

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner,

v.

SINGULAR COMPUTING LLC,
Patent Owner.

IPR2023-00396
Patent 11,169,775 B2

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

LEE, *Administrative Patent Judge*.

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

I. INTRODUCTION

Google LLC (“Petitioner”) filed a Petition requesting an *inter partes* review of claims 1–5, 7–14, 16–23 (“challenged claims”) of U.S. Patent No. 11,169,775 B2 (Ex. 1001, “the ’775 patent”) pursuant to 35 U.S.C. §§ 311–319. Paper 2 (“Pet.”). Singular Computing LLC (“Patent Owner”) filed a Preliminary Response. Paper 7 (“Prelim. Resp.”).

We have authority to determine whether to institute review under 35 U.S.C. § 314 and 37 C.F.R. § 42.4. Upon consideration of the arguments and evidence before us, we conclude Petitioner has not shown a reasonable likelihood that it would prevail in establishing unpatentability of any challenged claim of the ’775 patent.

The Petition is *denied*, and no *inter partes* review is instituted.

II. BACKGROUND

A. *Real Parties in Interest*

Petitioner identifies only itself as real party in interest. Pet. xvi.
Patent Owner identifies only itself as real party in interest. Paper 5, 1.

B. *Related Matters*

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

Patent Owner additionally identifies *Singular Computing LLC v. Google LLC*, 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.) as related matter. Paper 5, 1.

Petitioner states it is “filing a second petition for *inter partes* review of the ’775 patent” and “also filing petitions for *inter partes* review of related U.S. Patent No. 10,754,616.” Pet. xvii.

Patent Owner further identifies the following *inter partes* review proceedings before the Board as related matters: IPR2023-00394, IPR2023-00395, IPR2023-00397,¹ and IPR2023-00410.² Paper 5, 1.

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 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. The '775 patent recognizes, 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

¹ IPR2023-00357 involves a second petition by Petitioner seeking an *inter partes* review of the '775 patent.

² IPR2023-00410, however, is not a recognized proceeding in the Board’s P-TACTS (Patent Trial and Appeal Case Tracking System).

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 elements³ 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 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.

³ 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.

The '775 patent explains:

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:




FIG. 1

Figure 1 is an example of the overall design of a SIMD processor. *Id.* at 2:31–32. Computing system 100 of Figure 1 includes a collection of many 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 a 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:




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.

Claims 1, 7, and 16 are independent and reproduced below:⁴

- [1PRE] A computer 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;

⁴ The bracketed captions correspond to those used by Petitioner to reference the claim elements. *See* Pet. 83–90. We use them here for ease of reference, understanding, and consistency.

- [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;
- [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 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.

Id. at 30:21–31:25.

[7PRE] A computer system, comprising:

[7A] a host computer;

[7B] a computing chip comprising:

[7C] 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, wherein each of a first subset of the plurality of first processing elements is positioned at a first edge of the processing element array, and wherein each of a second subset of the plurality of first processing elements is positioned in the interior of the processing element array;

[7D] an input-output unit connected to each of the first subset of the plurality of first processing elements;

[7E] a plurality of processing element connections, each processing element connection connecting one of the plurality of first processing elements with another of the plurality of first processing elements, wherein each of the plurality of first processing elements is connected to at least one other of the plurality of first processing elements by at least one of the plurality of processing element connections;

[7F] a plurality of memory units, wherein each of the plurality of first processing elements is associated with a corresponding one of the plurality of memory units, and wherein each of the plurality of memory units is local to its associated one of the plurality of first processing elements;

- [7G] a plurality of first arithmetic units, wherein each of the plurality of first processing elements has positioned therein at least one of the plurality of arithmetic units;
- [7H] a plurality of second processing elements; and a plurality of second arithmetic units, wherein each of the plurality of second processing elements has positioned therein at least one of the plurality of second arithmetic units; and
- [7I] a host connection at least partially connecting the input-output unit with the host computer;
- [7J] wherein the plurality of first arithmetic units each comprises a first corresponding multiplier circuit adapted to receive as a first input to the first 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 first 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;
- [7K] wherein the first multiplier circuits corresponding to the plurality of first arithmetic units each comprises a first respective plurality of transistors and has no other transistors;
- [7L] wherein the plurality of second arithmetic units each comprises a second corresponding multiplier circuit adapted to receive as inputs to the second corresponding multiplier circuit two floating point values each of width at least 32 bits;
- [7M] wherein the second multiplier circuits corresponding to the plurality of second arithmetic units each comprises a second respective plurality of transistors;
- [7N] wherein each of the second respective pluralities of transistors of the second multiplier circuits corresponding to the plurality of second arithmetic units exceeds in number each of the first respective pluralities of transistors of the first multiplier circuits corresponding to the plurality of first arithmetic units.

Id. at 31:54–32:53.

[16PRE] A circuit board connected to a host computer, the circuit board comprising:

[16A] a plurality of processing element arrays, each comprising;

[16B] 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, wherein each of a first subset of the plurality of first processing elements of said processing element array is positioned at a first edge of said processing element array, and wherein each of a second subset of the plurality of first processing elements of said processing element array is positioned in the interior of said processing element array;

[16C] a plurality of processing element connections, each processing element connection connecting one of the plurality of first processing elements of said processing element array with another of the plurality of first processing elements of said processing element array, wherein each of the plurality of first processing elements of said processing element array is connected to at least one other of the plurality of first processing elements of said processing element array by at least one of the plurality of processing element connections of said processing element array;

[16D] a plurality of memory units, wherein each of the plurality of first processing elements of said processing element array is associated with a corresponding one of the plurality of memory units of said processing element array, and wherein each of the plurality of memory units of said processing element array is local to its associated one of the plurality of first processing elements of said processing element array; and;

[16E] a plurality of arithmetic units, wherein each of the plurality of first processing elements of said processing element array has positioned therein at least one of the plurality of arithmetic units of said processing element array;

[16F] wherein the plurality of arithmetic units of said processing element array each comprises a corresponding multiplier circuit adapted to receive as a first input to the first 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 first 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; and,

[16G] an input-output unit connected to at least one of the plurality of processing element arrays; and,

[16H] a host connection at least partially connecting the input-output unit with the host computer.

Id. at 33:30–34:16.

