

No. 21-

IN THE
Supreme Court of the United States

FG SRC LLC,

Petitioner,

v.

MICROSOFT CORPORATION,

Respondent.

ON PETITION FOR A WRIT OF CERTIORARI TO THE UNITED
STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT

PETITION FOR A WRIT OF CERTIORARI

JAY P. KESAN
CECIL KEY
HENNING SCHMIDT
DIMUROGINSBERG
DGKEYIP GROUP
1750 Tysons Boulevard,
Suite 1500
Tysons Corner, VA 22102
(703) 289-5118

MICHAEL W. SHORE
Counsel of Record
ARI B. RAFILSON
SHORE CHAN LLP
901 Main Street, Suite 3300
Dallas, TX 75202
(214) 593-9110
mshore@shorechan.com

Counsel for Petitioner

308635



COUNSEL PRESS

(800) 274-3321 • (800) 359-6859

QUESTION PRESENTED

This petition seeks only a Grant, Vacate, and Remand (GVR) Order.

This petition seeks a GVR Order regarding two judgments by the Court of Appeals for the Federal Circuit, Appeal 2020-1993, and Joint Appeal 2020-1925 and 2020-1926. The question presented is the same in both instances.

The single question presented is whether the Federal Circuit's summary decisions below should be vacated and these cases remanded to cure the constitutional violation in accordance with this Court's decision in *Arthrex*.

PARTIES TO THE PROCEEDINGS BELOW

Petitioner FG SRC was the patent owner in each of the proceedings before the Patent Trial and Appeal Board and the appellant in the Court of Appeals for the Federal Circuit. FG SRC is the successor-in-interest to Saint Regis Mohawk Tribe as to all rights, title, and interest in and to the '524, '324, and '800 Patents.

Respondent Microsoft Corporation was the petitioner in each of the proceedings before the Patent Trial and Appeal Board and appellee in the Court of Appeals for the Federal Circuit.

CORPORATE DISCLOSURE STATEMENT

Pursuant to this Court's Rule 29.6, Petitioner FG SRC states that it has no parent corporation and that no publicly held company owns 10% or more of its stock.

RELATED PROCEEDINGS

The following proceedings are directly related to this case within the meaning of Rule 14.1(b)(iii):

- *FG SRC, LLC v. Microsoft Corp.*, No. 20-1993, U. S. Court of Appeals for the Federal Circuit). Judgment entered May 6, 2021.
- *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01604, U. S. Patent and Trademark Office, Patent Trial and Appeal Board. Final Written Decision entered Apr. 30, 2020 (Paper 76).
- *FG SRC, LLC v. Microsoft Corp.*, No. 20-1925, 20-1926, U. S. Court of Appeals for the Federal Circuit). Judgment entered May 6, 2021.
- *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01601, U. S. Patent and Trademark Office, Patent Trial and Appeal Board. Final Written Decision entered Apr. 9, 2020 (Paper 72).
- *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01605, U. S. Patent and Trademark Office, Patent Trial and Appeal Board. Final Written Decision entered Apr. 9, 2020 (Paper 76).

TABLE OF CONTENTS

	<i>Page</i>
QUESTION PRESENTED	i
PARTIES TO THE PROCEEDINGS BELOW.....	ii
CORPORATE DISCLOSURE STATEMENT	iii
RELATED PROCEEDINGS	iv
TABLE OF CONTENTS.....	v
TABLE OF APPENDICES	vi
TABLE OF CITED AUTHORITIES	viii
PETITION FOR A WRIT OF CERTIORARI.....	1
OPINIONS BELOW.....	1
SYNOPSIS OF PROCEEDINGS BELOW.....	1
JURISDICTION.....	4
STATUTORY PROVISIONS INVOLVED.....	4
STATEMENT OF THE CASE	5
REASONS FOR GRANTING THE GVR PETITION	6
CONCLUSION	10

TABLE OF APPENDICES

	<i>Page</i>
APPENDIX A—JUDGMENT OF THE UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT, FILED MAY 6, 2021	1a
APPENDIX B—JUDGMENT OF THE UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT, FILED MAY 6, 2021	3a
APPENDIX C — OPINION OF THE UNITED STATES PATENT AND TRADEMARK OFFICE, BEFORE THE PATENT TRIAL AND APPEAL BOARD, FILED APRIL 30, 2020	5a
APPENDIX D — JUDGMENT OF THE UNITED STATES PATENT AND TRADEMARK OFFICE, BEFORE THE PATENT TRIAL AND APPEAL BOARD, DATED APRIL 9, 2020	65a
APPENDIX E — JUDGMENT OF THE UNITED STATES PATENT AND TRADEMARK OFFICE BEFORE THE PATENT TRIAL AND APPEAL BOARD, DATED APRIL 9, 2020	194a
APPENDIX F — DENIAL OF REHEARING OF THE UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT, FILED AUGUST 6, 2021	313a

Table of Appendices

	<i>Page</i>
APPENDIX G — DENIAL OF REHEARING OF THE UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT, FILED AUGUST 6, 2021	315a

TABLE OF CITED AUTHORITIES

	<i>Page</i>
CASES:	
<i>FG SRC, LLC v. Microsoft Corp.</i> , 847 F. App'x 916 (Fed. Cir. 2021)	5
<i>FG SRC, LLC v. Microsoft Corp.</i> , No. 2020-1925, 2021 WL 1808605 (Fed. Cir. May 6, 2021)	3, 6, 7
<i>Lawrence on Behalf of Lawrence v. Chater</i> , 516 U.S. 163 (1996)	8, 9
<i>Netherland v. Tuggle</i> , 515 U.S. 951 (1995)	9
<i>Robinson v. Story</i> , 469 U.S. 1081 (1984)	8, 9
<i>Stutson v. United States</i> , 516 U.S. 193 (1996)	8, 9
<i>United States v. Arthrex, Inc.</i> , 141 S. Ct. 1970 (2021)	<i>passim</i>
STATUTES:	
28 U.S.C. § 1254	4
35 U.S.C. § 315	2

Cited Authorities

	<i>Page</i>
U.S. Const. art. II, § 2, cl. 2	6, 4
RULES:	
U.S. Sup. Ct. R. 13.	4
OTHER AUTHORITIES:	
<i>Microsoft Corporation v. Saint Regis Mohawk Tribe et al</i> , No. IPR2018-01604, Paper 76 (P.T.A.B. Apr. 30, 2020) (IPR2018-01604)	5
<i>Microsoft Corporation v. Saint Regis Mohawk Tribe et al</i> , No. IPR2018-01601, Paper 72 (P.T.A.B. Apr. 9, 2020) (IPR2018-01601, IPR2018-01602, and IPR2018-01603)	5
<i>Microsoft Corporation v. Saint Regis Mohawk Tribe et al</i> , No. IPR2018-01605, Paper 76 (P.T.A.B. Apr. 9, 2020) (IPR2018-01605, IPR2018-01606, and IPR2018-01607)	5, 8

PETITION FOR A WRIT OF CERTIORARI

Petitioner FG SRC LLC respectfully petitions for a writ of certiorari to review the judgment of the United States Court of Appeals for the Federal Circuit.

