UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/749,379	06/24/2015	Alan Baker	A04715 (ALTR:0034)	1059

83515 7590 10/15/2020

FLETCHER YODER PC
P.O.BOX 692289
HOUSTON, TX 77269-2289

EXAMINER
LIN, ARIC

ART UNIT	PAPER NUMBER
2851	

NOTIFICATION DATE	DELIVERY MODE
10/15/2020	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docket@fyiplaw.com
hill@fyiplaw.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte ALAN BAKER, ANDREW CHAANG LING, and ANDREI MIHAI HAGIESCU MIRISTE

Appeal 2019-005572
Application 14/749,379
Technology Center 2800

Before ROMULO H. DELMENDO, LINDA M. GAUDETTE, and DEBRA L. DENNETT, *Administrative Patent Judges*.

DENNETT, *Administrative Patent Judge.*

DECISION ON APPEAL[1]

Pursuant to 35 U.S.C. § 134(a), Appellant[2] appeals from the Examiner's decision to reject claims 1–3, 6, 7, 9–14, 18, and 19 of Application 14/749,379. We have jurisdiction under 35 U.S.C. § 6(b).

For the reasons set forth below, we AFFIRM.

---

[1] In our Decision, we refer to the Specification ("Spec.") of Application No. 14/749,379 filed June 24, 2015; the Final Office Action dated May 17, 2018 ("Final Act."); the Advisory Action dated Aug. 9, 2018 ("Adv. Act."); the Appeal Brief filed Jan. 14, 2019 ("Appeal Br."); the Examiner's Answer dated May 16, 2019 ("Ans."); and the Reply Brief filed July 16, 2019 ("Reply Br.").

[2] We use the word "Appellant" to refer to "applicant" as defined in 37 C.F.R. § 1.42. Appellant identifies the real party in interest as Altera Corporation. Appeal Br. 2.

STATEMENT OF THE CASE

The '379 Application relates to dynamic sizing of channels used for kernel communication on integrated circuits (ICs) such as field programmable gate arrays (FPGAs). Spec. ¶ 1. The Specification explains that in ICs such as FPGAs, the programmable logic is typically configured using low level programming languages, but these may provide a low level of abstraction and, thus, a development barrier for programmable logic designers. *Id.* ¶ 4. The Specification contends that higher level programming languages enable more ease in programmable logic design, and are used to generate code corresponding to the low level programming languages. *Id.* The Specification uses "kernels" to refer to a digital circuit that implements a specific function and/or program. *Id.* Kernels may be useful to bridge the low level programming languages into executable instructions that may be performed by the integrated circuits. *Id.* Each kernel implemented on the IC may execute independently and concurrently from the other kernels on the IC. *Id.* Kernels may be individually balanced and data may flow from one kernel to another using one or more dataflow channels (e.g., First-in-first-out (FIFO) channels) between two kernels. *Id.*

According to the Specification, the dataflow channels may be varied in size to accept an appropriate amount of data to flow from one kernel to another. *Id.* ¶ 5. User-specified capacity for channels does not account for implementation details, because users typically only work with the higher level programs rather than the low level programming languages. *Id.*

The Specification indicates that the invention concerns systems, methods, and devices for enhancing performance of machine-implemented programs through automatic inter-kernel channel sizing based on one or

more factors. *Id.* ¶ 7. The automatic sizing may aim to increase data throughput between kernel executions. *Id.*

Claim 1 reproduced below from the Claims Appendix of the Appeal Brief illustrates the claimed subject matter:

> 1. A tangible, non-transitory, machine-readable-medium, comprising machine readable instructions to:
>
> access, via a compiler, a high level program comprising instructions to be programmed on an integrated circuit;
>
> convert, via the compiler, the high level program into a low level program to be implemented on the integrated circuit, wherein the low level program comprises a first kernel, a second kernel, and an inter-kernel channel that enables inter-channel communication between the first kernel and the second kernel; wherein the first kernel and the second kernel each comprise a digital circuit implementation that bridge the low level program and executable instructions to be performed by the integrated circuit;
>
> identify, via the compiler: a latency of the inter-kernel channel, predication between the first kernel and the second kernel, a scheduling imbalance between data processing of the first kernel and data processing of the second kernel, or any combination thereof;
>
> modify, via the compiler, a size of the inter-kernel channel, by:
>
> adding additional depth to the inter-kernel channel until a depth of the inter-kernel channel is greater than the latency;
>
> adding additional depth to the inter-kernel channel until the depth of the inter-kernel channel includes enough space to store the implemented channel capacity and additional data received during an amount of time equal to the latency;
>
> adjusting the depth of the inter-kernel channel based upon a calculation of a number of threads that need to be held in the inter-kernel channel, in the worst case,

when one of the first kernel or the second is able to consume more threads than the other;

or any combination thereof; and

provide, via the compiler, the low level program with the modified size of the inter-kernel channel to the integrated circuit for implementation on the integrated circuit.