D. Evidence relied on by Petitioner

Petitioner relies on the following document references:⁵

⁵ The '775 patent issued from Application No. 16/882,694, filed May 25, 2020, which is a continuation of Application No. 16/571,871, filed on September 16, 2019, now Patent No. 10,664,236, which is a continuation of Application No. 16/175,131, filed on October 30, 2018, now Patent No. 10,416,961, which is a continuation of Application No. 15/784,359, filed on October 16, 2017, now Patent No. 10,120,648, which is a continuation of Application No. 14/976,852, filed on December 21, 2015, now Patent No. 9,792,088, which is a continuation of Application No. 13/849,606, filed on March 25, 2013, now Patent No. 9,218,156, which is a continuation of Application No. 13/399,884, filed on February 17, 2012, now Patent No. 8,407,273, which is a continuation of Application No. 12/816,201, filed on June 15, 2010, now Patent No. 8,150,902. Ex. 1001, codes (21), (22), (63). The '775 patent also claims priority to Provisional Application No. 61/218,691, filed June 19, 2009. *Id.* at code (60).

Name	Reference	Exhibit
Schmidt-2000	Schmidt et al., <i>KPROC - An Instruction Systolic Architecture for Parallel Prefix Applications</i> , Parallel and Distributed Computing Practices Vol. 3, No. 2, pp. 163–178 (2000)	1007
Schmidt-1999	Schmidt et al., <i>A Parallel Accelerator Architecture for Multimedia Video Compression</i> , Proceedings of Euro-Par, Parallel Processing (Lecture Notes in Computer Science; vol. 1685), pp. 950–960 (1999)	1009
Schimmler-1996	Schimmler et al., <i>The Instruction Systolic Array in Image Processing Applications</i> , Proceedings of SPIE—The International Society for Optical Engineering, Vision Systems: Sensors, Sensor systems, and Components, vol. 2784, pp. 136–144 (1996)	1011
Kan ⁶	US Pat. No. 5,355,508	1015

Petitioner also relies on the Declaration of Paul Franzon, Ph.D.
Ex. 1003.⁷

⁶ Issued October 11, 1994, based on Application No. [07]793,353 and PCT No. PCT/JP91/00604 (PCT filed May 7, 1991). Ex. 1015, codes (21), (22), (45), (86).

⁷ Patent Owner's Preliminary Response does not rely on the testimony of any technical declarant.

E. Asserted Grounds of Unpatentability

Petitioner asserts that the challenged claims of the '775 patent are unpatentable based on the following grounds:

Claim(s) Challenged	35 U.S.C. §⁸	Reference(s)/Basis
1, 2, 5	103	Schmidt-2000, Kan
3, 4, 7–14	103	Schmidt-2000, Kan, Schmidt-1999 ⁹
16, 23	103	Schmidt-2000, Schimmler-1996
17–22	103	Schmidt-2000, Schimmler-1996, Schmidt-1999, Kan

III. ANALYSIS

A. Level of Ordinary Skill in the Art

Petitioner contends that a person of ordinary skill in the art at the time of the alleged invention “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” and that “[m]ore education could substitute for experience, and vice versa.” Pet. 4 (citing Ex. 1003 ¶¶ 35–36).

⁸ The Leahy-Smith America Invents Act, Pub. L. No. 112–29, 125 Stat. 284 (2011) (“AIA”), amended 35 U.S.C. §§ 102 and 103. We preliminarily determine that the '775 patent has an effective filing date prior to the effective date of the applicable AIA amendments (March 16, 2013). Therefore, we refer to the pre-AIA version of 35 U.S.C. § 103.

⁹ For claims 7–14, Petitioner adds the background knowledge of one of ordinary skill in the art. Pet. 48. But the background knowledge of one of ordinary skill in the art is implicitly included in an obviousness analysis without need to expressly state it in the specific identification of the ground.

Patent Owner has not proposed its own articulation of the level of ordinary skill in the art. Nor has Patent Owner expressed disagreement with Petitioner's statement of the level of ordinary skill. Patent Owner states: "For the purpose of this Preliminary Response only, Patent Owner utilizes Petitioner's proposed level of ordinary skill in the art." Prelim. Resp. 9.

On this record, we adopt Petitioner's proposed level of ordinary skill in the art. It appears consistent with what is reflected by the content of the applied prior art references. *Cf. Okajima v. Bourdeau*, 261 F.3d 1350, 1354–55 (Fed. Cir. 2001) (the applied prior art may reflect an appropriate level of skill).

B. Claim Construction

We use the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. § 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent. 37 C.F.R. § 42.100(b). The claim construction standard set forth in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc) is applicable.

Claim terms are generally given their ordinary and customary meaning as would be understood by one with ordinary skill in the art in the context of the specification, the prosecution history, other claims, and extrinsic evidence including expert and inventor testimony, dictionaries, and learned treatises, although extrinsic evidence is less significant than the intrinsic record. *Phillips*, 415 F.3d at 1312–17. Usually, the specification is dispositive, and it is the single best guide to the meaning of a disputed term. *Id.* at 1315.

The specification may reveal a special definition given to a claim term by the patentee, or the specification or prosecution history may reveal an intentional disclaimer or disavowal of claim scope by the inventor. *Id.* at 1316. If an inventor acts as his or her own lexicographer, the definition must be set forth in the specification with reasonable clarity, deliberateness, and precision. *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1249 (Fed. Cir. 1998). The disavowal, if any, can be effectuated by language in the specification or the prosecution history. *Poly-America, L.P. v. API Indus., Inc.*, 839 F.3d 1131, 1136 (Fed. Cir. 2016).

Only those claim terms that are in controversy need to be construed, and only to the extent necessary to resolve the controversy. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017); *Wellman, Inc. v. Eastman Chem. Co.*, 642 F.3d 1355, 1361 (Fed. Cir. 2011); *Vivid Techs., Inc. v. Am. Sci. & Eng'g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999).

Neither Petitioner nor Patent Owner has proposed a construction for any claim term. We determine that no express claim construction is necessary for this Decision, except for one term: “input-output unit” which is recited in each of independent claims 1, 7, and 16.

“input-output unit”

In its analysis of claims 1, 7, and 16, as applied to Schmidt-2000, Petitioner reads the claimed “input-out unit” onto a data bus. Pet. 27–28, 63, 74. Patent Owner argues, however, that “because Schmidt-2000 provides direct access to each processor’s internal memory, there is no input-output unit that is connected to both a first and second edge processing unit as claimed.” Prelim. Resp. 24. Here, we consider just whether a data bus 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 which 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 International, 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, *Board Decision*, 2016 Pat. App. LEXIS 3764, at *4, was erroneous.”).