OPINIONS BELOW

Appeal 20-1993: The Federal Circuit’s opinion, Pet. App. A, pp. 1a-2a, is reported at 847 F. App’x 916; the denial of the motion to remand is at Pet. App. F, pp. 313a-314a; and the Board’s Administrative Decision is at Pet. App. C, pp. 5a-64a.

Appeals 20-1925 and 20-1926: The Federal Circuit’s opinion, Pet. App. B, pp. 3a-4a, is reported at 847 F. App’x 915; the denial of the motion to remand is at Pet. App. G, pp. 315a-316a; and the Board’s Administrative Decisions are at Pet. App. D, pp. 65a-193a and Pet. App. E, pp. 194a-312a.

SYNOPSIS OF PROCEEDINGS BELOW

Appeal 20-1993: FG SRC LLC (“FG SRC”)¹ is the owner of U.S. Patent No. 7,421,524 B2 (“the ’524 Patent”), which was challenged by Microsoft Corporation at the Patent Trial and Appeal Board (“PTAB” or “Board”) in the U.S. Patent and Trademark Office (“USPTO”) in IPR2018-01604. The Final Written Decision (“FWD”) in IPR2018-01604 was entered on April 30, 2020. Pet. App.

1. FG SRC is the successor in interest to Saint Regis Mohawk Tribe as to all rights, title, and interest in and to the patents that are the subject to this Petition.

C, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01604, Paper 76, (P.T.A.B. Apr. 30, 2020). FG SRC appealed this FWD in Appeal 2020-1993.

Joint Appeal 20-1925 and 20-1926: FG SRC LLC (“FG SRC”) is the owner of U.S. Patent Nos. 7,225,324 B2 (“the ’324 Patent”) and 7,620,800 B2 (“the ’800 Patent”), which were challenged by Microsoft Corporation at the PTAB in IPR2018-01601, IPR2018-01602, and IPR2018-01603 (challenging the ’324 Patent), and IPR2018-01605, IPR2018-01606, and IPR2018-01607 (challenging the ’800 Patent). On April 12, 2019, the PTAB exercised its authority under 35 U.S.C. § 315(d) to consolidate IPR2018-01601, IPR2018-01602, and IPR2018-01603 and conduct the three proceedings as one trial captioned IPR2018-01601, and separately consolidate IPR2018-01605, IPR2018-01606, and IPR2018-01607 and conduct these three proceedings as a second trial captioned IPR2018-01605. The FWDs in both IPR2018-01601 and IPR2018-01605 were entered on April 9, 2020. Pet. App. D, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01601, Paper 72 (P.T.A.B. Apr. 9, 2020); Pet. App. E *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01605, Paper 76 (P.T.A.B. Apr. 9, 2020). FG SRC appealed these FWDs in Appeals 2020-1925 and 2020-1926, which were joined by the Court of Appeals for the Federal Circuit.

On May 6, 2021, the Federal Circuit summarily affirmed all three FWDs of the PTAB under Rule 36. Pet. App. A, *FG SRC, LLC v. Microsoft Corp.*, 847 F. App’x 916 (Fed. Cir. 2021) (Appeal 20-1993); Pet. App. B, *FG SRC, LLC v. Microsoft Corp.*, No. 2020-1925, 2021 WL 1808605 (Fed. Cir. May 6, 2021) (Joint Appeal 20-1925, 20-1926). On June 30, 2021, FG SRC filed a Motion To Dismiss

Appeal And Remand To The USPTO in light of the Court's recently issued decision in *United States v. Arthrex, Inc.*, 141 S. Ct. 1970 (2021) in both Appeals. *FG SRC, LLC v. Microsoft Corporation*, 20-1993, Dkt. 39 (citing *Arthrex*, 141 S. Ct. at 1986); *FG SRC, LLC v. Microsoft Corporation*, 20-1925, Dkt. 42 (same). On Aug. 6, 2021, the Federal Circuit summarily denied FG SRC's motions (Pet. App. F, *FG SRC, LLC v. Microsoft Corporation*, 20-1993, Dkt. 42; Pet. App. G, *FG SRC, LLC v. Microsoft Corporation*, 20-1925, Dkt. 45), thereby rejecting FG SRC's constitutional challenge that the Board's decisions were rendered by unconstitutional APJ panels under the Supreme Court's decision in *United States v. Arthrex*. *Arthrex*, 141 S. Ct. at 1986.

In both appeals, the Federal Circuit's original summary affirmance of the PTAB decisions was issued 45 days *before* this Court's opinion in *Arthrex*. Nine days after *Arthrex* issued, FG SRC asked the Federal Circuit to reconsider its summary affirmance in light of the *Arthrex* decision. The Federal Circuit, once again, summarily rejected FG SRC's challenge to the constitutionality of the PTAB's actions under the *Arthrex* decision. Pet. App. F, *FG SRC, LLC v. Microsoft Corporation*, 20-1993, Dkt. 42, Pet. App. G, *FG SRC, LLC v. Microsoft Corporation*, 20-1925, Dkt. 45. The Federal Circuit's holding improperly ignored this Court's holding in *Arthrex*.

In *Arthrex*, this Court found that administrative patent judges ("APJs") wielded unreviewable authority incompatible with their appointment as inferior officers and created a mechanism for review of APJ decisions by a presidentially appointed, Senate-confirmed principal officer. *Arthrex*, 141 S. Ct. at 1978-85. This Court held

the statutory restrictions on rehearing Board decisions unconstitutional insofar as they prohibited the Director from single-handedly reviewing a decision. *Id.* at 1986-88 (plurality); *id.* at 1997 (Breyer, J., concurring in judgment in part). The Court specifically held that “the appropriate remedy is a remand to the Acting Director for him to decide whether to rehear the [*inter partes* review].” *Id.* at 1987-88 (plurality); *id.* at 1997 (concurrence).