REFERENCES

The Examiner relies on the following references in rejecting the claims:

Name	Reference	Date
Fujiwara et al. (Fujiwara")	US 2005/0080874 A1	Apr. 14, 2005
Chen et al. ("Chen")	US 2013/0212365 A1	Aug. 15, 2013
Kumar et al. ("Kumar")	US 2014/0098683 A1	Apr. 10, 2014
Sundararajan et al. ("Sundararajan")	US 8,875,073 B1	Oct. 28, 2014

REJECTIONS[3]

The Examiner rejects claims 1–3, 6, 9–14, 16, 18, and 19 under 35 U.S.C. § 103 over Chen in view of Sundararajan and Fujiwara, and claim 7 under 35 U.S.C. § 103 over Chen in view of Sundararajan and Fujiwara, and further in view of Kumar. Final Act. 8–19.

[3] In the Answer the Examiner withdrew the rejections of claims 1–10 and 18–20 under 35 U.S.C. § 101; claims 1–20 under 35 U.S.C. § 112(a); and claims 2 and 14 under 35 U.S.C. § 112(d).

DISCUSSION

We review the appealed rejections for error based upon the issues identified by Appellant and in light of the arguments and evidence produced thereon. *Ex parte Frye*, 94 USPQ2d 1072, 1075 (BPAI 2010) (precedential), (cited with approval in *In re Jung*, 637 F.3d 1356, 1365 (Fed. Cir. 2011)) ("[I]t has long been the Board's practice to require an applicant to identify the alleged error in the [E]xaminer's rejections."). After considering the evidence presented in this Appeal and each of Appellant's arguments, Appellant fails to persuade us of reversible error in the Examiner's rejections.

Appellant argues for patentability of independent claims 1, 11, and 18 based on limitations in each claim relating to modification to an inter-kernel channel that enables inter-channel communication between first and second kernels. Appeal Br. 19–22. We select claim 1 as representative of the claims subject to the first ground of rejection, which covers all pending claims except claim 7. 37 C.F.R. § 42.37(c)(1)(iv). Appellant relies on its arguments against the first rejection for patentability of claim 7. Appeal Br. 24. Thus, claims 2, 3, 6, 7, 9–14, 18, and 19 stand or fall with claim 1.

The Examiner finds that Chen discloses a compilation process in which a program having executable instructions in a high level programming language is compiled into a low level program for implementation on the integrated circuit. Ans. 4–5 (citing Chen ¶¶ 9, 29–30). The Examiner finds that Chen's compiler, like the claimed invention, generates hardware implementations in the low level language for a plurality of kernels (functional blocks) described by executable instructions in the high level language. *Id.* at 5 (citing Chen ¶¶ 43, 57–58, Figs. 3, 4). The Examiner finds that the kernels receive inputs through an inter-kernel channel

implemented as a FIFO buffer, which are sized in accordance with implementation details. *Id.* (citing Chen ¶¶ 51, 52). According to the Examiner, similarly to Appellant's invention, Chen allows designers to generate low level implementations of kernels without knowledge of low level hardware description languages. *Id.* (comparing Chen ¶ 6 to Spec. ¶ 20).

Regarding claim 1, the Examiner finds that Chen teaches the limitations except for the "identify . . . the latency," "adding additional depth," and "adjusting the depth" limitations. Final Act. 8–10. In other words, the Examiner finds that Chen discloses sizing a communication channel in a low-level design, but does not disclose either identifying from the low level program the specific characteristics that the sizing is based on, or the claimed modifications performed. *Id.* at 11. The Examiner finds that Sundararajan and Fujiwara disclose these elements. *Id.* Specifically, the Examiner finds that Sundararajan discloses the "identify . . . the latency" and "adjusting the depth" limitations. *Id.* at 9–10. The Examiner finds that Fujiwara discloses identifying latency of the inter-kernel channel and adding additional depth to the inter-kernel channel until a depth of the inter-kernel channel is greater than the latency. *Id.*

The Examiner concludes that it would have been obvious to one of ordinary skill in the art at the time of the invention to combine the teachings of Chen, Sundararajan, and Fujiwara because it would have involved a routine combination of known elements according to known methods to produce predictable results. *Id.* at 10–11.

Appellant seeks reversal of the rejection over Chen in view of Sundararajan and Fujiwara. Appeal Br. 7, 18–24.