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:




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. 26–27). 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, 9:5–16 (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 conductor. 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.”

C. Alleged Obviousness of Claims 1, 2, and 5 over Schmidt-2000 and Kan

1. The Law on Obviousness

A claim is unpatentable under 35 U.S.C. § 103 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 to a person having ordinary skill in the art to which said subject matter pertains. *See 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) where in evidence, so-called secondary considerations, including commercial success, long-felt but unsolved needs, failure of others, and unexpected results. *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17–18 (1966).

When evaluating a combination of teachings, we must also “determine whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue.” *KSR*, 550 U.S. at 418 (citing *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)). “[T]here must be some

articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *Kahn*, 441 F.3d at 988.

2. Overview of Schmidt-2000

Schmidt-2000 discloses a parallel computer with 1024 floating-point processors on a single chip that Schmidt-2000 identifies as the “KPROC-chip” (“KPROC” for KiloPROCessor). Ex. 1007, 163. A “host computer” uses the KPROC-chip as a “coprocessor.” *Id.* at 163, 168–175. The KPROC-chip’s 1024 processors are arranged in a 32x32 “[i]nstruction systolic array” (ISA) in which each “identical processor[]” is “connected to its four direct neighbors.” *Id.* at 163.

Figure 2.1 of Schmidt-2000 is reproduced below:




FIG. 2.1. Control flow in an ISA .

Figure 2.1 depicts a control flow in an ISA. *Id.* at 164.

Schmidt-2000 explains:

The ISA is a quadratic array of identical processors, each connected to its four direct neighbors by single data wires. The array is synchronized by a global clock. The processors operate bit-serially, i.e., one bit per clock cycle. The size of the internal registers is 16 bits and each instruction produces a 16-bit result. Therefore, one period of 16 subsequent clock cycles is called an instruction clock cycle. The processors are controlled by instructions [and] row and column selectors. In every instruction

clock cycle one instruction, 32 row-selector bits and 32 column-selector bits enter the array.

Id. at 163–164.

The instructions are input in the upper left corner of the processor array, and from there they move step by step in horizontal and vertical direction through the array. *Id.* As shown in Figure 2.1, selectors move through the array: the row selectors move horizontally from left to right, and the column selectors move vertically from top to bottom. *Id.* The selectors mask the execution of the instructions within the processors, i.e. an instruction is executed if and only if both selector bits, currently in that processor, are equal to one. *Id.* According to Schmidt-2000:

During the execution of one instruction the processor receives the next instruction together with a column selector bit from the upper neighbour and a row selector bit from the left neighbour. The selectors are interpreted. If both are 1 then the execution of the instruction is prepared. If one of the selectors 0 then the execution of a no-operation is initialised. Instruction and column selector are propagated bit serially to the lower neighbour and the row selector to the right neighbour during the execution of the instruction.

Id. at 165–166.

Figure 3.1 of Schmidt-2000 is reproduced below:




FIG. 3.1. *Architecture of one processor.*

Figure 3.1 depicts the architecture of a single processor within the ISA. *Id.* at 165. Each processor has sixty-four general purpose registers serving as local memory, and other “units required for execution of the actual instruction,” such as a multiplier, shifter, floating point control, conditional unit, and ALU (arithmetic logic unit). *Id.* at 164–165. Each register is 16-bits wide and each instruction produces a 16-bit result. *Id.* at 163–164.

Schmidt-2000 describes that for its KPROC-chip having 1024 processors, “64-Bit and 32-Bit floating point [number] formats should be supplied.” *Id.* at 164. Schmidt-2000 explains:

In a conventional floating point processor the internal format is the one known as IEEE standard.^[10] The IEEE 64-Bit [double precision] format for example consists of one-sign bit, 11 exponent bits and 52 bits of the significant (mantissa). . . .

The solution in the KPROC processor is very simple and efficient: instead of the IEEE standard, internally an extended standard is used. The significant [mantissa] is extended to 64 bits (which even gives additional precision [than the IEEE double precision format]) and three bits are added to the exponent.

Id. at 167. Schmidt enlarges the mantissa length of both the IEEE an format (23-bit mantissa) and IEEE double precision format (52-bit mantissa) for floating point value representation to 64-bit mantissa. *Id.*

3. Overview of Kan

Kan relates to a parallel data processing system combining a SIMD (single-instruction multiple-data) unit with a MIMD (multiple-instruction multiple-data) unit and sharing a common bus, memory, and system controller. Ex. 1015, code (54). Kan proposes a processing architecture which utilizes the advantages of both SIMD and MIMD type units and mitigates the disadvantages of each across different types of processing tasks. *Id.* at 2:50–55. Figure 1 of Kan is reproduced below:

¹⁰ The IEEE standard single precision floating point representation has 32 bits, including one sign bit, 8 exponent bits, and 23 mantissa bits. Ex. 1003 ¶ 51.




FIG. 1

Figure 1 illustrates the configuration of Kan's parallel data processing system. *Id.* at 4:10–11. Figure 1 includes SIMD type parallel data processing unit 50 and MIMD type parallel data processing unit 51. *Id.* at 4:13–14. Figure 1 further includes common memory bus 41 shared by the units and system controller 43 for allocating processing tasks to a suitable one of the parallel data processing units. *Id.* at 4:15–20.

4. Independent Claim 1

a) Modifying the 16-Bit Architecture of Schmidt-2000

Petitioner begins its analysis by proposing to modify the 16-bit architecture of Schmidt-2000's KPROC-chip, e.g., 16-bit wide registers and 16-bit operands for execution units, to an 8-bit architecture. Pet. 7–10.

Petitioner asserts:

Schmidt-2000's KPROC-chip used a 16-bit architecture in which each processor has 16-bit registers and computing units that operate on 16-bit operands one bit per clock cycle. Schmidt-2000, 163–164. POSAs [persons of ordinary skill in the art] would have had motivation and reasonable expectation of success in modifying this to an 8-bit architecture to provide a high-performance coprocessor optimized for image-processing applications. Franzone, ¶¶59–60.