JURISDICTION

The judgments of the Court of Appeals were both entered on May 6, 2021. Pet. App. A, pp 1a-2a; Pet. App. B, pp. 3a-4a. The Federal Circuit denied FG SRC’s Motions To Dismiss Appeal And Remand To The USPTO in light of *Arthrex* in both appeals on Aug. 6, 2021. Pet. App. F, pp 313a-314a; Pet. App. G, pp. 315a-316a. This petition is timely as it was filed within 90 days of Aug. 6, 2021, the date of entry of the judgments or orders sought to be reviewed, per this Court’s Rule 13. U.S. Sup. Ct. R. 13. This Court has jurisdiction under 28 U.S.C. §1254(1).

STATUTORY PROVISIONS INVOLVED

The relevant statutory provision is the Appointments Clause of the U.S. Constitution, reproduced here:

He shall have Power, by and with the Advice and Consent of the Senate, to make Treaties, provided two thirds of the Senators present concur; and he shall nominate, and by and with the Advice and Consent of the Senate, shall appoint Ambassadors, other public Ministers and Consuls, Judges of the supreme Court,

and all other Officers of the United States, whose Appointments are not herein otherwise provided for, and which shall be established by Law: but the Congress may by Law vest the Appointment of such inferior Officers, as they think proper, in the President alone, in the Courts of Law, or in the Heads of Departments.

U.S. Const. art. II, § 2, cl. 2.

STATEMENT OF THE CASE

This case arises out of the final written decisions by the PTAB issued in IPR2018-01601 (joined with IPR2018-01602 and IPR2018-01603), IPR2018-01604, and IPR2018-01605 (joined with IPR2018-01606 and IPR2018-01607) which were entered prior to this Court's decision in *Arthrex*, in each of which, the PTAB improperly rejected FG SRC's constitutional argument that the APJs who decided the matter were unconstitutionally appointed in violation of the Appointments Clause. Pet. App. C, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01604, Paper 76 (P.T.A.B. Apr. 30, 2020) (IPR2018-01604); Pet. App. D, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01601, Paper 72 (P.T.A.B. Apr. 9, 2020) (IPR2018-01601, IPR2018-01602, and IPR2018-01603); Pet. App. E, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01605, Paper 76 (P.T.A.B. Apr. 9, 2020) (IPR2018-01605, IPR2018-01606, and IPR2018-01607). FG SRC sought appellate review and the Federal Circuit summarily affirmed the PTAB's administrative decisions. Pet. App. A, *FG SRC, LLC v. Microsoft Corp.*, 847 F. App'x 916 (Fed. Cir. 2021) (Appeal 20-1993 regarding IPR2018-

01604); Pet. App. B, *FG SRC, LLC v. Microsoft Corp.*, 847 F. App'x 915 (Fed. Cir. 2021) (20-1925, 20-1926) (Appeals 20-1925 and 20-1926 regarding IPR2018-01601, IPR2018-01602, and IPR2018-01603; and IPR2018-01605, IPR2018-01606, and IPR2018-01607); *see also* Pet. App. F, *FG SRC, LLC v. Microsoft Corporation*, 20-1993, Dkt. 42 (denying FG SRC's Motions To Remand For Director Review based on constitutional challenge that the Board's decisions were rendered by unconstitutional APJ panels) and Pet. App. G, *FG SRC, LLC v. Microsoft Corporation*, 20-1925, Dkt. 45 (same). Like *Arthrex*, FG SRC argued that Congress granted APJs final authority to revoke previously issued patents in violation of the Appointments Clause. U.S. Const. art. ii, §2, cl. 2. The Court of Appeals improperly rejected FG SRC's constitutional argument in its Aug. 6, 2021, orders (Pet. App. F, *FG SRC, LLC v. Microsoft Corporation*, 20-1993, Dkt. 42 and Pet. App. G, *FG SRC, LLC v. Microsoft Corporation*, 20-1925, Dkt. 45) and then issued its mandates on Aug. 13, 2021, in direct violation of this Court's *Arthrex* decision, which recognized that, given their power and independence, APJs indeed function as principal officers in violation of the Appointments Clause. *Arthrex*, 141 S. Ct. at 1986. The Federal Circuit's holding is contrary to this Court's decision in *Arthrex* and must, therefore, be vacated and this case remanded for proceedings consistent with *Arthrex*.

REASONS FOR GRANTING THE GVR PETITION

I. *Arthrex* overruled the appellate court's decisions. This Court's recent decision in *Arthrex* departs in a pivotal aspect from the PTAB's decisions and the Federal Circuit's decisions, and both bodies rejected FG SRC's constitutional argument regarding the unconstitutionality

of the appointment of APJs. The PTAB decisions on the constitutionality of the appointment of APJs (and the Federal Circuit's summary affirmance thereof) have been overruled by this Court in *Arthrex, United States v. Arthrex, Inc.*, 141 S. Ct. 1970 (2021) issued on June 21, 2021.

2. *In Arthrex*, this Court held that “the unreviewable authority wielded by APJs during *inter partes* review is incompatible with their appointment by the Secretary to an inferior office.” *Id.* at 1985. Having found this constitutional violation and in lieu of finding the IPR process unconstitutional entirely, this Court provided a narrowly tailored remedy for the specific constitutional violation, *i.e.*, a “remand to the Acting Director for him to decide whether to rehear the petition filed by [Petitioner].” *Id.* at 1987. FG SRC has been denied that remedy by both the PTAB and the Federal Circuit. Under these extraordinary circumstances, Appellant FG SRC's request to grant this Petition to vacate the Federal Circuit's orders and remand these matters for proceedings consistent with this Court's order in *Arthrex* is appropriate and just.

3. This Court should grant, vacate, and remand for proceedings consistent with *Arthrex*. A GVR order is appropriate here because both the PTAB's decisions and the Federal Circuit's decisions are directly overruled by this Court's *Arthrex* decision. Pet. App. F, *FG SRC, LLC v. Microsoft Corporation*, 20-1993, Dkt. 42; Pet. App. G, *FG SRC, LLC v. Microsoft Corporation*, 20-1925, Dkt. 45; Pet. App. C, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al.*, No. IPR2018-01604, Paper 76, (P.T.A.B. Apr. 30, 2020); Pet. App. D, *Microsoft Corporation v.*

Saint Regis Mohawk Tribe et al, No. IPR2018-01601, Paper 72 (P.T.A.B. Apr. 9, 2020) (IPR2018-01601, IPR2018-01602, and IPR2018-01603); and Pet. App. E, *Microsoft Corporation v. Saint Regis Mohawk Tribe et al*, No. IPR2018-01605, Paper 76 (P.T.A.B. Apr. 9, 2020) (IPR2018-01605, IPR2018-01606, and IPR2018-01607).

