 patent:

[T]oday’s CPU chips make inefficient use of their transistors. For example, a conventional CPU chip containing a billion transistors might enable software to perform merely a few operations per clock cycle. Although this is highly inefficient, those having ordinary skill in the art design CPUs in this way for what are widely accepted to be valid reasons. For example, such designs satisfy the (often essential) requirement for software compatibility with earlier designs. Furthermore, they deliver great precision, performing exact arithmetic with integers typically 32 or 64 bits long and performing rather accurate and widely standardized arithmetic with 32 and 64 bit

³ Patent Owner additionally identifies, as a related matter, Patent Trial and Appeal Board case IPR2023-00410. Paper 5, 1. That case, however, is not a recognized proceeding in the Board’s Patent Trial and Appeal Case Tracking System (P-TACTS).

floating point numbers. Many applications need this kind of precision. As a result, conventional CPUs typically are designed to provide such precision, using on the order of a million transistors to implement the arithmetic operations.

Id. at 2:62–3:10. According to the ’775 patent, however:

There are many economically important applications, however, which are not especially sensitive to precision and that would greatly benefit, in the form of application performance per transistor, from the ability to draw upon a far greater fraction of the computing power inherent in those million transistors. Current architectures for general purpose computing fail to deliver this power.

Id. at 3:11–17.

The ’775 patent discloses embodiments directed to computer processors or other devices which use low precision high dynamic range (LPHDR) processing elements⁴ to perform computations (such as arithmetic operations). *Id.* at 5:58–62. According to the ’775 patent, because LPHDR processing elements are small, a single processor or other device may include a very large number of LPHDR processing elements, adapted to operate in parallel with each other and therefore may constitute a massively parallel LPHDR processor or other device. *Id.* at 6:47–52.

The ’775 patent states that such a processor or device has not been described or practiced as a means of doing general purpose computing by those having ordinary skill in the art for two reasons: (1) it was commonly believed that LPHDR computation, in particular massive amounts of LPHDR computation, is not practical; and (2) it was commonly believed by

⁴ The ’775 patent states that “references herein to ‘processing elements’ within embodiments of the present invention should be understood more generally as any kind of execution unit, whether for performing LPHDR operations or otherwise.” Ex. 1001, 8:5–9.


those having ordinary skill in the art that massive amount of even high precision computation on a single chip or in a single machine is not useful without a corresponding increase in bandwidth between processing elements and into and out of the machine. *Id.* at 6:52–6:67.

The '775 patent also states:

Despite these views—that massive amounts of arithmetic on a chip or in a massively parallel machine are not useful, and that massive amounts of LPHDR arithmetic are even worse—embodiments of the present invention disclosed herein demonstrate that massively parallel LPHDR designs are in fact useful and provide significant practical benefits in at least several significant applications.

Id. at 7:1–7.

Figure 1 of the '775 patent is reproduced below:




Id. at Fig. 1. Figure 1 shows the design of a SIMD (single instruction stream/multiple data stream) processor. *Id.* at 2:31–32, 3:20–22.

Computing system 100 of Figure 1 includes a collection of processing elements (PEs) that the '775 patent refers to as “the Processing Element Array” (PEA). *Id.* at 8:10–17. Figure 1 also depicts control unit (CU) 106, I/O unit (IOU) 108, various peripheral devices 110, and host computer 102. *Id.*

In an embodiment, host 102 is responsible for overall control of the computing system 100 and directs PEA 104 to perform massive amounts of computations in a useful way by causing the PE's to perform computations, typically on data stored in each PE, in parallel with one another. *Id.* at 8:27–28; 8:35–38. Figure 1 depicts specialized CU 106 having the primary task of retrieving and decoding instructions from an instruction memory and issuing the partially decoded instructions to all the PEs in PEA 104. *Id.* at 8:47–52. Host 102 loads a program for execution by PEA 104 into the CU instruction memory and then instructs CU 106 to interpret the program's instructions and cause PEA 104 to execute the instruction. *Id.* at 8:59–63.

Figure 1 further depicts I/O unit (IOU) 108 to interface host 102 and various peripherals in the host system with CU 106 and PEA 104. *Id.* at 9:6–10. The '775 patent discloses that IOU 108 may be connected only to the PEs at the edges of PEA 104, and that to get data out of, and perhaps into, PEA 104, the data is read and written at the edges of the array and CU instructions are performed to shift data between the edges and interior of the PEA. *Id.* at 9:66–10:5. PEA 104 may be implemented such that data may move from PEs only to their nearest neighbors so that there are no long distance transfers. *Id.* at 9:39–42.

Figure 2 of the '775 patent is reproduced below:



Id. at Fig. 2. Figure 2 is an example of the Processing Element Array of a SIMD processor. *Id.* at 2:33–35. Figure 2 shows a subset of the PEs in PEA 104, namely PE 202, PE 204, PE 206, PE 208, and PE 210. *Id.* at 9:45–47. According to the '775 patent, when CU 106 issues data movement instructions, all the PEs access data from or send data to their respective specified nearest neighbor. *Id.* at 9:47–49. For instance, every PE might access a specified data value in its neighbor to the West and copy it into its own local storage. *Id.* at 9:49–51.

D. Illustrative Claim

Of the challenged claims of the '775 patent, claims 1, 7, and 16 are independent claims. Claim 1, reproduced below with bracketed annotations⁵ inserted, is illustrative.

1. [1PRE] A computing system, comprising:
 - [1A] a host computer;
 - [1B] a computing chip comprising:
 - [1C] a processing element array comprising a first edge processing element positioned at a first edge of the processing element array, a second edge processing element positioned at the first edge of the processing element array, a first interior processing element positioned at a first location in the interior of the processing element array, and a second interior processing element positioned at a second location in the interior of the processing element array;
 - [1D] a first processing element connection connecting the first edge processing element with the first interior processing element;
 - a second processing element connection connecting the second edge processing element with the second interior processing element;
 - [1E] an input-output unit connected to the first edge processing element and the second edge processing element;
 - [1F] a first memory local to the first edge processing element;
 - a second memory local to the second edge processing element;
 - a third memory local to the first interior processing element;
 - a fourth memory local to the second interior processing element; and,
 - [1G] a fifth arithmetic unit;

⁵ We utilize Petitioner's annotations for claim 1 but have retained the paragraph formatting from the issued patent.

[1H] wherein the first edge processing element comprises a first arithmetic unit;

wherein the second edge processing element comprises a second arithmetic unit;

wherein the first interior processing element comprises a third arithmetic unit; and

wherein the second interior processing element comprises a fourth arithmetic unit; and,

[1I] a host connection at least partially connecting the input-output unit with the host computer;

[1J] wherein the first, second, third and fourth arithmetic units each comprises a corresponding multiplier circuit adapted to receive as a first input to the corresponding multiplier circuit a first floating point value having a first binary mantissa of width no more than 11 bits and a first binary exponent of width at least 6 bits, and to receive as a second input to the corresponding multiplier circuit a second floating point value having a second binary mantissa of width no more than 11 bits and a second binary exponent of width at least 6 bits;

[1K] wherein the fifth arithmetic unit comprises a corresponding multiplier circuit adapted to receive as inputs to the corresponding multiplier circuit two floating point values each of width at least 32 bits;

[1L] wherein the multiplier circuit corresponding to the first arithmetic unit comprises a first plurality of transistors and has no other transistors, the multiplier circuit corresponding to the second arithmetic unit comprises a second plurality of transistors and has no other transistors, the multiplier circuit corresponding to the third arithmetic unit comprises a third plurality of transistors and has no other transistors, the multiplier circuit corresponding to the fourth arithmetic unit comprises a fourth plurality of transistors and has no other transistors, and [1M] the multiplier circuit corresponding to the fifth arithmetic unit comprises a fifth plurality of transistors; and,

[1N] wherein the fifth plurality of transistors exceeds in number each of the first plurality of transistors, the second plurality of transistors, the third plurality of transistors, and the fourth plurality of transistors.

Ex. 1001, 30:21–31:25.

E. Evidence

Petitioner relies on the following references:

Name	Reference	Exhibit(s)
Stuttard	WO 2005/037326 A2, published April 28, 2005	1006
Shirazi	Shirazi et al., “ <i>Quantitative Analysis of Floating Point Arithmetic on FPGA Based Custom Computing Machines</i> ,” PROCEEDINGS OF THE IEEE SYMPOSIUM ON FPGAS FOR CUSTOM COMPUTING MACHINES, April 19–21, 1995 ⁶	1036

Petitioner also relies on the declaration of Paul Franzon, Ph.D. (Ex. 1003) in support of its arguments. The parties also rely on other exhibits as discussed below.

F. Asserted Ground of Unpatentability

Petitioner asserts that the challenged claims are unpatentable on the following ground⁷:

Claim(s) Challenged	35 U.S.C. § ⁸	Reference(s)/Basis
1–5, 7–14, 16–23	103(a)	Stuttard, Shirazi

⁶ See Ex. 1037 (the Declaration of Gordon MacPherson indicating the date of the IEEE Symposium).