Each of Schmidt-2000's processors have 16-bit computing units, including a 16-bit integer multiplier and a 16-bit integer adder. Schmidt-2000, 166–167. Operands of more than 16 bits are partitioned into multiple 16-bit operands such that computations are performed iteratively, 16 bits at a time. Schmidt-2000, 166; Franzon, ¶ 60. Schmidt-2000's floating-point unit uses the 16-bit integer multiplier to multiply mantissas and the 16-bit integer adder to add exponents such that higher-bitwidth multiplication (i.e., more than 16-bits) are likewise performed iteratively, 16 bits at a time. In this way, Schmidt-2000's processors can perform floating-point multiplication of higher-precision numbers (e.g., 32-bit, 64-bit, etc.) by partitioning them into 16-bit operands and performing a sequence of 16-bit sub-operations. Schmidt-2000, 166–167; Franzon, ¶ 60.

Schmidt-2000 explains that the 16-bit architecture was chosen as a “compromise between the requirements of different application areas,” recognizing that the optimal bitwidth choice would be different for systems intended for different applications—e.g., an 8-bit architecture would be preferable for image processing; a 64-bit architecture for numeric applications. Schmidt-2000, 164–165 (“The wordlength . . . was chosen to 16 bits as a compromise between the requirements of different application areas like image processing (favourable 8 bits) and numeric (64 bits).”; Franzon, ¶ 61.

POSAs would have understood that because longer-bitwidth operands can be multiplied via iteration, the selection of the processor's architecture's native bitwidth—including the bitwidth of the processor's registers and computing units that determines the bitwidth of the inputs received by the multiplier after longer bitwidth operands are partitioned—was merely a design choice to optimize the trade-off between computation speed and the chip area and power needed for each processor, for whichever application the chip is designed to primarily perform. Franzon, ¶ 62.

Id. at 7–8. Based on the above-quoted explanation, Petitioner asserts:

POSAs would have recognized the desirability of implementing Schmidt-2000's KPROC-chip with 8-bit processors to achieve a

KPROC-chip optimized to primarily perform image processing applications in which a substantial percentage of the operands are contemplated to have fewer than 16 bits and for which Schmidt-2000 teaches 8 bits is “favourable”—and which rely on the types of computations for which Schmidt-2000 says its ISA is well-suited (matrix multiplication, convolution, FFT, motion estimation, histogram and filtered back-projection for tomographic image reconstruction, etc.)—with the capability of performing higher-bitwidth operations via iterative computations on 8-bit operands. Schmidt-2000, 164-165, 169-176. Franzon, ¶164.

Id. at 9.

We need not resolve whether, as Petitioner contends, one with ordinary skill in the art would have modified Schmidt-2000’s 16-bit architecture to an 8-bit architecture, to perform, primarily, image processing applications. Even assuming that one of ordinary skill in the art would have made the proposed modification to Schmidt-2000’s KPROC chip, and successfully in doing so, Petitioner still has not made a sufficient accounting for limitations [1E] and [1J] as we explain below. We proceed to analyze claim 1, assuming as Petitioner contends, that one of ordinary skill in the art would have modified Schmidt-2000’s 16-bit architecture to an 8-bit architecture, i.e., registers having 8 bits instead of 16 bits and all execution units, e.g., adders, multipliers, processing 8 bits at a time rather than 16 bits.

b) Combining MIMD and SIMD

Along with proposing to modify Schmidt-2000’s 16-bit architecture to 8-bit architecture, Petitioner additionally proposes to modify Schmidt-2000’s SIMD architecture¹¹ to include a MIMD unit. Pet. 14. Petitioner explains that SIMDS architectures were known to provide high performance

¹¹ Petitioner asserts that Schmidt-2000’s ISA structure is a type of SIMD. Pet. 11–12.

for only a limited class of computations, large-scale repetition of the same simple computations on different data. *Id.* at 11–12 (citations omitted).

With respect to Kan, Petitioner asserts:

Kan recognized that SIMD and MIMD (multiple-instruction multiple-data) architectures had respective advantages and disadvantages that made one well-suited to some application areas and not others (and vice versa). Kan, 1:50–68. Moreover, SIMD and MIMD architectures had respective advantages/disadvantages for different types of processing within a single application field. Kan, 2:7–49; Franzon [Ex. 1003], ¶ 75.

Kan discloses a parallel processing system combining both SIMD-type and MIMD-type units to take advantage of each architecture, while each compensates for the other’s drawbacks, thereby providing an improved parallel-processing system optimized for different types of processing within a given application field, such as image processing. Kan, 2:45–56, FIG. 1 (annotated below); Franzon [Ex. 1003], ¶ 76.




FIG. 1

[Figure 1 illustrates the configuration of a parallel data processing system according to Kan. Ex. 1015, 3:58–60.]

“SIMD type” parallel data processing unit 50 and “MIMD type” parallel processing unit 51 are both coupled to a system controller 43 that allocates processing of programs to each parallel-processing unit depending on which unit is more suitable for that type of processing. Kan, 4:10–32; Franzon [Ex. 1003],

¶ 77. System controller 43 can utilize SIMD-type unit 50 and MIMD-type unit 51 to execute separate programs operating on different data, execute portions of a program that operate on the same data (including data computed by the other unit) and/or execute the same program independently and compare the results. Kan, 4:24–5:4, 7:1–8:18; Franzon [Ex. 1003], ¶ 77.

Id. at 12–13.

Based on the foregoing, Petitioner asserts:

Given the known limitations of accelerating processing using a single type of parallel-processing architecture, POSAs would have been motivated to integrate aspects of Kan’s hybrid parallel data processing system on the same chip as the 8-bit KPROC ISA to provide a more flexible single-chip coprocessor capable of accelerating a wider array of computations and more types of processing within a given application area, such as image processing for which the 8-bit KPROC ISA is optimized (*supra* §V.A.1). Kan, 1:24–3:55, 9:31–10:16; Franzon [Ex. 1003], ¶ 78.

In particular, POSAs would have been motivated to integrate a MIMD-type unit like Kan’s on the 8-bit KPROC-chip to provide a hybrid single-chip coprocessor (“Schmidt-Kan coprocessor”) that benefits from acceleration of not only large-scale parallelism via the KPROC array, but also processing tasks for which the KPROC array is not well suited but for which Kan’s MIMD-type unit is optimized, such as course-grain parallelism, task parallelism (e.g., repetition of routines with substantial sequential or irregular computations), or complex data processing (e.g., complex data flow, frequent branch conditions or data dependency), etc., that are present in image processing applications. Kan, 1:50–2:55, 4:22–32; Franzon [Ex. 1003], ¶¶ 79–80 (also citing Schmidt-2000, 163, discussing goal of increasing chip’s processing power).