4. The standard of review for a GVR order is substantially lower than the plenary review standard in Rule 10. GVR orders have “become an integral part of this Court’s practice.” *Lawrence on Behalf of Lawrence v. Chater*, 516 U.S. 163, 166 (1996) (*per curiam*). Granting GVRs “conserve[s] the scarce resources of this Court that might otherwise be expended on plenary consideration, assist the court below by flagging a particular issue that it does not appear to have fully considered, assist this Court by procuring the benefit of the lower court’s insight before we rule on the merits, and alleviate the ‘[p]otential for unequal treatment’ that is inherent in our inability to grant plenary review of all pending cases raising similar issues.” *Id.* at 167.

5. A GVR order is appropriate when, as here, a relevant Supreme Court decision was rendered only shortly before the summary affirmance by the court of appeals that is the subject of the petition for certiorari. *Robinson v. Story*, 469 U.S. 1081 (1984) (GVR granted for further consideration in light of a Supreme Court decision rendered almost three months before the summary affirmance by the court of appeals that was the subject of the petition for certiorari); *Stutson v. United States*, 516 U.S. 193, 196 (1996) (GVR appropriate where the court of appeals wrote no opinion to show whether or how it considered a [Supreme Court] precedent that

the lower tribunal had no opportunity to consider). Here, the PTAB had no opportunity to implement a procedure consistent with the constitutional requirements this Court recognized in *Arthrex*, which issued *after* the Federal Circuit issued its summary affirmances and only one week before the Federal Circuit issued its summary denial of FG SRC's Motion to Remand. Given the summary nature of the affirmance and denial, the Federal Circuit may well have relied on a prior standard that has been altered by this Court's *Arthrex* decision, which would have affected the outcome. The fact that the impact of *Arthrex* was briefed by the parties is not sufficient to determine whether the correct standard was applied by the Federal Circuit. *Lawrence on Behalf of Lawrence v. Chater*, 516 U.S. 163, 170 (1996) (“we have never held lower court briefing to bar our review and *vacatur* where the lower court's order shows no sign of having applied the precedents that were briefed”); *see also Netherland v. Tuggle*, 515 U.S. 951 (1995) (*per curiam*) (vacating court of appeals' summary order for probable failure to apply a 12-year-old Supreme Court precedent that the parties briefed to the court of appeals). A GVR order is needed in this instance because it is the only method that “guarantees to the petitioner full and fair consideration of his rights in light of all pertinent considerations.” *Stutson v. United States*, 516 U.S. 190, 197 (1996).

6. This Court need not be certain that the case would be resolved differently on remand. GVR is appropriate if there is a “reasonable probability” of a different outcome. *Lawrence on Behalf of Lawrence v. Chater*, 516 U.S. 163, 167 (1996) (*per curiam*). Petitioner clearly meets this standard.

CONCLUSION

For the reasons stated above, this Petition for a Writ of Certiorari should be granted, the decisions below should be vacated, and these cases remanded for proceedings consistent with *Arthrex*.

Date: November 4, 2021 Respectfully submitted,

JAY P. KESAN
CECIL KEY
HENNING SCHMIDT
DIMUROGINSBERG
DGKEYIP GROUP
1750 Tysons Boulevard,
Suite 1500
Tysons Corner, VA 22102
(703) 289-5118

MICHAEL W. SHORE
Counsel of Record
ARI B. RAFILSON
SHORE CHAN LLP
901 Main Street, Suite 3300
Dallas, TX 75202
(214) 593-9110
mshore@shorechan.com

Counsel for Petitioner

APPENDIX

1a

**APPENDIX A — JUDGMENT OF THE UNITED
STATES COURT OF APPEALS FOR THE
FEDERAL CIRCUIT, FILED MAY 6, 2021**

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

FG SRC, LLC,

Appellant,

v.

MICROSOFT CORPORATION,

Appellee.

2020-1993

Appeal from the United States Patent and Trademark
Office, Patent Trial and Appeal Board in No. IPR2018-
01604.

JUDGMENT

THIS CAUSE having been heard and considered, it is

ORDERED and ADJUDGED:

PER CURIAM (PROST, *Chief Judge*, LOURIE and
O'MALLEY, *Circuit Judges*).

2a

Appendix A

AFFIRMED. See Fed. Cir. R. 36.

ENTERED BY ORDER OF THE COURT

May 6, 2021
Date

/s/ Peter R. Marksteiner
Peter R. Marksteiner
Clerk of Court

3a

**APPENDIX B — JUDGMENT OF THE UNITED
STATES COURT OF APPEALS FOR THE
FEDERAL CIRCUIT, FILED MAY 6, 2021**

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

FG SRC, LLC,

Appellant,

v.

MICROSOFT CORPORATION,

Appellee.

2020-1925, 2020-1926

Appeals from the United States Patent and Trademark
Office, Patent Trial and Appeal Board in Nos. IPR2018-
01601, IPR2018-01602, IPR2018-01603, IPR2018-01605,
IPR2018-01606, IPR2018-01607.

JUDGMENT

THIS CAUSE having been heard and considered, it is

ORDERED and ADJUDGED:

PER CURIAM (PROST, *Chief Judge*, LOURIE and
O'MALLEY, *Circuit Judges*).

4a

Appendix B

AFFIRMED. See Fed. Cir. R. 36.

ENTERED BY ORDER OF THE COURT

May 6, 2021
Date

/s/ Peter R. Marksteiner
Peter R. Marksteiner
Clerk of Court

5a

**APPENDIX C — OPINION OF THE UNITED
STATES PATENT AND TRADEMARK OFFICE,
BEFORE THE PATENT TRIAL AND APPEAL
BOARD, FILED APRIL 30, 2020**

UNITED STATES PATENT
AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL
AND APPEAL BOARD

Paper No. 76
Filed: April 30, 2020

MICROSOFT CORPORATION,

Petitioner,

v.

FG SRC LLC,

Patent Owner.