⁷ Although there is only one ground in the Petition, Petitioner analyzes the challenged claims as two groups under the headings “Ground 1A” and “Ground 1B.” See, e.g., Pet. i–vi (Table of Contents).

⁸ For purposes of this Decision, we assume the challenged claims have an effective filing date prior to March 16, 2013, the effective date of the

II. 35 U.S.C. § 314(a)

On the same day, Petitioner filed two parallel petitions challenging the patentability of claims of the '775 patent. In the Petition before us, Petitioner challenges claims 1–5, 7–14, and 16–23 as obvious over Stuttard and Shirazi. Pet. 1. In IPR2023-00396, Petitioner challenges claims 1–5, 7–14, and 16–23 as obvious over other references not relied on in the Petition before us. IPR2023-00396, Paper 2, 1. On December 22, 2022, in accordance with the Board’s Consolidated Trial Practice Guide (“Consolidated TPG”),⁹ Petitioner filed a Notice Ranking Petitions and Explaining Material Differences. Paper 3 (“Notice”). In its Notice, Petitioner requests that we consider the Petition in IPR2023-00396 first and the Petition in this proceeding second. *Id.* at 1. Patent Owner did not respond to the Notice in either its Preliminary Response or in a separate paper, as the Consolidated TPG allows. *See Consolidated TPG*, 60.

We need not decide the issue of discretion under 35 U.S.C. § 314(a). For the reasons detailed in the Decision on Institution in IPR2023-00396, institution is denied in that case because the information presented in the petition does not show a reasonable likelihood of prevailing.

relevant aspects of the Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (“AIA”), and we apply the pre-AIA version of 35 U.S.C. § 103. *See* Pet. 1, 5 (applying pre-AIA provisions).

⁹ Patent Trial and Appeal Board Consolidated Trial Practice Guide (Nov. 2019), *available at* <https://www.uspto.gov/TrialPracticeGuideConsolidated>, 59–61 (explaining that the Board may exercise discretion under 35 U.S.C. § 314(a) to deny a petition if it determines that more than one petition challenging claims of the same patent is not warranted).

III. ANALYSIS OF PETITIONER'S CHALLENGE

A. Principles of Law

Petitioner bears the burden of persuasion to prove unpatentability of the claims challenged in the Petition, and that burden never shifts to Patent Owner. *Dynamic Drinkware, LLC v. Nat'l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015).

A patent claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4), if present, any objective evidence of obviousness or non-obviousness. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

B. The Level of Ordinary Skill in the Art

In determining the level of ordinary skill in the art, various factors may be considered, including the “type of problems encountered in the art; prior art solutions to those problems; rapidity with which innovations are made; sophistication of the technology; and educational level of active workers in the field.” *In re GPAC Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995) (quoting *Custom Accessories, Inc. v. Jeffrey–Allan Indus., Inc.*, 807 F.2d 955, 962 (Fed.Cir.1986)).

Petitioner contends that:

A [person of ordinary skill in the art] would have had a bachelor's degree in Electrical Engineering, Computer Engineering, Applied Mathematics, or the equivalent, and two years of academic or industry experience in computer architecture. More education could substitute for experience, and vice versa.

Pet. 4 (citing Ex. 1003 ¶¶ 33–37). Patent Owner states that, “[f]or the purposes of this Preliminary Response only, Patent Owner utilizes Petitioner’s proposed level of ordinary skill in the art.” Prelim. Resp. 9.

Petitioner’s definition is consistent with the level of ordinary skill reflected in the prior art references of record. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001) (recognizing that the prior art itself may reflect an appropriate level of skill in the art). For purposes of this decision, we apply Petitioner’s definition of the person of ordinary skill in the art.

C. Claim Construction

We apply the same claim construction standard used in district court actions under 35 U.S.C. § 282(b), namely that articulated in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). *See* 37 C.F.R. § 42.100(b). In applying that standard, claim terms generally are given their ordinary and customary meaning as would have been understood by a person of ordinary skill in the art at the time of the invention and in the context of the entire patent disclosure. *Phillips*, 415 F.3d at 1312–13. “In determining the meaning of the disputed claim limitation, we look principally to the intrinsic evidence of record, examining the claim language itself, the written description, and the prosecution history, if in evidence.” *DePuy Spine, Inc. v. Medtronic Sofamor Danek, Inc.*, 469 F.3d 1005, 1014 (Fed. Cir. 2006) (citing *Phillips*, 415 F.3d at 1312–17).

Petitioner asserts that, in the Petition, all claim terms have been given their ordinary and customary meaning, and does not propose any constructions under the “Claim Interpretation” heading. Pet. 6. Patent Owner states that it “agrees that no claim construction is necessary,” and impliedly asserts that it applies the “plain meaning” of all claim terms. Prelim. Resp. 9. We determine that no express claim construction is necessary for purposes of this Decision, except for “input-output unit” and the recitations regarding the interior processing element connections.

1. Input-Output Unit


Independent claims 1, 7, and 16 each recite an “input-output unit.” Ex. 1001, 30:39, 31:66, 34:13. In its analysis of, for example, independent claim 1, as applied to Stuttard, Petitioner reads the claimed “input-output unit” onto “[t]he bus (or alternatively, the bus plus the I/O engines and channels).” Pet. 30. Patent Owner argues, however, that a bus simply is a set of conductors and that a person of ordinary skill in the art would not understand a bus, in light of the Specification, to be an input-output unit. See Prelim. Resp. 26–27. Here, we consider whether a bus as in Stuttard reasonably can be regarded as an “input-output unit” in the context of the specification of the ’775 patent. We determine that it would be unreasonable and not consistent with the specification of the ’775 patent to do so.

There is no express definition of “input-output unit” in the specification of the ’775 patent. But that does not mean “input-output unit” automatically can be construed as encompassing anything that has some tangential involvement in getting data into and out of a component that accepts and provides data. The construction still has to be reasonable and consistent with the specification, even under a broadest reasonable interpretation rule of claim construction that is more lenient than the *Phillips*

standard which applies here. *See In re Smith Int'l, Inc.*, 871 F.3d 1375, 1382–83 (Fed. Cir. 2017). A particular reading is not deemed consistent with the specification simply because the specification does not expressly preclude that reading. *Id.* at 1382 (“Therefore, the Board’s reasoning that because the specification does not ‘in and of itself proscribe the Examiner’s construction,’ the examiner’s interpretation was reasonable, was erroneous.” (citation omitted)).

Although an input-output unit must be connected to a system component to function as an input-output unit with respect to that component, it would be unreasonable to regard that connection, such as a data bus, as itself an input-output unit. A data bus carries data from place to place but does not have a significant role on when to place data onto the bus or when to off load data from the bus, or what data to place on the bus.

Figure 1 of the '775 patent is reproduced below:



Ex. 1001, Fig. 1. Figure 1 is an example of an overall design according to one embodiment of the '775 patent. Ex. 1001, 2:31–32. Microsoft Computer Dictionary, Fifth Edition (2002), provides a meaning for “bus” as: “A set of hardware lines (conductors) used for data transfer among the components of a computer system.” Ex. 2001, 77 (cited at Prelim. Resp. 27). As is evident from the figure, the wired connections to and from the “Input/Output Unit,” depicted as bidirectional arrows, are not the “Input/Output Unit.”

The '775 patent also explains the “Input/Output Unit” as follows:

In order to get data into and out of the CU 106 and PEA 104, the I/O Unit 108 may interface the CU 106 and PEA 104 with the Host 102, the Host's memory (not shown in FIG. 1), and the system's Peripherals 110, such as external storage (e.g., disk drives), display devices for visualization of the computational results, and sometimes special high bandwidth input devices (e.g., vision sensors). *The PEA's ability to process data far faster than the Host 102 makes it useful for the IOU 108 to be able to completely bypass the Host 102 for some of its data transfers.*

Ex. 1001, 9:5–14 (emphasis added). We determine that this description of an ability of the “Input/Output Unit” to account for or accommodate the different processing rates of the components it serves indicates a capability beyond merely serving as a connector. Petitioner does not explain how a data bus has that capability and we do not, on this record, see how it would.

For the foregoing reasons, we determine it is unreasonable and not consistent with the specification of the '775 patent to regard a data bus as an “input-output unit.”

2. Interior Processing Element Connections

According to Patent Owner, the '775 patent recognized a data transfer bottleneck issue in the art and addresses this “by expressly avoiding both

(i) a shared bus and (ii) any direct connections between the vast majority of PEs in the PE array (PEA), and the rest of the computer (*i.e.*, the IO unit [or IOU]).” Prelim. Resp. 20–21 (citing Ex. 1001 (the ’775 patent), 9:60–10:11. Patent Owner argues that “[t]his teaching in the ’775 Patent is manifested in both the edge PE and interior PE limitations of the Challenged Claims (*i.e.*, connecting the IOU to the PEs only at the edges of the PEA, and *not connecting the IOU to the interior PEs without going through at least one edge PE*).” *Id.* at 21 (emphasis added); *see id.* (“Thus, the ’775 Patent connects its IO unit to only the edge PEs.”); *id.* at 24 (“A plain reading of the claims shows that the edge processing elements are connected to the input-output unit (and thus, at the logical edge of the array), and that the interior elements are not.”). In other words, Patent Owner contends that the independent claims preclude the interior PE elements from being directly connected to the input-output unit and preclude the use of a shared bus connecting the interior PEs to the I/O unit. We understand Patent Owner to identify—as embodying these preclusion aspects—limitations 1C, 1D, and 1E of independent claim 1 and the corresponding limitations of independent claims 7 and 16. *See id.* at 16 n.1, 21, 24–27, 30–31.