Id. at 14.

We need not resolve whether, as Petitioner contends, one with ordinary skill in the art would have been motivated to integrate a MIMD-type unit like Kan on the 8-bit KPROC-chip of Schmidt-2000 to provide a

hybrid single-chip coprocessor. Even assuming that one of ordinary skill in the art would have made the proposed modification to Schmidt-2000's KPROC chip, to include a MIMD-type unit like that of Kan, and would have been successful in doing so, Petitioner still has not made a sufficient accounting for limitations [1E] and [1J] as we explain below. We proceed to analyze claim 1, assuming as Petitioner contends, that one of ordinary skill in the art would have integrated a MIMD-type unit like that of Kan on the modified 8-bit KPROC-chip of Schmidt-2000 to provide a hybrid single-chip coprocessor.

c) Recitation [1PRE]

Claim 1 recites "A computing system." Ex. 1001, 30:21. Petitioner asserts Schmidt-2000's KPROC-chip, modified to have an 8-bit architecture, together with Schmidt-2000's host computer, and further combined with Kan's MIMD-Type Unit 51 including a plurality of separate processors, is a computing system. Pet. 20 (citing Ex. 1003 ¶¶ 93–94). The assertion is supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that the Schmidt-2000 and Kan combination proposed by Petitioner discloses a computing system. It is not necessary to determine whether the preamble of claim 1 is limiting.

d) Limitation [1A]

Claim 1 recites in limitation [1A]: "a host computer." Ex. 1001, 30:22. Petitioner asserts Schmidt-2000's host computer 102 is a host computer. Pet. 20 (citing Ex. 1003 ¶ 95). The assertion is supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that Schmidt-2000 discloses a host computer.

e) Limitation [1B]

Claim 1 recites in limitation [1B]: “a computing chip.” Ex. 1001, 30:23. Petitioner asserts Schmidt-2000’s KPROC-chip, modified to have an 8-bit architecture and including 1024 processors, is a computing chip. Pet. 20 (citing Ex. 1007, 163). Further, Petitioner asserts that the Schmidt-Kan coprocessor structure according to its combination of Schmidt-2000 and Kan is also a computing chip. *Id.* (citing Ex. 1003 ¶ 96). The assertion is supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that Petitioner’s proposed combination of Schmidt-2000 and Kan discloses a computing chip.

f) Limitation [1C]

Claim 1 recites in limitation [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;

Ex. 1001, 30:24–32.

On page 23 of the Petition, Petitioner presents an annotated version of Figure 2.1 of Schmidt-2000, reproduced below:




FIG. 2.1. Control flow in an ISA .

Figure 2.1 illustrates the control flow in Schmidt-2000's ISA (Instruction Systolic Array). Ex. 1007, 163–164.

Petitioner asserts that any two of the red-outlined processors in the top row or left-most column (or bottom row or right-most column not shown in Fig. 2.1) are first and second edge processing elements positioned at a first edge of the processing element array, and any two processor in the blue-outlined interior that are adjacent to the two edge processors are first and second interior processing elements. Pet. 22 (citing Ex. 1003 ¶ 100). As example, Petitioner has shaded illustrative first and second edge processing elements red and first and second interior processing elements blue. *Id.*

Petitioner's assertion is supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that Schmidt-2000 discloses limitation [1C].

g) Limitation [1D]

Claim 1 recites in limitation [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.” Ex. 1001, 30:33–38.

Petitioner asserts that in Schmidt-2000's system, each processor's communication register is connected to that of each of its four neighbors via a pair of wires for exchanging data to and from the adjacent processor. Pet. 23 (citing Ex. 1007, 163–164, Ex. 1003 ¶¶ 102–103). On page 24 of the Petition, Petitioner provides an annotated version of Figure 3.1 of Schmidt-2000, reproduced below:





FIG. 3.1. Architecture of one processor.

Figure 3.1 depicts the architecture of a single processor within Schmidt-2000's ISA. Ex. 1007, 165. Petitioner explains that the four pairs of wires connecting a processor's communication registers to the four adjacent

neighbor processors, shown as north, south, east, and west data links (highlighted purple), provide processing element connections. Pet. 24 (citing Ex. 1003 ¶ 104). Petitioner further explains (1) that the first edge processing element has a “data link south” that meets the “first processing element connection connecting the first edge processing element with the first interior processing element, and (2) that the second edge processing element has a “data link south” that meets the “second processing element connection connecting the second edge processing element with the second interior processing element.” *Id.* at 24–25 (citing Ex. 1003 ¶¶ 105–106). On page 25 of the Petition, Petitioner provides an annotated combination of Figures 2.1 and 3.1 of Schmidt-2000, reproduced below:



[the combined annotated figures illustrate the first processing element connection from a first edge processing element (Pet. 24–25)]

Petitioner explains that the second edge processing element would have the same “data link south” meeting the claimed “second processing element connection connecting the second processing element with the second interior processing element.” *Id.* at 25 (citing Ex. 1003 ¶¶ 105–106).

Petitioner’s assertions are supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that Schmidt-2000 discloses limitation [1D].

h) Limitation [1E]

Claim 1 recites in limitation [1E]: “an input-output unit connected to the first edge processing element and the second edge processing element.” Ex. 1001, 30:39–41. On page 26 of the Petition, Petitioner presents an annotated partial version of Figure 3.1 of Schmidt-2000, reproduced below:




Figure 3.1 depicts the architecture of a single processor within Schmidt-2000’s ISA. Ex. 1007, 165. Petitioner explains that in Schmidt-2000, every processing element has two memory banks, from which one bank (R48–R63, highlighted green) is connected to the outside and the other (R0–R47, highlighted blue) serves as local storage. Pet. 26 (citing Ex. 1003 ¶¶ 107–108).

Petitioner further explains that one of ordinary skill in the art would have understood that registers R48–R63 are “connected to the outside” for “data exchange with the external memory” of the host computer by input-output connections that allow off-chip communication between the host computer and the KPROC-chip. *Id.* at 27 (citing Ex. 1003 ¶ 109). Petitioner

asserts that one of ordinary skill in the art would have understood that R48–R63’s “‘external’ read/write” is what connects to the outside by connecting to the host computer’s external memory via the 64-bit global bus, as illustrated by red arrows in the above-annotated Figure. *Id.* Petitioner also asserts that one of ordinary skill in the art would have understood that the obvious way to implement the functionality of data exchange with external memory and connect R48–R63 to the outside world is to couple R48–R63 to the global data bus. *Id.* (citing Ex. 1003 ¶¶ 110–113).