IPR2018-01604
Patent 7,421,524 B2

Before KALYAN K. DESHPANDE, JUSTIN T. ARBES,
and CHRISTA P. ZADO, *Administrative Patent Judges.*

DESHPANDE, *Administrative Patent Judge.*

Appendix C

JUDGMENT

Final Written Decision

Determining All Challenged Claims Unpatentable

35 U.S.C. § 318(a)

I. INTRODUCTION

A. Background

Microsoft Corporation (“Petitioner”) filed a Petition requesting an *inter partes* review of claims 1, 2, and 13–15 of U.S. Patent No. 7,421,524 B2 (Ex. 1001, “the ’524 patent”). Paper 2 (“Pet.”). FG SRC LLC (“Patent Owner”) filed a Preliminary Response pursuant to 35 U.S.C. § 313. Paper 15 (“Prelim. Resp.”).¹

On May 3, 2019, we issued a Decision ordering that “an *inter partes* review is hereby instituted with respect to all grounds set forth in the Petition.” Paper 21 (“Dec.”), 29. After institution, Patent Owner filed a Patent Owner’s Response (Paper 38, “PO Resp.”). In reply, Petitioner filed a Petitioner’s Reply to Patent Owner’s Response (Paper 56, “Pet. Reply”). In response, Patent Owner filed a Patent Owner’s Sur-Reply (Paper 62, “PO Sur-Reply”). Petitioner and Patent Owner also filed Motions to Exclude Evidence (Papers 63 (“Pet. Mot.”), 64 (“PO Mot.”)), Oppositions to the Motions (Papers 65 (“Pet. Opp. Mot.”), 66) and Replies to the Oppositions (Papers 68, 69). Petitioner and Patent Owner presented oral arguments on February 3, 2020,

1. Saint Regis Mohawk Tribe, originally named as Patent Owner, assigned the ’524 patent to DirectStream, LLC on May 21, 2019. Paper 27, 1. DirectStream, LLC assigned the ’524 patent to FG SRC LLC on January 22, 2020. Paper 73, 1.

Appendix C

and a transcript has been entered into the record. Paper 75 (“Tr.”).

The Board has jurisdiction under 35 U.S.C. § 6. In this Final Written Decision, after reviewing all relevant evidence and arguments, we determine that Petitioner has met its burden of showing, by a preponderance of the evidence, that claims 1, 2, and 13–15 of the ’524 patent are unpatentable. Petitioner’s Motion to Exclude is *denied-in-part* and *dismissed-in-part*. Patent Owner’s Motion to Exclude is *denied-in-part* and *dismissed-in-part*.

B. Related Proceedings

The parties indicate that the ’524 patent currently is involved in *SRC Labs, LLC et al. v. Microsoft Corp.*, Civil Action No. 2-18-cv-00321 (W.D. Wash.), which was transferred from *SRC Labs, LLC et al. v. Microsoft Corp.*, Civil Action No. 1-17-cv-01172 (E.D. Va.). Pet. 3; Prelim. Resp. 4–5. The following proceedings, before the Board, also involve the same parties: IPR2018-01594, IPR2018-01599, IPR2018-01600, IPR2018-01601, IPR2018-01602, IPR2018-01603, IPR2018-01605, IPR2018-01606, and IPR2018-01607.²

C. The ’524 Patent (Ex. 1001)

The ’524 patent is directed to the field of computer architecture, and, specifically, “a switch/network adapter port (‘SNAP’) for clustered computers employing a chain

2. IPR2018-01602 and IPR2018-01603 have been consolidated with IPR2018-01601. IPR2018-01606 and IPR2018-01607 have been consolidated with IPR2018-01605.

Appendix C

of multi-adaptive processors ('MAP[]') . . . in a dual in-line memory module ('DIMM') format." Ex. 1001, 1:29–37.

A block diagram of an exemplary MAP element is disclosed in Figure 3:




Fig. 3

120

Figure 3 illustrates memory bank 120 for a MAP element-based system computer architecture (not depicted in Figure 3). *Id.* at 4:60–64. “Each memory bank 120 includes a bank control logic block 122 bi-directionally coupled to the computer system trunk lines, for example, a 72 line bus 124.” *Id.* at 4:64–66. Memory array 130 is coupled to bank control logic 122 via bi-directional data bus 126 and address bus 128. *Id.* at 4:67–5:3. “MAP element 112 comprises a control block 132 coupled to the address bus 128.” *Id.* at 5:5–6. “[C]ontrol block 132 is also bi-directionally coupled to a user field programmable gate array ('FPGA),” and “user FPGA 134 is coupled directly to the data bus 126.” *Id.* at 5:6–10. The '524 patent discloses

Appendix C

that MAP element 112 has direct memory access (DMA) capability, which permits it to write to memory, and “it is possible to allow a MAP element 112 to feed results to another MAP element 112 through use of a chain port” because MAP element 112 receives operands via writes to memory. *Id.* at 5:50–54.

Computer system 200 including MAP element 212 in DIMM format is depicted in Figure 5:




Fig. 5

Figure 5 illustrates “computer system 200 includ[ing] one or more processors 202₀ and 202₁ which are coupled to an associated PC memory and I/O controller 204.” *Id.* at 7:42–45. “[C]ontroller 204 is . . . conventionally coupled to a number of DIMM slots 214 by means of a much higher bandwidth DIMM bus 216 capable of data transfer rates

Appendix C

of 2.1 GB/sec. or greater.” *Id.* at 7:56–59. DIMM MAP element 212 is associated with, or physically located within, one of DIMM slots 214, which may be coupled to another clustered computer MAP element by a cluster interconnect fabric connection 220 that is connected to MAP chain ports. *Id.* at 7:59–62, 7:65–8:1.

The ’524 patent discloses that because DIMM MAP element 212 is placed in one of DIMM slots 214, FPGA 134 of DIMM MAP element 212 accepts normal memory “read” and “write” transactions, and converts them to a format used by an interconnect switch or network. *Id.* at 8:13–17. According to the ’524 patent, however, “the electrical protocol of the DIMMs is such that once the data arrives at the receiver, there is no way for a DIMM module within the DIMM slots 214 to signal the microprocessor 202 that it has arrived.” *Id.* at 8:29–33. The “efforts of the processor 202 would have to be synchronized through the use of a continued polling of the DIMM MAP elements 212 to determine if data has arrived.” *Id.* at 8:29–36. According to the ’524 patent, this would consume the processor and much of its bandwidth, thereby stalling all other bus agents. *Id.* at 8:36–38.

To address this issue, the ’524 patent discloses connection 218 interconnecting DIMM MAP element 212 and PCI bus 210 such that DIMM MAP element 212 may generate communications packets and send them via PCI bus 210 to processor 202. *Id.* at 7:62–65; 8:39–43. According to the ’524 patent, because these packets would account for a very small percentage of the data moved, low bandwidth effects of PCI bus 210 would be minimized. *Id.* at 8:43–47.