Petitioner, in its analysis of independent claims 1, 7, and 16, proposes a combination of Stuttard and Shirazi where the interior PE elements are directly connected to Stuttard’s bus, a portion of the structure mapped to the recited input-output unit. *See* Pet. 22 (demonstrative figure depicting the proposed combination), 30 (“The bus (or alternatively, the bus plus the I/O engines and channels) . . . is an input-output unit . . . [and] is connected to all the poly PEs[.]”). Thus, Petitioner implicitly construes the independent

claims as not precluding direct connections between interior PE elements and the input-output unit and as not precluding the use of a shared bus.

For the reasons that follow, we construe the independent claims as *not* precluding the interior PE elements from being connected directly to the input-output unit and as *not* precluding the use of a shared bus. We focus our discussion here on the limitations of independent claim 1.

Limitation 1C of independent claim 1 recites “edge processing element[s] positioned at a first edge of the processing element array” and “interior processing element[s] positioned at a . . . location in the interior of the processing element array.”¹⁰ Ex. 1001, 30:24–32. This limitation defines the physical positions and locations of the PEs, and does not address the connections.

Limitation 1D, however, does address the connections, reciting “a first processing element connection connecting the first edge processing element with the first interior processing element,” and “a second processing element connection connecting the second edge processing element with the second interior processing element.” *Id.* at 30:33–38. Thus, each respective pair of an edge PE and an interior PE must be connected.

Limitation 1E also addresses connections, reciting “an input-output unit connected to the first edge processing element and the second edge processing element.” *Id.* at 30:39–41. This limitation requires a connection

¹⁰ Petitioner asserts that the claim’s recitation of “a first location in the interior of the processing element array” means “not along an outer edge of the array.” Pet. 22. Patent Owner similarly asserts that “[a]n ‘interior processing element’ is not at the ‘edge’ and thus is inside the array.” Prelim. Resp. 24; *but see id.* (impliedly arguing that any PE connected to the I/O unit is, regardless of physical location, “at the logical edge of the array”).


between the I/O unit and the edge PEs, but does not address any connection or lack thereof between the I/O unit and the interior PEs.

Accordingly, none of the limitations identified by Patent Owner recites language precluding an interior PE from being connected to another component.

Patent Owner contends that the Specification confirms Patent Owner's proposed construction, asserting that "[t]he '775 Patent explains that only edge processing elements are connected to the input-output unit." Prelim. Resp. 25 (citing Ex. 1001, 9:9:55–10:5). The cited portion of the Specification states:

FIG. 3 shows a PE 302 that includes data connections to the IOU 108. PE 302 is connected at the North to PE 304, at the East to PE 306, at the South to PE 308, and at the West to PE 310. However, driving signals from inside the PEA 104 out to the IOU 108 usually requires a physically relatively large driving circuit or analogous mechanism. Having those at every PE may absorb much of the available resources of the hardware implementation technology (such as VLSI area). In addition, having independent connections from every PE to the IOU 108 means many such connections, and long connections, which also may absorb much of the available hardware resources. For these reasons, the connections between the PEs and the IOU 108 *may be limited to those PEs at the edges of the PE array 104*. In this case, to get data out of, and perhaps into, the PEA 104, the data is read and written at the edges of the array and CU instructions are performed to shift data between the edges and interior of the PEA 104.

Ex. 1001, 9:55–10:5 (emphasis added). This describes the embodiment of Figure 3 as depicting PE 302 as an interior PE (in that it has neighbors on all four sides) and as that PE 302 being connected to IOU 108. Figure 3 is reproduced below.



Id. at Fig. 3. Figure 3 “is an example of how a Processing Element in a Processing Element Array communicates data with other parts of the processor according to one embodiment of the present invention.” *Id.* at 2:36–39.

Patent Owner points to the statements in the block quote above indicating that there are certain disadvantages to the configuration of the Figure 3 embodiment. Prelim. Resp. 25. Patent Owner quotes the Specification’s statement that, in light of the disadvantages, “the connections between the PEs and the IOU 108 *may be limited* to those PEs at the edges of the PE array 104” as support for the proposition that the Specification explains that *only* edge PEs are connected to the I/O unit and that the claim preclude interior PEs from being directly connected to the I/O unit. *See id.* at 24–25 (quoting Ex. 1001, 9:66–10:1)) (emphasis added).

We do not agree with Patent Owner’s arguments. The Specification uses the permissive term “may” in stating that connections to the I/O unit “may be limited” to the edge PEs. Thus, the Specification suggests that is

one alternative embodiment of the disclosed invention. Although such a configuration may have advantages over the Figure 3 embodiment, Figure 3 also is identified as “one embodiment of the present invention.” Ex. 1001, 2:36–39.

Additionally, we note that Patent Owner omits from the block quote above the concluding sentence in that paragraph. That sentence provides: “The design may permit data to be pushed *from the IOU 108 inward to any PE in the array using direct connections*, but may require readout to occur by using the CU 106 to shift data to the edges where it can be read by the IOU 108.” Ex. 1001, 10:5–9 (emphasis added). This suggests a hybrid embodiment where data flows to any PE in the array via direct connection from the I/O unit but data flows out to the I/O unit only via an edge PE. *See id.* at 10:1–3 (the Specification indicating that the data in and data out may be treated differently in stating, “to get data out of, *and perhaps into, the PEA 104*, the data is read and written at the edges of the array”) (emphasis added).

Patent Owner has not explained adequately why the claims should be construed as limited to one disclosed embodiment.

Accordingly, we construe the independent claims as *not* precluding the interior PE elements from being connected directly to the input-output unit and as *not* precluding the use of a shared bus.

D. The Alleged Obviousness of Claims 1–5 Over Stuttard and Shirazi (Ground 1A)

Petitioner alleges that claims 1–5 of the ’775 patent would have been obvious over Stuttard and Shirazi. *See* Pet. 11–41 (addressing claim 1). Petitioner relies on Stuttard for much of the claimed subject matter, and turns to Shirazi for teachings related to the multiplier circuit (limitation 1J).

See, e.g., id. at 11 (“POSAs [persons of ordinary skill in the art] would have had reasons to adapt the PEs in Stuttard’s poly unit to perform floating-point math using Shirazi’s custom reduced-bitwidth floating-point values.”).

Patent Owner argues, *inter alia*, that a person of ordinary skill in the art would not have had a reason to make the proposed modification and that the proposed combination lacks the recited interior processing elements and the input-output unit. *See, e.g.,* Prelim. Resp. 1–3, 26–27.

1. Stuttard (Ex. 1006)

Stuttard discloses a “Multi-Threaded Array Processing (MTAP) architecture” with “parallel data processing capability,” “developed to address a number of problems in high performance, high data rate processing.” Ex. 1006, 8:23–9:14. Stuttard discloses the MTAP processor, which “has a standard, RISC-like, control unit with instruction fetch, caches and I/O mechanisms” coupled to “a highly parallel execution unit which provides . . . performance and scalability.” *Id.* at 9:9–11.

Figure 4 of Stuttard is reproduced below:




Fig. 4

Figure 4 is a schematic diagram of an MTAP execution unit. *Id.* at 6:32.

The execution unit consists of a number of processing elements (PEs) which allow it to process data elements in parallel. *Id.* at 10:2–3. Each PE consists of an ALU, registers, memory and I/O. *Id.* at 10:3–4. According to Stuttard:

The execution unit can be thought of as two, largely independent, parts (see Figure 4). One PE forms the *mono* execution unit; this is dedicated to processing mono (i.e. scalar or non-parallel) data. The mono execution unit also handles program flow control such as branching and thread switching. The rest of the PEs form the *poly* execution unit which processes parallel (poly) data.

The poly execution unit may consist of tens, hundreds or even thousands of PEs. This array of PEs operates in a synchronous manner, similar to SIMD, where every PE executes the same instruction on its piece of data. Each PE also has its own independent local memory; this provides fast access to the data being processed.

Id. at 10:6–15.

Figure 7 of Stuttard is reproduced below:




Fig. 7

Figure 7 is a block schematic diagram depicting an exemplary MTAP processor. *Id.* at 7:1–2. Figure 7 shows an MTAP “Processor Core” including a single control unit and two execution units: a SIMD (poly) execution unit and a scalar (mono) execution unit. *See id.* at 2:29–34.