Based on the foregoing, Petitioner specifically identifies (1) registers R48–R63 of each processor, and (2) the global data bus of the KPROC-chip, both individually or together, as the claimed “input-output unit” of limitation [1E]. *Id.* at 27–28. For reasons discussed below, we disagree and determine Petitioner has not sufficiently shown, for purposes of institution of review, that Schmidt-2000 discloses the claimed “input-output unit.”

First, we explained above in Claim Construction Section III.B that it is unreasonable and not consistent with the specification of the ’775 patent to regard a data bus as an “input-output unit.” The analysis need not be reiterated here. Second, because limitation [1E] recites “an input-output unit connected to the first edge processing element and the second edge processing element,” it is unreasonable to treat registers within the first or second edge processing element as the input-output unit. Patent Owner correctly notes (Prelim. Resp. 26) that that would regard an internal component of the processor as making a connection to the processor. That reading is illogical and Petitioner points to nothing in the specification of the ’775 patent to indicate that the ’775 patent itself used the “connected to” term in this manner. Third, because neither the global data bus nor registers R48–R63 within each processor can be the claimed “input-output unit,”

Petitioner provides no adequate reasoning how the two in combination can be the claimed “input-output unit.”

We determine that Petitioner has not sufficiently shown that Schmidt-2000 discloses limitation [1E].

i) *Limitation [1F]*

Claim 1 recites limitation [1F] which requires each of the first and second edge processing elements and each of the first and second interior processing element to include a local memory. Ex. 1001, 30:42–49. On page 29 of the Petition, Petitioner provides an annotated Figure 3,1 of Schmidt-2000, reproduced below:




FIG. 3.1. Architecture of one processor.

Figure 3.1 depicts the architecture of a single processor within Schmidt-2000’s ISA. Ex. 1007, 165. Petitioner explains that Schmidt-2000 describes that each processing element has read and write access to “its own memory,” and that each processing element’s “local memory was limited to 64 general-

purpose registers [highlighted blue] plus one communication register.”
Pet. 29 (citing Ex. 1007, 164).

Petitioner identifies all sixty-four registers, R0–R63 in the above-reproduced illustration, as the local memory in each of the first and second edge processing element and each of the first and second interior processing element. *Id.* at 30. Alternatively, Petitioner identifies only forty-eight of these registers, R0–R47 in the above reproduced illustration, which are accessible only internally, as the local memory in each of the first and second edge processing element and each of the first and second interior processing element. *Id.*

Petitioner’s assertions are supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that Schmidt-2000 discloses limitation [1F].

j) *Limitation [1G]*

Claim 1 recites limitation [1G]: “a fifth arithmetic unit.” Ex. 1001, 30:50. Petitioner asserts that any one of the eight PEs in the MIMD-type unit of the Schmidt-2000 and Kan combination is a “fifth arithmetic unit.” Pet. 36 (citing Ex. 1003 ¶ 134). The assertion is supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that the Schmidt-2000/Kan combination discloses limitation [1G].

k) *Limitation [1H]*

Claim 1 recites limitation [1H] which requires each of the first and second edge processing elements and each of the first and second interior processing element to include an arithmetic unit. Ex. 1001, 30:51–58. On page 30 of the Petition, Petitioner provides an annotated Figure 3.1 of Schmidt-2000, reproduced below:




FIG. 3.1. Architecture of one processor.

Figure 3.1 depicts the architecture of a single processor within Schmidt-2000's ISA. Ex. 1007, 165.

Petitioner identifies the computation units within each processing element (highlighted yellow) as the “arithmetic unit” of limitation [1H]. Pet. 31. Petitioner explains that these computation units include a full adder and logic gates implementing a standard ALU (arithmetic logic unit), plus an integer multiplier and a “floating point unit” that utilizes the integer multiplier with floating point control and a “shift/adder” to perform floating point operations. *Id.* (citing Ex. 1007, 166–167; Ex. 1003 ¶¶ 121–122).

Petitioner's assertions are supported by the cited evidence and not disputed by Patent Owner. We are sufficiently persuaded that Schmidt-2000 discloses limitation [1H].

l) Limitation [1I]

Claim 1 recites in limitation [1I]: “a host connection at least partially connecting the input-output unit with the host computer.” Ex. 1001, 30:59–60. Petitioner attempts to account for this limitation on page 32 of the Petition. Petitioner’s explanation, however, depends on Petitioner’s accounting for limitation [1E] which identifies Schmidt-2000’s registers R48–R63 as a part of the input-output unit. Pet. 32. On that basis, Petitioner identifies the I/O pins on Schmidt-2000’s KPROC-chip as the claimed “host connection.” *Id.*

Because we determine, as explained above in the discussion of limitation [1E], that Petitioner has not sufficiently shown that Schmidt-2000 discloses an “input-out unit,” we also determine Petitioner has not sufficiently shown that Schmidt-2000 discloses limitation [1I].

m) Limitation [1J]

Claim 1 recites in limitation [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.

Ex. 1001, 30:61–31:3.

Petitioner asserts:

Schmidt-2000’s processors (when not modified to an 8-bit architecture) have a 16-bit architecture with 16-bit registers providing 16-bit operands to computing units that operate on 16-bit operands on each instruction cycle. Schmidt-2000, 163-164, 166; Franzon [Ex. 1003], ¶ 130. To perform computations on

operands of more than 16-bits, the operands are partitioned into multiple 16-bit words for storage in the 16-bit registers, and the computing units receive and each operate on 16-bit operands at a time.

As discussed above, the processor multiplies two floating-point numbers by adding their exponents in the 16-bit adder and multiplying their mantissas in the 16-bit multiplier. If the mantissas have more than 16 bits, the processor multiplies them by first loading 16 bits of each mantissa into the multiplier, multiplying those 16-bit mantissas, and then subsequently loading the remaining mantissa bits to perform “the same procedure...more than one times,...add[ing] the intermediate results” to arrive at the full product. Schmidt-2000, 166 (description of “Multiplier”); Franzon [Ex. 1003], ¶131.