First, Appellant argues that the cited portion of Chen does not relate to an inter-kernel channel, and thus, the Examiner wrongly relies on Chen as disclosing "modify, via the compiler, a size of the inter-kernel channel." Appeal Br. 22 (referencing Chen ¶ 52). Appellant contends that Chen relates to data pipeline registers at the inputs of each function block, which are not analogous to the claimed inter-kernel channels. *Id.* According to Appellant, Chen's data pipeline registers are not disclosed as enabling inter-channel communication between the first kernel and the second kernel. Reply Br. 2.

Claim 1 recites that a first and second kernel "each comprise[s] a digital circuit implementation that bridge[s] the low level program and executable instruction to be performed by the integrated circuit." Appeal Br. 26 (Claims App.). As the Examiner notes, Appellant defines "kernel" as "a digital circuit that implements a specific function and/or program." Ans. 7 (quoting Spec. ¶ 4). According to the Specification, kernels written in a high level programming language will have a corresponding hardware implementation in the low level program when the high level program is compiled into the low level program. Spec. ¶ 4. The Examiner explains that programs are sequences of executable instructions, so a circuit design written in OpenCL, for example, expresses the functions performed by the circuit as executable instructions. Ans. 8. Chen discloses the "bridging" limitation in teaching writing design functionality in a programming language that has a corresponding implementation in a low level program. *See id.* Chen's function blocks satisfy the "kernel" requirement of claim 1, as

> [T]hey are hardware circuit implementations of specific functions (Fig. 4, blocks 301, 202, and 401 implement

> trigonometric (cosine, tangent), arithmetic (add, multiply), and microprocessor functions, respectively) compiled from high level program representations (¶7, ¶43) that bridge low level hardware circuit implementations to executable instructions in a high level user program (¶ 57–58, ¶64).

Ans. 8.

Claim 1 recites that an inter-kernel channel enables inter-channel communication between the first kernel and the second kernel. Appeal Br. 26 (Claims App.). The inter-kernel channels may be implemented as FIFO buffers. Spec. ¶ 28.

Chen discloses FIFO buffers connected to inputs of function blocks to handle data flows from upstream function blocks on the data path to a given functional block. Chen ¶¶ 34 (block processes data according to instructions within block and products output that be read by other blocks), 43 (interconnected network of function blocks representing combinations of logic elements implement functions), 51 (data pipeline registers at inputs of function block preferably are FIFOs to balance pipelines). "Since data input to a functional block from other functional blocks is handled by the FIFO buffers on the inputs of the receiving functional block, the FIFO buffers necessarily constitute an inter-kernel channel that communicatively couples together kernels (functional blocks)." Ans. 8–9. The depth (size) of each FIFO in Chen may be selected based on pipeline imbalances. Chen ¶ 52.

Our review of the record supports the Examiner's finding that Chen discloses "kernels (functional blocks performing specific functions, written in high level virtual fabrics that are compiled into low level implementations) that are communicatively coupled by inter-kernel channels (FIFO buffers handling input to a functional block by other functional blocks

in a virtual fabric), the channels being sized based on implementation details (pipeline imbalances)." *See* Ans. 9.

Appellant's second argument is that none of the cited art includes a compiler that looks at implementation details to auto-size inter-kernel channels. *See* Appeal Br. 22. According to Appellant, Chen at best discloses manual user specification of FIFO buffers based upon the maximum expected pipeline imbalance. *Id.* However, Appellant argues, the manual user specification "is precisely the input avoided by the current application, instead relying on data specifically available to the compiler to derive the sizings." *Id.* (citing Spec. ¶ 5).

The Examiner counters that claim 1 does not require auto-sizing inter-kernel channels, but, nonetheless, Sundararajan discloses a compiler that looks at implementation details to auto-size inter-kernel channels. Ans. 11.

Claim 1 recites a tangible, non-transitory, machine-readable medium to, *inter alia*," identify, via the compiler: a latency of the inter-kernel channel, predication between the first kernel and the second kernel, a scheduling imbalance between data processing of the first kernel and data processing of the second kernel, or any combination thereof." Appeal Br. 26 (Claims App.).

Sundararajan discloses generation of internal data interfaces that is substantially automated. Sundararajan 3:17–18, 47–49. Sundararajan discloses automatic analysis for data propagation, and determining data propagation may be used for determining data input width, data output width, and depth of a buffer entry, to generate FIFO buffers of indicated sizes. *Id.* 6:4–5, 15–24; 9:46–52; 10:17–18. Sundararajan discloses that latency may be automatically generated. *Id.* 7:4–5. Therefore, contrary to

Appellant's argument, Sundararajan discloses the compiler auto-sizes channels based on implementation details.