2. *Shirazi (Ex. 1036)*

Shirazi relates to quantitative analysis of floating point arithmetic on FPGA based custom computing machines. Ex. 1036, 155. Shirazi teaches that “[m]any algorithms rely on floating point arithmetic for . . . dynamic range” and “require millions of calculations per second.” *Id.* 2-D FFTs (fast Fourier transform) and FIR (finite impulse response) filtering computations are examples of “[i]mage and digital signal processing applications [that] typically require high calculation throughput.” *Id.* According to Shirazi, such computationally intensive algorithms are “candidates for acceleration using custom computing machines (CCMs) . . . tailored for the application.” *Id.* Shirazi states that “[a]n advantage of using a CCM for floating point implementation is the ability to customize the format and algorithm data flow to suit the application’s needs.” *Id.*

Figures 2 and 3 of Shirazi are reproduced below:




Figure 2: 18 Bit Floating Point Format.




Figure 3: 16 Bit Floating Point Format.

Id. at 156. Figures 2 and 3 show two custom formats of Shirazi: (1) an 18-bit format developed for a 2-D FFT (fast Fourier transform) application, and (2) a 16-bit format developed for the FIR (finite impulse response) filter application. *Id.*

3. *The Alleged Obviousness of Independent Claim 1 Over of
Stuttard and Shirazi*

For reasons discussed below, Petitioner has shown a reasonable likelihood that it would prevail in establishing unpatentability of independent claim 1 as obvious over Stuttard and Shirazi.

- a. *[IPRE] A computing system, comprising:*
[1A] a host computer;
[1B] a computing chip comprising:

Petitioner, as discussed further below in the context of the pertinent limitations, proposes a modification with “the PEs in Stuttard’s poly unit [adapted] to perform floating-point math using Shirazi’s custom reduced-bitwidth floating-point values” (Pet. 11), and with the poly unit PEs arranged in a two-dimensional array (*id.* at 19). Petitioner refers to this as the “Stuttard-Shirazi processor.” *See id.* at i (Table of Contents with the heading “The Stuttard-Shirazi Processor” over the descriptions of the two references and of the “Stuttard-Shirazi Combination).

Petitioner contends that “[t]he Stuttard-Shirazi processor communicates with a ‘host system’ including a ‘host processor’ that runs an ‘application or operating system,’” and that Stuttard’s “host” is a host computer, as recited in limitation 1A. Pet. 20 (citing, *inter alia*, Ex. 1003 ¶ 105; Ex. 1006, 1:11–24, 21:13–31). Petitioner further contends that Stuttard discloses that the processor is for use in a number of computing applications by integration into systems, and that “[d]ata and instructions are provided by a host computer,” and, “[t]hus, the Stuttard-Shirazi processor and host computer are integrated in a computing system,” as recited in the preamble. *Id.* (citing Ex. 1003 ¶ 104; Ex. 1006, 1:3–5, 1:21–24, 5:12–14, 9:15–16, 13:21–25). Petitioner additionally contends that “[t]he Stuttard-

Shirazi processor would have been fabricated as a computing chip [as recited in limitation 1B], following Stuttard's teachings that 'illustrate the use of an MTAP processor in a real device' manufactured as a 'chip.'" *Id.* at 21 (citing Ex. 1003 ¶ 109; Ex. 1006, 7:3–4, 22:3–7, 22:12–19) (emphasis omitted).

Patent Owner does not raise any arguments regarding these contentions.

- b. [1C] a processing element array comprising a first edge processing element positioned at a first edge of the processing element array, a second edge processing element positioned at the first edge of the processing element array, a first interior processing element positioned at a first location in the interior of the processing element array, and a second interior processing element positioned at a second location in the interior of the processing element array;*


Limitation 1C calls for a processing element array comprising two edge PEs located at the same (the first) edge of the array and two interior PEs in the interior of the array. Ex. 1001, 30:24–32. The parties characterize this configuration as a two-dimensional, or 2D, array. *See* Pet. 21–23; PO Resp. 2–3.

Petitioner maps the processing elements of Stuttard's "Poly" execution unit to the PEs of the challenged claims. *See* Pet. 14–15, 21–23. Petitioner asserts that Stuttard discloses that the poly unit PEs are oriented linearly, i.e., one-dimensionally. *See id.* at 15. For example, Petitioner asserts that "Stuttard's Figures 4 and 7 illustrate poly PEs side-by-side because each PE communicates with its left and right neighbors in what

Stuttard calls a ‘swazzle’ path.”¹¹ *Id.* (citing Ex. 1006, 20:31–21:2; Ex. 1003 ¶¶ 96–97). Petitioner contends that this disclosure reflects a conceptional programming model and that persons of ordinary skill in the art “would have understood that it does not limit the physical layout of the processor on a chip.” *Id.* (citing Ex. 1006, 13:2–5; Ex. 1003 ¶ 97). Petitioner further contends that persons of ordinary skill in the art “would have understood that the physical layout was a matter of design choice and . . . had reasons to manufacture a physical chip using a two-dimensional PE layout.” *Id.* at 15–16 (citing Ex. 1003 ¶ 97).

Petitioner first contends that persons of ordinary skill in the art “knew that a two-dimensional layout reduces the overall length of wiring in the chip, such as Stuttard’s ‘channels’ that transfer data between PEs and the host.” *Id.* at 16 (citing Ex. 1006, 17:20–23, 23:26–32; Ex. 1003 ¶ 98). According to Petitioner, persons of ordinary skill in the art “would have understood that physically arranging PEs in a two-dimensional array reduces the overall length of each channel, beneficially reducing (i) signal delays, (ii) resistances (and thus heat) and (iii) undesirable ‘parasitic’ capacitances.” *Id.* (citing Ex. 1003 ¶ 98). To assist in explaining its position, Petitioner provides the demonstrative figures reproduced below.

¹¹ Stuttard explains that a “*swazzle* path . . . connects the register file of each PE with the register files of its left and right neighbours.” Ex. 1006, 20:31–34.



Id. at 16–17 (citing Ex. 1003 ¶ 101). Above are depicted linear and two-dimensional layouts of PEs with “[t]he red line in each figure represent[ing] Stuttard’s channel connected to each PE [and] the green lines represent[ing] Stuttard’s ‘swazzle’ connections.” *Id.* at 16. Petitioner contends that “both layouts allow for a swazzle path connecting each PE to its linearly adjacent neighbors, but the channel’s length in the 2D array can beneficially be approximately half as long as in the 1D line.” *Id.* (citing Ex. 1003 ¶ 99).


Petitioner next contends that persons of ordinary skill in the art “knew it was desirable and conventional for computing chips to have aspect ratios (width-to-length ratios) closer to 1, because very high or very low aspect ratios cause ‘mechanical stress[es]’ and ‘adversely change the electrical properties of the IC circuitry.’” *Id.* at 17 (citing Ex. 1067 ¶ 7; Ex. 1003 ¶ 100) (alteration in original).

Petitioner, relying on the testimony of Dr. Franzon, next contends that “other related prior art further corroborates that a two-dimensional interconnect topology (and corresponding two-dimensional physical


topology) was a conventional alternative to a linear topology.” *Id.* Petitioner points to, as an example, Exhibit 1108 and asserts that it “corroborates the POSA’s knowledge that ‘as the number of PE cores increases **other (than linear)** interconnect topologies could be introduced.”” *Id.* (quoting Ex. 1108, 35; citing Ex. 1003 ¶ 101). Petitioner also points to Exhibit 1081, asserting that it teaches a processor “having ‘a processing element (PE) array’ with ‘PE interconnections’ that move data ‘along the PE array (hereinafter swazzled),’ and explains the array’s ‘PE interconnection’ can be arranged in ‘a 1D ring/line array **or a 2D grid array.**”” *Id.* at 18 (quoting Ex. 1081, 5:51–65); *see also id.* (quoting 1081, 8:40–47 (“Two interconnection networks are generally used: a 1D array/ring and a 2D grid.”)). According to Petitioner, Exhibit 1081 “further taught that ‘**a 2D array** with a two dimensional bi-directional (up, down, left, right) nearest-neighbor configuration, often known as a mesh’ was a ‘common configuration’ and sometimes ‘preferable to exploit the 2D connectivity.’” *Id.* (quoting Ex. 1081, 21:10–22). Petitioner reasons that, “[t]hus, using a 2D-array in Stuttard-Shirazi would have been using a known technique to improve similar devices (parallel processor having multiple PEs) in the same way (by facilitating applications that benefited from full-2D connectivity between PEs).” *Id.* (citing *KSR*, 550 U.S. at 417; Ex. 1003 ¶ 102).

Petitioner contends that applying these principles to its proposed combination of Stuttard and Shirazi would result in a two-dimensional layout of the Stuttard poly PEs. *Id.* at 19 (citing Ex. 1003 ¶ 103). Reproduced below is Stuttard’s Figure 7 alongside Petitioner’s demonstrative figure depicting its proposed modification having a two-dimensional layout.

Stuttard, FIG. 7



Stuttard-Shirazi



Pet. 22; *see also id.* at 19 (the same or similar figures without the color annotations). Above, on the left, is Stuttard's Figure 7, which is a block schematic diagram depicting an exemplary MTAP processor (Ex. 1006, 7:1–2) and, on the right, is Petitioner's demonstrative figure depicting a two-

dimensional array of the poly PEs of the Stuttard-Shirazi processor (Pet. 19, 21). The figure on the right contains Petitioner’s annotations depicting, according to Petitioner, first and second edge PEs (red and orange, respectively) and first and second interior PEs (blue and green, respectively). *Id.* at 22–23.