* * *

As discussed *supra* §§V.A.1-V.A.2, the Schmidt-Kan coprocessor modifies Schmidt-2000 to use an 8-bit architecture. Thus, the floating-point multiplier circuit of each PE’s arithmetic unit is adapted to receive as first and second inputs two floating-point values, each having an 8-bit exponent (added by the 8-bit adder) and a mantissa of up to 8 bits (the first mantissa partition multiplied by the 8-bit multiplier), meeting [1J]. Franzon [Ex. 1003], ¶133.

Pet. 34–36.

In the above-quoted explanation of how limitation [1J] is met, Petitioner equates Schmidt-2000’s 16-bit architecture to receiving the 16-bit mantissa of a floating point value having a 16-bit mantissa and the 8-bit modified architecture it proposes to receiving the 8-bit mantissa of a floating point value having an 8-bit mantissa. That is not only incorrect but also an oversimplification of what Schmidt-2000 describes. Petitioner’s expert, Dr. Franzon, acknowledges that in the 16-bit architecture, the internal multiplier component in each processor is adapted to receive as first and second inputs only the first 16-bit portion of a larger mantissa in floating

point values, and will subsequently receive additional mantissa portions of the same floating point values as inputs to perform the multiplication on floating point values. Ex. 1003 ¶ 131.

As recognized by the Petitioner, bits that constitute only a portion of a floating point value mantissa are loaded into the 8-bit multiplier (modified by Petitioner from 16-bit multiplier) in Schmidt-2000. Pet. 35. The size of the operand that is loaded into the 8-bit multiplier is not the size of the full mantissa of the floating point value on which the multiplication is performed, but only the size of a portion of that mantissa. The actual mantissa for floating point values in Schmidt-2000 has more than 11 bits. Ex. 1007, 164 (“64-Bit and 32-Bit floating point formats of IEEE should be supplied”). If the 8-bit operands are really that of a floating point value having just an 8-bit mantissa, then the multiplier need not be activated and applied multiple times to perform multiplication on floating point values as is recognized by Petitioner. See Pet. 35–36.

The IEEE single precision floating point format has 32 bits, including one sign bit, 8 exponent bits, and 23 mantissa bits. Ex. 1003 ¶ 51. The IEEE double precision floating point format has 64 bits, including one sign bit, 11 exponent bits, and 52 mantissa bits. Ex. 1007, 167. Schmidt-2000 expressly states that “64-Bit and 32-Bit floating point formats of IEEE should be supplied.” *Id.* at 164. That is unchanged even when Schmidt-2000’s 16-bit architecture is modified by Petitioner to an 8-bit architecture. With an 8-bit architecture, the multiplier would have to be activated and applied more times to complete a floating point operation than with a 16-bit architecture, but the full mantissa width of the floating point value being operated on is still the same, larger than that required by claim 1.

Schmidt-2000 proposes to make the mantissa width of floating point formats even larger than those according to the IEEE single precision format and IEEE double precision format. *Id.* at 167. Specifically, Schmidt-2000 proposes to extend the mantissa size of floating point numbers to 64 bits and add 3 bits to the exponent. *Id.* (“The solution in the KPROC processor is very simple and efficient: instead of the IEEE standard, internally an extended standard is used. The significant is extended to 64 bits (which even gives additional precision) and three bits are added to the exponent.”). Thus, all of the floating point formats disclosed in Schmidt-2000 have a mantissa size of more than 11 bits, i.e., 23 bits (IEEE single precision), 52 bits (IEEE double precision), and 64 bits (Schmidt-2000’s proposed extension).

Schmidt-2000’s multiplier, modified to have an 8-bit architecture, is not configured to operate with operands which are the mantissa of floating point values with an 8-bit mantissa. Petitioner’s analysis inappropriately focuses on only a portion of the operation of the multiplier. The characterization by Petitioner of the 8-bits received as an operand of the multiplier as 8-bits of the full mantissa of a floating point value is without support in Schmidt-2000. Rather, we find on this record that that is inappropriately based on hindsight in light of the Specification of the ’775 patent. Nowhere does Schmidt-2000 describe working with or operating on floating point values having a mantissa size of just 8 bits or no more than 11 bits.

In addition, we observe an internal discrepancy on Petitioner’s accounting for limitation [1J]. As discussed above, Petitioner’s explanation focuses on the fact that the internal multiplier component of Schmidt-2000, the one appearing just beneath the shifter/adder in Figure 3.1, operates only

on 8-bit operands in the modified 8-bit architecture. However, Petitioner actually has identified something larger and more inclusive as the “multiplier circuit” recited in claim 1.

Specifically, Petitioner states: “The integer adder, integer multiplier [the internal multiplier component], ‘floating point control’ and ‘shift/add’ units that are used to perform floating-point multiplication thus meet the claimed ‘*multiplier circuit*’ and are adapted to receive first and second ‘*floating point value[s]*’ as first and second inputs, as claimed.” Pet. 34 (second alteration in original). From that perspective, a floating point value that is a first operand for multiplication by that larger multiplier circuit is not simply an individual segment provided to the integer multiplier component. Neither is a floating point value that is a second operand for multiplication by that larger multiplier circuit. Rather, each of the first and second floating point values being multiplied by the larger multiplier circuit would be what those values are prior to any partitioning or segmentation that is performed as a part of the multiplication operation.

We are not sufficiently persuaded by Petitioner that Schmidt-2000 discloses limitation [1J].

n) Conclusion

For the foregoing reasons, Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of claim 1 over Schmidt-2000 and Kan.

5. Dependent Claims 2 and 5

Claims 4 and 5 each depends from claim 1. Ex. 1001, 31:26–31, 31:46–49. Thus, each incorporates all limitations of claim 1. The deficiencies of Petitioner’s arguments discussed above, with respect to claim 1, equally apply to claims 2 and 5. Thus, Petitioner has not shown a

reasonable likelihood that it would prevail in establishing obviousness of claims 2 and 5 over Schmidt-2000 and Kan.

D. Alleged Obviousness of Claims 3, 4, and 7–14 over Schmidt-2000, Kan, and Schmidt-1999

1. Overview of Schmidt-1999

Schmidt-1999 relates to a parallel architecture for a variety of algorithms for video compression. Ex. 1009, 950. Similar to Schmidt-2000, Schmidt-1999 describes an ISA comprising a quadratic array of identical processors, each connected to its four direct neighbors. *Id.* at 951.