*Appendix C**D. Illustrative Claim*

Petitioner challenges claims 1, 2, and 13–15 of the '524 patent. Pet. 12–80. Claim 1 is the only independent claim at issue. Claim 1 is illustrative of the challenged claims and is reproduced below:

1. A processor element for a memory module bus of a computer system, said processor element comprising:

a field programmable gate array configurable to perform an identified algorithm on an operand provided thereto and operative to alter data provided directly thereto on said memory module bus; and

a direct data connection coupled to said field programmable gate array for providing said altered data directly from said memory module bus to an external device coupled thereto.

Ex. 1001, 9:42–10:4.

E. The Alleged Grounds of Unpatentability

The information presented in the Petition sets forth grounds of unpatentability of claims 1, 2, and 13–15 of the '524 patent as follows (*see* Pet. 12–80):^{3,4}

3. Petitioner supports its challenge with the Declaration of Scott Hauck, Ph.D. Ex. 1003.

4. All references to 35 U.S.C. §§ 102, 103 herein are pre-AIA.

Appendix C

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
1, 2, 13–15	§ 102	Tsutsui ⁵
1, 2, 13–15	§ 103	Tsutsui
1, 2, 13–15	§ 103	Tsutsui, Tsutsui II ⁶
2	§ 103	Tsutsui, Stone, ⁷ with or without Tsutsui II
14	§ 103	Tsutsui, Collins, ⁸ with or without Tsutsui II
15	§ 103	Tsutsui, Hayashi, ⁹ with or without Tsutsui II

5. Akihiro Tsutsui et al., *YARDS: FPGA/MPU Hybrid Architecture for Telecommunication Data Processing*, Proceedings of Association for Computing Machinery / Special Interest Group for Design Automation (ACM/SIGDA) FPGA '97, 93–99, (1997) (Ex. 1007, “Tsutsui”).

6. A. Tsutsui et al., *Special Purpose FPGA for High-speed Digital Telecommunications Systems*, 1995 IEEE International Conference on Computer Design: VLSI in Computers & Processors, pp. 486–491 (1995) (Ex. 1009, “Tsutsui II”).

7. Harold S. Stone, MICROCOMPUTER INTERFACING 1–41 (Tom Robbins et al. 2d ed. 1983) (Ex. 1010, “Stone”).

8. U.S. Patent No. 5,671,355, issued September 23, 1997 (Ex. 1008, “Collins”).

9. K. Hayashi et al., *Reconfigurable Real-Time Signal Transport System using Custom FPGAs*, IEEE Symposium on FPGAs for Custom Computing Machines, IEEE (1995) (Ex. 1013, “Hayashi”).

Appendix C

II. ANALYSIS

A. Motions to Exclude

The party moving to exclude evidence bears the burden of proof to establish that it is entitled to the relief requested—namely, that the material sought to be excluded is inadmissible under the Federal Rules of Evidence. *See* 37 C.F.R. §§ 42.20(c), 42.62(a). For the reasons discussed below, Petitioner’s Motion is *denied-in-part* and *dismissed-in-part*, and Patent Owner’s Motion is *denied-in-part* and *dismissed-in-part*.

*1. Petitioner’s Motion to Exclude**a. Exhibits 2067 and 2086*

Petitioner moves to exclude the declaration of one of the named inventors of the ’524 patent, Jon Huppenthal (Exhibit 2067), “in its entirety as not being relevant to any issue on which trial has been instituted, and for lacking foundation, containing hearsay, and/or causing undue prejudice.” Pet. Mot. 3–6. Petitioner additionally moves to exclude portions of “Mr. Huppenthal’s declaration (Ex. 2067 ¶¶ 80, 82–86) due to his refusal to answer questions concerning those portions of the declaration.” *Id.* at 1–3 (citing Paper 51, 7–8).

Petitioner also moves to exclude a transcript (Exhibit 2086) of a deposition of Petitioner’s declarant from other *inter partes* reviews as “not being relevant to any issue on which trial has been instituted, for containing hearsay,

Appendix C

and/or causing undue prejudice.” *Id.* at 6–7. Petitioner argues that allowing the transcript in the record would be “highly prejudicial as it presents itself with the indicia of expert testimony while being totally devoid from the necessary context of the matter from which it originates.” *Id.* at 6. Patent Owner cites Exhibits 2067 and 2086 in its Response and Sur-Reply.

Petitioner’s Motion is dismissed as moot, as we do not rely on the testimony in a manner adverse to Petitioner in this Decision. As explained below, even if the testimony is considered, we are not persuaded by Patent Owner’s arguments regarding the state of the art or alleged nonobviousness of the challenged claims, and Patent Owner has not shown proof of secondary considerations that would support a conclusion of nonobviousness. *See supra* Sections II.C–II.D.

b. Exhibit 2102

Petitioner moves to exclude paragraph 133 of the declaration of Patent Owner’s declarant, Houman Homayoun, Ph.D., which refers to Exhibit 2067. Pet. Mot. 8. Because we do not exclude that exhibit, and do not rely on paragraph 133 of the Homayoun Declaration, we also dismiss as moot Petitioner’s Motion with respect to Exhibit 2102.

Appendix C

- c. Exhibits 2060, 2062–2064, 2066, 2072–2080, 2085, 2088, 2092, 2094, 2096–2101, 2103, 2104, 2106–2145, 2147, 2148, 2157, and 2160*

Petitioner moves to exclude a number of exhibits as “not being relevant to any issues on which trial has been instituted, lacking foundation, and/or causing undue prejudice” because the exhibits were not discussed substantively and/or cited in Patent Owner’s Response and Sur-Reply. Pet. Mot. 7–8. Petitioner’s Motion is dismissed as moot, as we do not rely on the exhibits in a manner adverse to Petitioner in this Decision. We note, however, that in evaluating Petitioner’s asserted grounds of unpatentability, we only consider substantive arguments made by the parties in their papers during trial (i.e., the Petition, Response, Reply, and Sur-Reply). To the extent a document is filed in the record but never discussed in a paper, there is no substantive argument pertaining to that document that can be considered.

d. Patent Owner’s Response

Petitioner moves to exclude portions of Patent Owner’s Response referring to the exhibits that Petitioner seeks to exclude. Pet. Mot. 8. Patent Owner’s Response is a paper with attorney arguments, not evidence that may be excluded.¹⁰ Further, we do not exclude any of the exhibits

10. Petitioner did not seek authorization to file a motion to strike Patent Owner’s Response. *See* Patent Trial and Appeal Board Consolidated Trial Practice Guide (Nov. 2019), 80–81, *available at* <https://www.uspto.gov/TrialPracticeGuideConsolidated> (“Trial Practice Guide”).