Patent Owner asserts that Petitioner has not set forth an adequate reason as to why a person of ordinary skill in the art would have modified Stuttard from a one-dimensional array to a two-dimensional array. *See* Prelim. Resp. 16–24.

Patent Owner first argues that Petitioner’s cited passages from Exhibits 1081 and 1108 “show nothing about whether a POSA would modify Stuttard *in particular* into such [a two-dimensional] arrangement.” Prelim. Resp. 17–18 (citing Pet. 17–18). Patent Owner argues, for example, that it is speculative that “other than linear” suggests a two-dimensional array or that “[p]referably, a 1D line array is provided” suggests the use of a 2D array. *See id.* (citing Pet. at 17 (citing Ex. 1108, 35; Ex. 1081, 5:65)). We do not find those arguments persuasive on the current record. Dr. Franzon opines that “a two-dimensional interconnect topology (and corresponding two-dimensional physical topology) would have been a conventional alternative to a linear topology,” and testifies that Exhibits 1081 and 1108 further corroborate that opinion. Ex. 1003 ¶¶ 101–102; *see, e.g.*, Ex. 1081, 8:46–47 (“Two interconnection networks are generally used: a 1D array/ring and a 2D grid.”). At this stage, that opinion is un rebutted.

Patent Owner notes that David Stuttard is listed as an author on Exhibit 1108 and as an inventor on Petitioner’s primary reference

(Ex. 1006). Prelim. Resp. 18. Patent Owner argues that Petitioner failed to address why, if Mr. Stuttard was aware of 2D arrays, he did not mention a 2D arrangement in the primary reference. *Id.* We do not find this persuasive because the obviousness analysis seeks to determine what a person of ordinary skill in the art would have found to be obvious based on the teachings in the references, not what David Stuttard thought. Also, we know of no requirement that everything known to an inventor, for instance Mr. Stuttard here, should be disclosed or described in a patent specification.

Patent Owner next argues:

Google also does not address that even its own evidence at times teaches away from using a 2D arrangement. Pet. at [17–18] (citing Ex. 1108, 1081). Indeed, generalizing the array to 2D requires much engineering, rewiring, area, and power increases, and all this would require significant experimentation which Google does not flag, let alone address. *See* Pet. at 14–19. A POSA would not be motivated to conduct this exercise.

Prelim. Resp. 18. We do not, at this stage, find these arguments to be persuasive. Patent Owner does not identify with any specificity how Petitioner’s evidence purportedly teaches away from using a two-dimensional arrangement. *Id.*; *see In re Ethicon, Inc.*, 844 F.3d 1344, 1351 (Fed. Cir. 2017) (A reference’s statement, “absent clear discouragement from use, does not compel a finding that [the reference] teaches away from using [a technical feature].”). Patent Owner’s assertion that “generalizing the array to 2D . . . would require significant experimentation” is, at this point, merely attorney-argument on a highly technical subject.

As mentioned above, Petitioner, after discussing Exhibits 1081 and 1108, contends that “using a 2D-array . . . would have been using a known technique to improve similar devices . . . in the same way (by facilitating applications that benefited from full-2D connectivity between

PEs).” Pet. 18. Patent Owner asserts that, “[i]n [Petitioner’s] words, ‘full-2D connectivity’ means that every PE is *directly* connected to the four PEs surrounding it, *i.e.*, there exist ‘up, down, left, right’ connections between PEs.” Prelim. Resp. 19 (citing Pet. 18) (emphasis added). Patent Owner also reproduces Petitioner’s demonstrative figure depicting a two-dimensional layout with channels between up and down neighbors, and Patent Owner characterizes that figure as Petitioner’s “proposed alteration.” *Id.* at 19–20 (citing Pet. 17); *see* Pet. 16 (“The red line in each figure represents Stuttard’s channel connected to each PE.”). Patent Owner argues that Petitioner’s “combination, having only a shared bus to effect ‘up’ and ‘down’ connections, would lack the *direct* ‘up’ and ‘down’ connections that form its primary reason for making the alteration.” Prelim. Resp. 20. Patent Owner’s arguments are not persuasive on the current record. It is not clear that Patent Owner is correctly characterizing Petitioner’s reasoning because the Petition, on the cited page 18, does not contain the word “direct,” and Patent Owner does not explain adequately why it understands Petitioner’s reasoning to hinge on direct connections.

Patent Owner argues “even though Google states that a POSA would have made the modification of adopting a 2D topology to facilitate applications that ‘benefited from full-2D connectivity,’ Google fails to identify any such applications.” Prelim. Resp. 20 (citing Pet. at 18; Ex. 1003 ¶ 102). Although it is not clear what Patent Owner would expect of Petitioner in articulating an obviousness challenge, we note that Petitioner’s mention of “facilitating applications” follows the discussion of Exhibits 1081 and 1108. *See* Pet. 17–18; *see also* Ex. 1003 ¶ 102 (Dr. Franzon: “*Thus, EX1108 and EX1081 demonstrate that using a 2D-*

array in Stuttard-Shirazi would have been using a known technique (2D-mesh) to improve similar devices (parallel processor having an array of PEs) in the same way (*by facilitating applications that benefited from full-2D connectivity between PEs*).” (emphasis added)).

As mentioned, Patent Owner asserts that the ’775 patent recognized a data transfer bottleneck problem. *See* Prelim. Resp. 20. Patent Owner further asserts that the ’775 patent provides a solution by avoiding a shared bus and by connecting the I/O unit to only the edge PEs, i.e., precluding interior PEs from being connected directly to the I/O unit. *See id.* at 20–21. Patent Owner argues that Petitioner has failed to explain how to prevent its proposed combination, having a shared bus, from creating a bottleneck. *See id.* at 20. This argument is not persuasive as we, for the reasons discussed above, construe the independent claims as *not* precluding the interior PE elements from being connected directly to the input-output unit and as *not* precluding the use of a shared bus. *See supra* Section III.C.2. To the extent that Patent Owner argues that Petitioner’s proposed combination would contain a bottleneck that limits the speed of data transfers in and out of the PE array or would require “absorbing large amounts of hardware resources” (Prelim. Resp. 20), we do not find that persuasive on the record as it now stands. “[A] given course of action often has simultaneous advantages and disadvantages, and this does not necessarily obviate motivation to combine.” *Medichem, S.A. v. Rolabo, S.L.*, 437 F.3d 1157, 1165 (Fed. Cir. 2006). It is not necessary that Petitioner’s proposed modification be the best way, in all respects, to operate an array of PEs. The existence of a “bottleneck” does not render Petitioner’s proposed array of PEs inoperable. Further, we find that on the record before us at this stage of the proceeding, the evidence is

insufficient to show that the resulting bottleneck would be so severe that it outweighs Petitioner's reasons to make the modification and thus one of ordinary skill in the art would not make the modification.

The Petition contains two demonstrative figures (reproduced *supra*) that are used to illustrate Petitioner's contention that that "physically arranging PEs in a two-dimensional array reduces the overall length of each channel" as compared to a one-dimensional array. Pet. 16–17. Patent Owner argues that there is no evidence that the demonstrative figures "accurately represent the relative scale and relationship between the PEs."

Prelim. Resp. 22. Patent Owner further argues that:

[Petitioner] Google apparently asks the Board to take its pictorial representation of a Stuttard chip as somehow representative of the actual dimensions of the PEs in the chip itself, measure features in its not-to-scale Figure (*i.e.*, as opposed to in the chip) to ascertain dimensions on the chip, and finally compare the unspecified lengths of the channels illustrated within those Figures to one another. The Board should decline to do so.

Id. We do not understand Petitioner to ask us to perform such acts. Further, Dr. Franzon testifies that the demonstrative figures illustrate his opinions regarding the benefits of a two-dimensional array over a one-dimensional array, and specifically that "the overall length of the channel in the 2D array can be approximately half as long as the length of the channel in the 1D line." Ex. 1003 ¶ 99. Dr. Franzon also testifies that a person of ordinary would know of these benefits and thus would be motivated to use a 2D array. *See id.* ¶ 98. At this stage of the proceeding, that testimony is un rebutted.

Patent Owner argues that Petitioner's reasoning based on the desirability for computing chips to have aspect ratios closer to one is flawed

and does not provide a motivation to use a two-dimensional PE array. Prelim. Resp. 22–23 (citing Pet. 17; Ex. 1003 ¶ 100). Patent Owner contends that a discussion of a chip’s aspect ratio “says nothing about the topology of an array of PEs” because a chip can hold multiple functional blocks other than PEs. *Id.* Even if Patent Owner is correct in arguing that Petitioner’s position regarding aspect ratios is flawed, that does not undermine Petitioner’s other reasoning, including that a person of ordinary skill in the art was aware of a two-dimensional array as an alternative to a one-dimensional array and that the use of a two-dimensional array beneficially reduces the overall wiring length. *See* Pet. 16–18.