Figure 4 of Schmidt-1999 is reproduced below:




Fig. 4: Video accelerator architecture and data paths to the environment.

Figure 4 depicts a cascaded memory concept implemented on-chip that forms a fast input and output environment for a parallel processing unit. *Id.* at 955–956. Each processor has two memory banks, each memory bank containing 8 interface registers. *Id.* One of these banks is always assigned to the corresponding processor, the other to the internal RAM by means of a fast data channel. *Id.* Data transfer is controlled by an on-chip controller that supplies the processor array with instructions and selectors. *Id.*

2. *Dependent Claims 3 and 4*

Claim 3 depends from claim 1 and claim 4 depends from claim 3.

Ex. 1001, 31:32–45. By reason of their dependency from claim 1, each of claims 3 and 4 includes all of the limitations of claim 1. The deficiencies of Petitioner’s arguments discussed above, with respect to claim 1 and in the context of alleged obviousness over Schmidt-2000 and Kan, equally apply to claims 3 and 4. Schmidt-1999, as applied by Petitioner, does not cure those deficiencies. Thus, Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of claims 3 and 4 over Schmidt-2000, Kan, and Schmidt-1999.

3. *Independent Claim 7*

In key parts, claim 7 is substantially similar to claim 1. Specifically, limitation [7J] corresponds to limitation [1J], and recites as follows:

wherein the plurality of arithmetic units each comprises a first corresponding multiplier circuit adapted to receive as a first input to the first 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 first 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.

Ex. 1001, 32:26–35. Limitation [7D] corresponds to limitation [1E] and recites as follows: “an input-output unit connected to each of the first subset of the plurality of first processing elements.” *Id.* at 31:66–67.

For limitation [7J], Petitioner relies on the same arguments it asserts for limitation [1J] in the context of alleged obviousness over Schmidt-2000 and Kan. Pet. 67. For limitation [7D], Petitioner relies on the same argument it asserts for limitation [1E] in the context of alleged obviousness over Schmidt-2000 and Kan. *Id.* at 63. Schmidt-1999, as applied by

Petitioner, does not cure those deficiencies. Thus, the same deficiencies discussed above, for limitation [1J] and limitation [1E], respectively apply to limitation [7J] and limitation [7D].

Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of claim 7 over Schmidt-2000, Kan, and Schmidt-1999.

4. *Dependent Claims 8–14*

Claims 8–14 each depend, directly or indirectly, from claim 7. Ex. 1001, 32:54–33:25. By reason of their dependency from claim 7, each of claims 8–14 includes all of the limitations of claim 7. The deficiencies of Petitioner’s arguments discussed above, with respect to claim 7, equally apply to claims 8–14. Thus, Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of any of claims 8–14 over Schmidt-2000, Kan, and Schmidt-1999.

E. *Alleged Obviousness of Claims 16 and 23 over Schmidt-2000 and Schimmler-1996*

1. *Overview of Schimmler-1996*

Schimmler-1996 relates to a parallel computer implementing an ISA. Ex. 1011, 136. Schimmler-1996 discloses integration of 1024 processors on an add-on-board for PCs achieved by the development of a low-power/low-area processor architecture whose instruction set is suited particularly for image processing applications. *Id.*

2. *Independent Claim 16*

In key parts, claim 16 is substantially similar to claim 1. Specifically, limitation [16F] corresponds to limitation [1J], and recites as follows:

wherein the plurality of arithmetic units of said processing element array each comprises a first corresponding multiplier circuit adapted to receive as a first input to the first 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 first 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.

Ex. 1001, 34:1–12. Limitation [16G] corresponds to limitation [1E] and recites as follows: “an input-output unit connected to at least one of the plurality of processing element arrays.” *Id.* at 34:13–14.

For limitation [16F], Petitioner relies on the same arguments it asserts for limitation [1J] in the context of alleged obviousness over Schmidt-2000 and Kan. Pet. 67, 74. For limitation [16G], Petitioner relies on the same argument it asserts for limitation [1E] in the context of alleged obviousness over Schmidt-2000 and Kan. *Id.* at 63, 74. The same deficiencies discussed above, for limitation [1J] and limitation [1E], respectively apply to limitation [16F] and limitation [16G]. Petitioner’s application of Schimmler-1996 does not cure those deficiencies.

Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of claim 16 over Schmidt-2000 and Schimmler-1996.

3. *Dependent Claim 23*

Claim 23 depends from claim 16. Ex. 1001, 35:23–26. Thus, claim 23 incorporates all of the limitations of claim 16. The deficiencies of Petitioner’s arguments discussed above, with respect to claim 16, equally apply to claim 23. Schimmler-1996, as applied by Petitioner, do not cure those deficiencies. Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of claim 23 over Schmidt-2000 and Schimmler-1996.

F. Alleged Obviousness of Claims 17–22 over Schmidt-2000, Schimmler-1996, Schmidt-1999, and Kan

Claims 17–22 each depend, directly or indirectly, from claim 16.

Ex. 1001, 34:17–35:22. Thus, each of claims 17–22 incorporates all of the limitations of claim 16. The deficiencies of Petitioner’s arguments discussed above, with respect to claim 16 in the context of alleged obviousness over Schmidt-2000 and Schimmler-1996, equally apply to claims 17–22.

Schmidt-1999 and Kan, as applied by Petitioner, do not cure those deficiencies. Petitioner has not shown a reasonable likelihood that it would prevail in establishing obviousness of any of claims 17–22 over Schmidt-2000, Schimmler-1996, Schmidt-1999, and Kan.

IV. CONCLUSION

For the foregoing reasons, we conclude that Petitioner has not shown a reasonable likelihood that it would prevail in establishing unpatentability of any of claims 1–5, 7–14, and 16–23 of the ’775 Patent.

V. ORDER

It is

ORDERED that the Petition is denied, and no *inter partes* review is instituted on any claim over any alleged ground of unpatentability.

IPR2023-00396
Patent 11,169,775 B2

For Petitioner:

Elizabeth H. Hunt
Marc S. Johannes
Gregory S. Nieberg
WOLF, GREENFIELD & SACKS, P.C.
EHunt-PTAB@wolfgreenfield.com
MJohannes-PTAB@wolfgreenfield.com
GNieberg-PTAB@wolfgreenfield.com

For Patent Owner:

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