Third, Appellant argues that neither Sundararajan nor Fujiwara disclose identifying a latency of the inter-kernel channel, predication between the first kernel and the second kernel, a scheduling imbalance between data processing of the first kernel and data processing of the second kernel, or any combination thereof. *Id.* at 23. The cited portion of Sundararajan, according to Appellant, relates to "an estimate of latency of a module from input of data at an input interface to output of corresponding data at an output interface after being processed by a kernel of such a module." *Id.* (citing Sundararajan 7:5–9). Appellant contends that latency thus is not of an inter-kernel channel, but rather from an input to an output of a module post kernel processing. *Id.* Appellant argues that the cited portion of Fujiwara relates to a "latency of the network path, which is a period of time between transmission of data from the transmission-side node and a reception of an acknowledgement notice from the reception-side mode," and this is not analogous to the latency of an inter-kernel channel. *Id.* (citing Fujiwara ¶ 7).

Claim 1 recites "identify, via the compiler a latency of the inter-kernel channel, predication between the first kernel and the second kernel, a scheduling imbalance between data processing of the first kernel and data processing of the second kernel, or any combination thereof." Appeal Br. 26 (Claims App.). The use of commas between the "latency," "predication," and "scheduling imbalance" elements, plus the word "or" in the limitation means that any one of the listed elements or any combination of two or all three of the listed elements is required—but there is no requirement for *each* of the listed elements to be present.

The Specification defines "predication" as "the channel read and/or writes are not executed every execution cycle." Spec. ¶ 34. The Examiner explains that Sundararajan discloses

> 1/m and 1/n source/sink data rates (indicating one data sample written to the channel by the data source per *m* cycles, and read from the channel by the data sink per *n* cycles). The 1/*m* data rate against the 1/*n* data rate with burst *k* and idle *p* constitutes a schedule imbalance between data processing of the first and second kernels, since data is not only processed at different rates, but there are periods where one side is idle and the other side is processing at the constant 1/m rate.

Ans. 15. Thus, the record supports that predication and schedule imbalance are disclosed in Sundararajan. Appellant, however, fails to address these findings or the Examiner's citation to column 9, line 46 to column 10, line 15 of Sundararajan. *See generally* Appeal Br., Reply Br. Appellant thus fails to identify any error in the Examiner's finding that Sundararajan discloses the claim limitation. *See In re Crish*, 393 F.3d 1253, 1259 (Fed. Cir. 2004) ("[W]hen the prior art evidence reasonably allows the PTO to conclude that a claimed feature is present in the prior art, the evidence 'compels such a conclusion if the applicant produces no evidence or argument to rebut it.'") (quoting *In re Spada*, 911 F.2d 705, 708 n.3 (Fed. Cir. 1990)).

Appellant's contention that the cited portion of Fujiwara relates to a latency of the network path that is not analogous to the latency of an inter-kernel channel is not persuasive of reversible error for two reasons. *See* Appeal Br. 23. As an initial matter, given Sundararajan's disclosure discussed above, whether Fujiwara teaches latency of an inter-kernel channel is irrelevant to the Examiner's prima facie case of obviousness. In addition, Fujiwara discloses digital circuits that implement specific functions

(meeting the Specification's definition of "kernel") and a network path that communicatively couples processing devices with storage devices with latency that varies in the course of the communication. Fujiwara ¶¶ 2, 7. One of ordinary skill in the art would reasonably understand "inter-kernel channel," which is not defined in the Specification, to mean "anything that communicatively couples kernels." *See In re Am. Acad. of Sci. Tech. Ctr.*, 367 F.3d 1359, 1364 (Fed. Cir. 2004) (During prosecution, an application's claims are given their broadest reasonable scope consistent with the specification.). Fujiwara's disclosures support the Examiner's finding that the latency of the network path is the latency of an inter-kernel channel. Appellant's mere assertion that this latency is "not analogous" to latency of an inter-kernel channel is insufficient to rebut the Examiner's finding. *See Crish*, 393 F.3d at 1259.

Finally, Appellant argues that Sundararajan does not mention inter-kernel channels, and Fujiwara does not mention kernels at all. *Id.* This argument is not persuasive of reversible error.

We agree that Sundararajan does not use the term "inter-kernel channels," but the reference teaches communicatively coupled kernels, as discussed above, thus discloses the concept. Fujiwara does not use the term "kernel," but discloses digital circuits that implement a specific function and/or program (*see* Specification ¶ 4). *See* Fujiwara Abst.

We sustain the rejection of claim 1 over Chen in view of Sundararajan and Fujiwara.

DECISION SUMMARY

In summary:

Claims Rejected	35 U.S.C. §	Reference(s)/Basis	Affirmed	Reversed
1–3, 6, 9–14, 16, 18, 19	103	Chen, Sundararajan, Fujiwara	1–3, 6, 9–14, 16, 18, 19	
7	103	Chen, Sundararajan, Fujiwara, Kumar	7	
Overall Outcome			**1–3, 6, 7, 9–14, 16, 18, 19**	

TIME PERIOD FOR RESPONSE

No time period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a).

AFFIRMED