Lastly, Patent Owner argues that Petitioner’s proposed combination fails to contain any interior processing elements as recited in the claims. Prelim. Resp. 25–26. Patent Owner contends that Petitioner’s “asserted interior PEs are still connected to what [Petitioner] Google identifies as the input-output unit” and, because “an interior PE cannot be connected to the input-output unit[,] . . . all of the PEs in Google’s combination, even if reconfigured as a 2D array, are ‘edge processing elements;’ indeed, none of them are ‘interior’ in any respect.” *Id.* For the reasons discussed above, we construe the independent claims as *not* precluding the interior PE elements from being connected directly to the input-output unit and as *not* precluding the use of a shared bus. *See supra* Section III.C.2.

Limitation 1C requires the recited first and second interior processing elements to each be in a specified position—respectively, “at a first location in the interior of the processing element array, and . . . at a second location in the interior of the processing element array.” Ex. 1001, 30:28–32. Petitioner has explained adequately how the PEs mapped to the interior PEs

are so positioned in the proposed combination. Pet. 22. Specifically, whether a PE is in the interior of the array does not depend on whether it is directly connected to an input-output unit. Rather, a PE is in the interior of the array if it is not located in the outer most column (an edge) or the outer most row (also an edge) of the array.

We determine that Petitioner has articulated a sufficient reason, for purposes of this Decision, why one of ordinary skill in the art would have modified Stuttard to have a two-dimensional array configuration, and that Petitioner has demonstrated sufficiently, for purposes of this Decision, that the proposed modification has edge and interior processing elements positioned as recited in limitation 1C.

- c. [1D] a first processing element connection connecting the first edge processing element with the first interior processing element;*
- a second processing element connection connecting the second edge processing element with the second interior processing element;*

Petitioner notes that Stuttard discloses that processing elements “communicate with one another via what is known as the *swazzle* path that **connects** the register file of **each PE with** the register files of **its left and right neighbours.**” Pet. 23 (quoting Ex. 1006, 20:31–33). Petitioner maps these swazzle paths to the recited connections between pairs of edge and interior PEs. *Id.* at 23–24.

Patent Owner does not present arguments regarding limitation 1D beyond those addressed above. *See, e.g.,* Prelim. Resp. 24–26 (Patent Owner arguing that there is no motivation to use a 2D array and, because all the PEs in Petitioner’s proposed modification have certain connections in addition to the swazzle paths, there are no interior PEs).

We determine that Petitioner has demonstrated sufficiently, for purposes of this Decision, that the proposed modification has the processing element connections between edge and interior processing elements as recited in limitation 1D.

d. [1E] an input-output unit connected to the first edge processing element and the second edge processing element;

Petitioner contends, *inter alia*, that Stuttard's bus is the recited input-output unit. Pet. 30. Patent Owner argues that a bus simply is a set of conductors, and that a person of ordinary skill in the art would not understand a bus to be an input-output unit within the meaning of the '775 patent. Prelim. Resp. 26–27. As discussed above in the claim construction section, the Specification indicates that the input-output unit has functionality beyond merely serving as a connector between components. *See supra* III.C.1. Petitioner does not explain adequately how Stuttard's bus has such functionality beyond being a connector.

However, Petitioner contends, in the alternative, that Stuttard's "bus plus the I/O engines and channels" is the recited "input-output unit." Pet. 30. Petitioner asserts that "Stuttard's Figure 7 shows that [Programmed I/O (PIO) and Streaming I/O (SIO)] interfaces connect the bus to respective PIO/SIO 'engines,' which connect to every PE in the poly-unit array via the 'channels.'" *Id.* at 27–28 (citing Ex. 1003 ¶¶ 117–118; Ex. 1006, 13:6–16, 19:12–20, 19:22–31, Fig. 7). Petitioner also asserts that the PIO and SIO interfaces "are for PEs to exchange data with off-chip devices (e.g., the host)." *Id.* at 27. Patent Owner, acknowledges Petitioner's alternative contention but does not argue against it (beyond, for example, its argument, addressed above, that the claims preclude the use of a shared bus). *See*

Prelim. Resp. 26–27 (arguing that a bus alone lacks certain abilities but not making the same argument for the engines and channels).

Stuttard explains that “[t]he [PIO and SIO] engines are basically DMA engines which *manage the actual data transfer*.” Ex. 1006, 19:22–29 (emphasis added). This suggests that the structure mapped to the I/O unit—the bus plus the I/O engines and channels—has a capability beyond merely serving as a connector and has a role in inputting and outputting data. Accordingly, we determine that Petitioner has demonstrated sufficiently, for purposes of this Decision, that the proposed modification has the input-output unit recited in limitation 1E.

- e. [1F] a first memory local to the first edge processing element;*
a second memory local to the second edge processing element;
a third memory local to the first interior processing element;
a fourth memory local to the second interior processing element; and,

Petitioner contends that “Stuttard’s Figure 4 . . . illustrates that ‘[e]ach PE [] has its own independent **local memory**’ and addressable ‘register file,’ each of which is memory local to the PE, as claimed.” Pet. 30 (citing Ex. 1003 ¶ 121; Ex. 1006, 5:10–11, 5:29–30, 6:18–20, 7:14–16, 10:1–4, 10:14–19, 17:12–20, FIGS. 3–4, 7) (alterations in original except ellipses).

Patent Owner does not raise any arguments regarding this contention.

- f. [1G] a fifth arithmetic unit;*

Petitioner contends that “Stuttard’s mono unit has an ALU (**arithmetic logic unit**), which is a fifth arithmetic unit, as claimed.” Pet. 37 (citing Ex. 1003 ¶¶ 122, 132; Ex. 1006, 3:15–29, 7:14–24, 16:13–20, Figs. 3, 4, 7).

Patent Owner does not raise any arguments regarding this contention.

- g. *[1H] wherein the first edge processing element comprises a first arithmetic unit;*
wherein the second edge processing element comprises a second arithmetic unit;
wherein the first interior processing element comprises a third arithmetic unit; and
wherein the second interior processing element comprises a fourth arithmetic unit; and,

Petitioner contends that “[a]s Stuttard’s Figure 4 . . . illustrates, ‘*[e]ach PE*’ comprises ‘an ALU’ (*arithmetic logic unit*).” Pet. 31 (citing Ex. 1003 ¶ 124; Ex. 1006, 5:28–29, 7:14–16, 10:1–4, 10:14–19, 17:12–23, claims 9–10, Figs. 3–4, 7) (last alteration in original).

Patent Owner does not raise any arguments regarding this contention.

- h. *[1I] a host connection at least partially connecting the input-output unit with the host computer;*



Petitioner contends that Stuttard’s bus has external ports (Ports 0 and 1 in Figure 8) “connect[ed] to an FPGA [field programmable gate array] to provide other functions, such as . . . a *host interface*.” Pet. 32 (quoting Ex. 1006, 22:8–19 (alterations in original); citing Ex. 1003 ¶ 124; Ex. 1006, 21:13–15 (processor has “interfaces to the host system”)). Petitioner further contends that “the external ports . . . are one or more host connections at least partially connecting the input-output unit (including the bus) with the host computer, as claimed.” *Id.* at 33 (citing Ex. 1003 ¶ 128).

Patent Owner does not raise any arguments regarding these contentions.

- i. *[1J] wherein the first, second, third and fourth arithmetic units each comprises a corresponding multiplier circuit adapted to receive as a first input to the corresponding multiplier circuit a first floating point value having a first binary mantissa of width no more than 11 bits and a first binary exponent of width at least 6 bits, and to receive as a second input to the corresponding multiplier circuit a second floating point value having a second binary mantissa of width no more than 11 bits and a second binary exponent of width at least 6 bits;*

Petitioner contends that each PE in Stuttard's poly unit includes an arithmetic logic unit (ALU) and that "each ALU has an FPU (floating-point unit), which comprises circuits for performing floating-point operations, including floating-point multiplication." Pet. 34 (citing Ex. 1003 ¶ 129; Ex. 1006, 12:7–15, 13:21–28, 17:24–33).

Petitioner relies on Shirazi for the recited floating point values (Pet. 11), which the parties refer to as reduced bitwidth floating point values or formats. *See id.*; Prelim. Resp. 13. Petitioner contends that Shirazi discloses two custom reduced bitwidth formats using the annotated figures reproduced below.



Pet. 10. Above are Petitioner’s annotated versions of Shirazi’s Figures 2 and 3, which show two formats: (1) an 18-bit format developed for a 2-D FFT (fast Fourier transform) application; and (2) a 16-bit format developed for the FIR (finite impulse response) filter application. Ex. 1036, 156.

According to Petitioner, annotated Figures 2 and 3, respectively, show:

- (1) an 18-bit format developed for “2-D FFT application,” having one sign bit, 7 exponent bits [blue], and 10 mantissa bits [yellow]; and
- (2) a 16-bit format developed for “the FIR filter application,” having one sign bit, 6 exponent bits [blue], and 9 mantissa bits [yellow].