Appendix C

referred to in the identified portions of the Response. Petitioner's Motion is denied as to Patent Owner's Response.

2. *Patent Owner's Motion to Exclude*

a. *Exhibits 1007, 1036 and 1038*

Patent Owner first objected to the authenticity of Exhibit 1007 in its Patent Owner Response. PO Mot. 5 (citing PO Resp. 48). Now, Patent Owner moves to exclude Exhibits 1007, 1036, and 1038 as unauthenticated under Federal Rule of Evidence 901 because Petitioner's declarant, Stephen Trimberger, Ph.D., "admits that his original declaration did not properly authenticate EX1007." PO Mot. 5-7 (citing Ex. 1034 ¶ 7). Patent Owner argues that Petitioner attempts to authenticate Exhibit 1007 using evidence submitted for the first time in Petitioner's Reply, which violates Patent Owner's due process and administrative rights. *Id.* at 6. Patent Owner further argues that Exhibits 1036 and 1038, which are identical copies of Exhibit 1007, should be excluded "for the same reasons" as Exhibit 1007. *Id.* at 5, 7.

"Any objection to evidence submitted during a preliminary proceeding must be filed within ten business days of the institution of the trial. . . . The objection must identify the grounds for the objection with sufficient particularity to allow correction in the form of supplemental evidence." 37 C.F.R. § 42.64(b)(1); *see* 37 C.F.R. § 42.2 (defining a "preliminary proceeding" as the time period beginning "with the filing of a petition for

Appendix C

instituting a trial” and ending “with a written decision as to whether a trial will be instituted”). This process allows the party that originally submitted the evidence to attempt to cure the objection by serving supplemental evidence. 37 C.F.R. § 42.64(b)(2). If the submitting party does not serve supplemental evidence, or if the supplemental evidence does not cure the objection, “[a] motion to exclude evidence must be filed to preserve [the] objection. The motion must identify the objections in the record in order and must explain the objections.” 37 C.F.R. § 42.64(c).

“To satisfy the requirement of authenticating or identifying an item of evidence, the proponent must produce evidence sufficient to support a finding that the item is what the proponent claims it is.” Fed. R. Evid. 901(a). Certain evidence, though, is “self-authenticating” and “require[s] no extrinsic evidence of authenticity in order to be admitted.” Fed. R. Evid. 902.

Petitioner asserts that Patent Owner waived its objection by failing to timely object to Exhibit 1007 as lacking authenticity. Pet. Opp. Mot. 1. We agree. Patent Owner admits that it objected to the authenticity of Exhibit 1007 in its Response. PO Mot. 5. Patent Owner submitted its Response on August 5, 2019, more than ten business days after we instituted this proceeding on May 3, 2019. Therefore, Patent Owner’s objection to Exhibit 1007 based on authenticity is untimely.

Petitioner further asserts that Exhibits 1007, 1036, and 1038 are authentic. Pet. Opp. Mot. 1–4. First, Petitioner asserts Exhibit 1007 is authentic under Federal Rule of Evidence 901(a) because the record evidence

Appendix C

unambiguously shows that Exhibit 1007 is a copy of the Tsutsui reference. *Id.* at 2. Second, Petitioner asserts that Exhibits 1036 and 1038 are self-authenticating as ancient documents under Federal Rule of Evidence 901(b)(8) because each is “in a condition that creates no suspicion about its authenticity,” “was in a place where, if authentic, it would likely be,” and “is at least 20 years old when offered.” *Id.* at 2–3. Because Exhibits 1036 and 1038 are self-authenticating, and Exhibits 1007, 1036, and 1038 are identical, then Exhibit 1007 is also self-authenticating according to Petitioner. *Id.* Third, Petitioner argues that Exhibits 1007, 1036, and 1038 are Association for Computing Machinery (ACM) publications that each include “an ACM trade inscription, copyright symbol, and [International Standard Book Number (ISBN)],” and therefore, these documents self-authenticate under Federal Rules of Evidence 902(6) and 902(7). *Id.* at 3–4. Finally, Petitioner asserts that Exhibits 1007, 1036, and 1038 are also authenticated under Federal Rule of Evidence 901(b)(4) based on the totality of the circumstances and because Patent Owner does not identify anything to suggest the exhibits are not authentic. *Id.* at 4.

We agree with Petitioner that Exhibits 1007, 1036, and 1038 are self-authenticating under Federal Rules of Evidence 902(6) and 902(7), and that Exhibits 1036 and 1038 are ancient documents under Federal Rule of Evidence 901(b)(8) for the reasons stated by Petitioner. Exhibits 1007, 1036, and 1038 include an ACM trade inscription, copyright symbol, and ISBN number. Exhibits 1036 and 1038 are authentic as ancient documents, as they meet the requirements of Federal Rule of Evidence 901(b)(8).

Appendix C

Accordingly, we deny Patent Owner's Motion to Exclude Exhibits 1007, 1036, and 1038.

b. Exhibits 1034, 1035, and 1039

Patent Owner moves to exclude certain portions of Exhibits 1034, 1035, and 1039 because many portions of the reply declarations allegedly “are not based on the personal knowledge of the witness, are impermissible hearsay, and/or are impermissible expert opinions from a witness unqualified to provide such opinions.” PO Mot. 7–8. We are not persuaded. Patent Owner does not identify any particular “statement” in any of the exhibits that is being offered “to prove the truth of the matter asserted in the statement,” and thus fails to meet its burden to prove inadmissibility as hearsay. *See* Fed. R. Evid. 801(c); 37 C.F.R. § 42.20(c). Petitioner further establishes that the declarations were each based on the declarant's personal knowledge and qualifications. Pet. Opp. Mot. 5–7 (citing Ex. 1034 ¶¶ 2, 4, 9, 16, 30; Ex. 1035 ¶¶ 7–8, 17 (citing Ex. 1039 ¶¶ 2, 5–6)); *see* Ex. 1034 ¶ 2 (“I am submitting this declaration based on my own personal knowledge of the facts stated here . . .”). Accordingly, we deny Patent Owner's Motion to Exclude Exhibits 1034, 1035, and 1039.

c. Exhibit 1032

Patent Owner moves to exclude Exhibit 1032 because it is “irrelevant to this case” and because Petitioner “failed to properly authenticate” Exhibit 1032. PO Mot. 8–9. Patent Owner also asserts that, to the extent Petitioner attempts to meet its burden of proof using Exhibit 1032,

Appendix C

Petitioner's use of the document is impermissible hearsay. *Id.* at 9. Petitioner argues that Exhibit 1032 is the specific copy of the Tsutsui reference that Dr. Hauck relied on for purposes of his declaration, and was not initially filed with the Petition (but was filed with the Reply). Pet. Opp. Mot. 7–8. Petitioner also asserts that “extensive testimony entered into the record by Dr. Trimberger” authenticates Exhibit 1032. *Id.* at 8.