Pet. 9–10 (citing Ex. 1003 ¶ 80); *see* Prelim. Resp. 15 (“Shirazi seeks to use reduced bit width floating point formats.”). Petitioner contends that, in the proposed combination, “the multiplier circuit in each poly PE [of Stuttard] is adapted to receive as inputs two of Shirazi’s reduced-bitwidth floating-point values.” Pet. 36 (citing Ex. 1003 ¶ 131).

Petitioner contends that persons of ordinary skill in the art “had reasons to adapt the floating-point units in Stuttard’s poly unit to have circuitry that accepts as inputs, and operates on, floating-point values having bitwidths no greater than one of Shirazi’s reduced-precision floating-point formats.” *Id.* at 13 (citing Ex. 1003 ¶ 94). Petitioner asserts that:

Shirazi teaches that “[m]any algorithms rely on floating point arithmetic for the[ir] dynamic range” and “require millions of calculations per second.” Shirazi, 155. For example, 2-D FFTs (Fast Fourier Transforms) and FIR (finite-impulse-response) filtering computations are commonly used in “[i]mage and digital signal processing applications [that] typically require high calculation throughput.” *Id.* These applications are, therefore, “candidates for acceleration using custom computing machines (CCMs) [] tailored for the application” (*id.*)—i.e., the type of machine Stuttard describes.

Pet. 9 (citing Ex. 1003 ¶ 80) (alterations in original). Petitioner also contends that “Shirazi taught that its floating-point formats having reduced-bitwidth mantissas provided sufficient precision for common high-data-rate applications, while also allowing the arithmetic circuitry to be implemented with reduced chip ‘area’ compared to full-precision circuits.” *Id.* at 11 (citing Ex. 1003 ¶ 83; Ex. 1036, 155–156). Petitioner asserts that this was well known in the art, and that, “[f]or example, EX1034 (‘Tong’) similarly taught that ‘[i]t has **long been known** that many . . . applications can get by **with less precision**’ (Tong, 273), and thus the IEEE’s single-precision (32-bit) format often ‘is not essential’ (Tong, 279).” *Id.* (citing Ex. 1003 ¶ 83) (alteration in original). Petitioner additionally asserts that “Tong also corroborates that POSAs knew that minimizing the number of mantissa bits reduces each PE’s power consumption because a ‘reduced precision FP unit’ has ‘small size’ compared to higher-precision units.” *Id.* (citing Ex. 1003 ¶ 83; Ex. 1034, 282).

Petitioner contends that one motivation to make the proposed combination is that “Stuttard’s MTAP processor architecture is for ‘high performance, high data rate processing’ applications,” and persons of ordinary skill in the art “would have understood this architecture to be well-suited for signal-processing applications like the 2-D FFTs and FIR filtering computations that Shirazi taught were compatible with reduced-precision floating-point formats.” *Id.* at 12–13 (citing Ex. 1003 ¶¶ 91–92; Ex. 1006, 8:24–28; Ex. 1036, 155–156). Petitioner continues: “Stuttard’s poly unit ‘gets its performance largely from the number of PEs brought to bear on the problem rather the processing power of the individual PEs,’ which ‘need to be efficiently implemented,” and “[t]his would have further motivated POSAs to pursue techniques like Shirazi’s reduced-bitwidth floating-point formats, which would have allowed for more PEs on the chip.” *Id.* at 13 (citing Ex. 1003 ¶ 93; Ex. 1006, 7:19–24; Ex. 1036, 155–156).

Petitioner, relying on the testimony of Dr. Franzon, contends that a person of ordinary skill in the art would have been able to make the proposed combination with a reasonable expectation of success. *Id.* at 14 (citing Ex. 1003 ¶ 95).

Patent Owner argues that Petitioner has failed to demonstrate a reason why a person of ordinary skill in the art would have been motivated to make the proposed combination. Prelim. Resp. 11. Patent Owner first asserts that Petitioner has failed to show that the “discovery” of the ’775 patent was known by persons of ordinary skill in the art prior to its discovery by the named-inventor. *See id.* at 11–13. According to Patent Owner, “a central advance of the ’775 Patent” was “the discovery that large amounts of LPHDR [Low Precision High Dynamic Range] arithmetic performed by

multiple low precision processing elements in a massively parallel architecture are, in fact, useful and provide significant practical benefits.” *Id.* at 11 (citing Ex. 1001, 7:3–38); *see id.* at 12 (the discovery “that large numbers of low precision and high dynamic range calculations could provide useful results for *any application*.”). The heading for this section of the Preliminary Response indicates that Patent Owner regards 5,000 PEs as the referenced “large amounts of LPHDR” PEs. *See id.* at 11. Patent Owner argues that Petitioner’s “references say nothing about scale in the number of processing elements performing imprecise LPHDR arithmetic in parallel, with all processing elements repeatedly performing imprecise LPHDR arithmetic.” *Id.* at 12.

Patent Owner asserts that the architecture related to this discovery is “covered by the Challenged Claims.” *Id.* at 12. However, Patent Owner, in this section of the Preliminary Response, cites only to the Specification and does not explain adequately how the “discovery” is embodied in the limitations of any claim. For example, although independent claims 7 and 16 recite an array of at least 5,000 PEs (Ex. 1001, 31:57–60, 33:33–35), Patent Owner does not point to any limitation in independent claim 1 directed to a large number of PEs or of a massively parallel architecture. Accordingly, we do not find Patent Owner’s argument about the “discovery” persuasive at least as to independent claim 1.

Patent Owner next argues that it “would not have been obvious for a POSA to adapt Stuttard to use Shirazi’s reduced bit width formats because doing so would not have furthered any expressed goal of the Stuttard system.” Prelim. Resp. 13. In characterizing that goal, Patent Owner asserts that “the Stuttard architecture is intended and designed to address

programming problems, not the sort of performance problems that could be addressed by Shirazi.” *Id.* at 14 (citing Ex. 1006, 1:25–2:8). Patent Owner argues that Shirazi’s reduced bitwidth formats “in no way make the Stuttard system more programmable, meaning a [person of ordinary skill in the art] would not have been motivated to look to Shirazi to improve the Stuttard system.” *Id.* at 15. Rather, according to Patent Owner, “[i]n keeping with the stated purpose of the Stuttard system, a POSA looking to improve it would have sought out improvements that make its SIMD array controller more programmable.” *Id.* at 14–15.

On the current record, we do not find to be persuasive Patent Owner’s implied argument that a person of ordinary skill in the art would look only to possible solutions to Stuttard’s programming problems. “A person of ordinary skill is also a person of ordinary creativity, not an automaton.” *KSR*, 550 U.S. at 421; *see id.* at 420–421 (holding that the lower court erred in concluding that, because “[t]he primary purpose of Asano was solving the constant ratio problem,” “an inventor considering how to put a sensor on an adjustable pedal would have no reason to consider putting it on the Asano pedal”). It is improper to assume the person of ordinary skill in the “will be led only to those elements of prior art designed to solve the same problem.” *Id.* at 420. On the record as it now stands, we are not persuaded that a person of ordinary skill in the art would fail to see the applicability of Shirazi’s reduced bitwidth teachings to Stuttard’s system, which, according to Petitioner, is used in high performance, high data rate processing. *See* Pet. 12–13. Dr. Franzon testifies that persons of ordinary skill in the art would have “understood that Stuttard’s processor was well-suited to the types of computations Shirazi taught were well-suited to reduced-bitwidth

floating point units (e.g., FFTs).” Ex. 1003 ¶ 92. Dr. Franzon discusses exhibits that are characterized as corroborating that opinion. *Id.* ¶¶ 91–92. At this stage of the proceeding, that testimony is unrebutted.

We determine that Petitioner has articulated a sufficient reason, for purposes of this Decision, why one of ordinary skill in the art would have modified Stuttard to use Shirazi’s reduced bitwidth floating point formats, and that Petitioner has demonstrated sufficiently, for purposes of this Decision, that the proposed modification has the multiplier circuit recited in limitation 1J.

- j. [1K] wherein the fifth arithmetic unit comprises a corresponding multiplier circuit adapted to receive as inputs to the corresponding multiplier circuit two floating point values each of width at least 32 bits;*

As mentioned above, Petitioner maps the recited “fifth arithmetic unit” (limitation 1G) to the arithmetic logic unit (ALU) of Stuttard’s mono unit. *See* Pet. 37. Petitioner contends that “[t]his ALU performs ‘floating point...multiplies’ (Stuttard [Ex. 1006], 13:26-28) using ‘extensions’ that ‘include a **multiplier**’ (Stuttard, 16:13-20), and can be ‘a 32-bit ALU’ (Stuttard 3:15-18).” *Id.* at 38. Petitioner further contends that persons of ordinary skill in the art “would have understood that Stuttard thus teaches providing a 32-bit floating-point multiplier circuit in the mono unit.” *Id.* (citing Ex. 1003 ¶ 132). Petitioner additionally contends that persons of ordinary skill in the art would have had reason “to include in Stuttard-Shirazi’s mono unit an FPU [floating-point unit] that supported higher bitwidths than the FPUs in the poly unit,” and would have reason to use the IEEE’s 32-bit values. *Id.* at 38–39 (citing Ex. 1003 ¶¶ 133–134; Ex. 1006, 3:15–18, 7:14–24).

Patent Owner does not raise any arguments regarding these contentions.

- k. [1L] wherein the multiplier circuit corresponding to the first arithmetic unit comprises a first plurality of transistors and has no other transistors, the multiplier circuit corresponding to the second arithmetic unit comprises a second plurality of transistors and has no other transistors, the multiplier circuit corresponding to the third arithmetic unit comprises a third plurality of transistors and has no other transistors, the multiplier circuit corresponding to the fourth arithmetic unit comprises a fourth plurality of transistors and has no other transistors, and*

This limitation 1L recites that each of the first four multiplier circuits “comprises a . . . plurality of transistors and has no other transistors.”

Ex. 1001, 31:8–21. Petitioner contends that “[b]y the 2000s, floating-point multiplier circuits were conventionally implemented using transistors.

Pet. 39 (citing Ex. 1003 ¶ 136). Petitioner additionally contends:

Because each PE in the Stuttard-Shirazi poly unit has a floating-point multiplier for multiplying values no wider than Shirazi’s 18-bit (or 16-bit) floating-point values . . . , POSAs would have designed each multiplier circuit to have sufficient transistors to support inputting and multiplying two of Shirazi’s 18-bit (or 16-bit) values. Franzon [Ex. 1003], ¶137. Only the transistors necessary to support those maximum input widths supported by the multiplier would have been included in the circuit. Franzon, ¶137

Id.

Patent Owner does not raise any arguments regarding these contentions.

- l. [1M] the multiplier circuit corresponding to the fifth arithmetic unit comprises a fifth plurality of transistors; and, [1N] wherein the fifth plurality of transistors exceeds in number each of the first plurality of transistors, the second plurality of transistors, the third plurality of transistors, and the fourth plurality of transistors.*

Petitioner contends that persons of ordinary skill in the art knew that the fifth ALU, like that of the first through fourth ALUs (limitation 1L), comprises a plurality of transistors. Pet. 40 (citing Ex. 1003 ¶ 138). Petitioner reiterates its assertion that “Stuttard-Shirazi’s first through fourth multiplier circuits comprise the number of transistors needed to multiply two input values no wider than Shirazi’s 18-bit (or 16-bit) floating-point numbers.” *Id.* (citing Ex. 1006, 155–156). Petitioner further contends:

Similarly, the fifth multiplier circuit (in the mono unit) has the number of transistors needed to receive and multiply two floating-point values at least 32 bits wide. . . . POSAs would have understood that a multiplier circuit adapted to receive and multiply two 32-bit-wide floating-point values would conventionally require more transistors than a similarly architected multiplier circuit adapted to receive and multiply two lower-bitwidth floating-point values. Franzon, ¶141 (citing EX1076, 5:21-26 (corroborating POSA’s knowledge that “the size of the multiplier is proportional to the number of bits multiplied”)).

Id. at 40–41.

Patent Owner does not raise any arguments regarding these contentions.

m. Conclusion as to Independent Claim 1

Having considered the parties’ arguments and evidence, we determine that Petitioner has demonstrated a reasonable likelihood of prevailing in its challenge to claim 1 as being obvious over Stuttard and Shirazi.

4. The Alleged Obviousness of Dependent Claims 2–5 Over Stuttard and Shirazi

Claims 2 through 5 depend directly or indirectly from independent claim 1. Ex. 1001, 31:26–49. Petitioner presents contentions as to how each limitation of those claims is taught or suggested by the combination of Stuttard and Shirazi. Pet. 41–50. Patent Owner does not raise any arguments regarding these contentions. We have reviewed Petitioner’s arguments and evidence in support of its contentions and determine that Petitioner has demonstrated a reasonable likelihood of prevailing in its challenge to dependent claims 2–5 as being obvious over Stuttard and Shirazi.

E. The Alleged Obviousness of Claims 7–14 and 16–23 Over Stuttard and Shirazi (Ground 1B)

Independent claim 7 recites “a computing chip” with that chip comprising “a processing element array comprising a plurality of first processing elements, wherein the plurality of first processing elements is no less than 5000 in number.” Ex. 1001, 31:56–60. Independent claim 16 recites “[a] circuit board” with each of its plurality of processing element arrays comprising “a plurality of first processing elements, wherein the plurality of first processing elements of said processing element array is no less than 5000 in number.” *Id.* at 33:30–35.

Petitioner asserts that persons of ordinary skill in the art “had reasons to scale the Stuttard-Shirazi processor in a manner that would have resulted in over 5000 reduced-bitwidth PEs on one computing chip.” Pet. 51 (citing Ex. 1003 ¶ 166). Petitioner does not point to any disclosure in the prior art of the use of 5,000 PEs. *See* Pet. 51–55, 59–65. The closest evidence that Petitioner points to is Stuttard’s statements that “[t]he poly execution unit

may consist of tens, hundreds or even thousands of PEs” and “the core can be scaled well beyond this number of [64] PEs, to 256 or even thousands of PEs.” Ex. 1006, 10:12, 12:8–16 (cited at Pet. 51). Petitioner contends that a person of ordinary skill in the art would have reasonably expected success in implementing over 5,000 PEs. *See* Pet. 64. Petitioner also contends that “the number of PEs recited in claim 7 carries no patentable weight, because ‘[i]t is well settled that the mere duplication of parts has no patentable significance unless a new and unexpected result is produced.’” *Id.* at 53–54 (quoting *In re Harza*, 274 F.2d 669, 671 (CCPA 1960)). Petitioner additionally contends that “selecting the number of PEs in a Stuttard-Shirazi chip would have been routine optimization of a result-effective variable, such that arriving at 5000+ PEs would have been obvious.” *Id.* at 55 (citing, *inter alia*, *In re Applied Materials, Inc.*, 692 F.3d 1289, 1295 (Fed. Cir. 2012)).

Patent Owner argues, *inter alia*, that Petitioner’s contentions are based on full-precision PEs, not the large number of *low precision* PEs as claimed. *See* Prelim. Resp. 28–29. Patent Owner also disagrees with Petitioner’s assertion that the limitation gets no weight, asserting that the Petition lacks affirmative evidence that no new and unexpected results are produced and merely incorrectly relies on the alleged silence in the Specification of the ’775 patent. *See id.* (citing Pet. 53–54; Ex. 1003 ¶ 170).

We question the strength of Petitioner’s case. Petitioner extrapolates from Stuttard’s seemingly aspirational statement of scaling to “even thousands,” appears to focus on the obviousness of scaling Stuttard’s full precision PEs, and contains little or no discussion of a motivation to scale up

to at least 5,000 *low precision* PEs or of a reasonable expectation of success of scaling up to 5,000 low precision PEs.

However, for the reasons discussed above in the context of Ground 1A, Petitioner has shown a reasonable likelihood that at least one claim of the '775 patent is unpatentable. Having decided to institute as to at least one challenged claim, we institute on all challenged claims and the entirety of the sole asserted ground (designated as Grounds 1A and 1B). *SAS Inst. Inc. v. Iancu*, 138 S. Ct. 1348, 1352–53, 1359–60 (2018), 37 C.F.R. § 42.108(a) (“When instituting . . . review, the Board will authorize the review to proceed on all of the challenged claims and on all grounds of unpatentability asserted for each claim.”). The parties may address further during trial the disputed issues pertaining to independent claims 7 and 16, including the obviousness of, for example, a chip of at least 5,000 PEs having the recited reduced bitwidth multiplier circuit. Claims 8–14 and 17–23 each depend, directly or indirectly from either claim 7 or claim 16, and thus also include the requirement of no less than 5,000 processing elements. During trial, the parties may address the limitation for those claims as well.

IV. CONCLUSION

Petitioner has demonstrated a reasonable likelihood of prevailing in showing the unpatentability of at least one of the challenged claims of the '775 patent. At this stage of the proceeding, we have not made a final determination with respect to the patentability of any of the challenged claims.

V. ORDER

For the foregoing reasons, it is

ORDERED that, pursuant to 35 U.S.C. § 314(a), an *inter partes* review of claims 1–5, 7–14, and 16–23 of the ’775 patent is instituted with respect to the sole ground of unpatentability set forth in the Petition; and

FURTHER ORDERED that, pursuant to 35 U.S.C. § 314(c) and 37 C.F.R. § 42.4(b), *inter partes* review of the ’775 patent shall commence on the entry date of this Order, and notice is hereby given of the institution of a trial.

For PETITIONER:

Elisabeth Hunt
Marc Johannes
Greg Nieberg
WOLF, GREENFIELD & SACKS, P.C.
ehunt-ptab@wolfgreenfield.com
mjohannes-ptab@wolfgreenfield.com
gnieberg-ptab@wolfgreenfield.com

For PATENT OWNER:

Peter Lambrianakos
Vincent Rubino
Enrique Iturralde
Richard Cowell
FABRICANT LLP
plambrianakos@fabricantllp.com
vrubino@fabricantllp.com
eiturralde@fabricantllp.com
rcowell@fabricantllp.com