We agree with Petitioner that Patent Owner has not set forth sufficient argument and evidence to exclude Exhibit 1032 as irrelevant. We are also persuaded that Dr. Trimberger's testimony authenticates Exhibit 1032. *See, e.g.*, Ex. 1034 ¶¶ 7 (“I understand that the version of the *Tsut[su]i* paper which I refer to as Exhibit 1007 (included here as EX1032 and referred to as *Tsutsui-1032*) was not included with my original declaration, and a different version of the same paper was included in its place.”), 5 (“*Tsutsui-1007* and *Tsutsui-1032* are substantively identical.”). Accordingly, Patent Owner's Motion to Exclude Exhibit 1032 is denied.¹¹

d. Exhibits 1033, 1036, 1038, 1040, 1041, 1043–1050, 1053–1057

Patent Owner moves to exclude Exhibits 1040, 1041, 1043, and 1045–1049 because none of these exhibits “are cited, discussed, or relied upon by any expert witness or

11. To be clear, Petitioner's asserted grounds of unpatentability are based on the version of Tsutsui filed as Exhibit 1007, not Exhibit 1032. *See* Dec. 7–8, 29; Pet. 4; Pet. Opp. Mot. 8.

Appendix C

fact witness in this case.” PO Mot. 9. Patent Owner also moves to exclude Exhibits 1033, 1036, 1038, 1044, 1050, and 1053–1057 because they are “not cited in any of Petitioner’s papers.” *Id.* at 12. We see no basis to exclude the exhibits for that reason and dismiss Patent Owner’s request, but again note that in evaluating Petitioner’s asserted grounds of unpatentability, we only consider substantive arguments made by the parties in their papers during trial (i.e., the Petition, Response, Reply, and Sur-Reply). To the extent a document is filed in the record but never discussed in a paper, there is no substantive argument pertaining to that document that can be considered.

Patent Owner further moves to exclude Exhibits 1040, 1043, and 1049 as unauthenticated under Federal Rule of Evidence 901. *Id.* at 10–11. Petitioner argues that “Patent Owner identifies nothing about the documents themselves that brings into question their authenticity,” and “Patent Owner bears the burden as movant to demonstrate these documents are not authentic.” Pet. Opp. Mot. 10. Petitioner further asserts that Exhibits 1040 and 1049 are IEEE publications that each include “a trade inscription, copyright symbol, and ISBN,” and, therefore, these documents self-authenticate under Federal Rules of Evidence 902(6) and 902(7). *Id.* at 10–11. Petitioner further argues that Exhibits 1040, 1043, and 1049 are authenticated under Federal Rule of Evidence 901(b)(4) “based on the totality of the circumstances based on their appearance alone, as nothing about those exhibits suggests that [they] are not what the[y] purport to be, and Patent Owner never identifies anything to suggest otherwise.” *Id.* at 11. Petitioner also argues that Exhibit

Appendix C

1040 is an ancient document because it is over twenty years old and meet the requirements of Federal Rule of Evidence 901(b)(8). *Id.* at 11. We agree with Petitioner that Patent Owner has not set forth sufficient argument and evidence to exclude Exhibits 1040, 1043, and 1049 as unauthenticated, and we further agree that Exhibit 1040 is self-authenticating and an ancient document for the reasons stated by Petitioner.

Patent Owner further moves to exclude Exhibits 1040, 1041, 1043, and 1045–1049 as containing inadmissible hearsay under Federal Rule of Evidence 802. PO Mot. 10–11. Patent Owner states that Petitioner in its Reply “cites each of these documents to prove the truth of technical matters allegedly asserted in such documents, *i.e.* to support Petitioner’s specific factual assertions regarding a technical issue.” *Id.* at 10. We are not persuaded. Patent Owner does not identify any particular “statement” in any of the exhibits that is being offered “to prove the truth of the matter asserted in the statement,” and thus fails to meet its burden to prove inadmissibility as hearsay. *See* Fed. R. Evid. 801(c); 37 C.F.R. § 42.20(c). Even if Patent Owner had done so, Petitioner cites the exhibits to show what a person of ordinary skill in the art would have known at the time of the ’524 patent about the technical features and developments in the pertinent art. Pet. Opp. Mot. 12 (citing Ex. 1040 and 1041). The exhibits are not being offered for the truth of any particular matter discussed in the references. Accordingly, we deny Patent Owner’s Motion to Exclude Exhibits 1040, 1041, 1043, and 1045–1049.

Appendix C

Patent Owner appears to move to exclude uncited portions of deposition testimony from Mr. Huppenthal and Dr. Homayoun (Exhibits 1033 and 1055–1057) because “Petitioner has failed to establish [their] relevance to this case.” PO Mot. 13. We see no basis to exclude the exhibits for that reason and dismiss Patent Owner’s request, but again note that in evaluating Petitioner’s asserted grounds of unpatentability, we only consider substantive arguments made by the parties in their papers during trial (i.e., the Petition, Response, Reply, and Sur-Reply). To the extent a document is filed in the record but never discussed in a paper, there is no substantive argument pertaining to that document that can be considered.

e. Exhibits 1051 and 1052

Patent Owner moves to exclude certain portions of the transcripts of two depositions of Steven A. Guccione, Ph.D., because the questions asked were “vague and ambiguous.” PO Mot. 11–12 (citing Ex. 1052, 51:1–7; Ex. 1051, 38:14–39:6). Patent Owner’s Motion is dismissed as moot, as we do not rely on the disputed portions of the testimony in rendering our Decision.

Additionally, Patent Owner appears to move to exclude uncited portions of deposition testimony from Dr. Guccione because “Petitioner has failed to establish [their] relevance to this case.” *Id.* at 13. We see no basis to exclude the exhibits for that reason and dismiss Patent Owner’s request, but again note that in evaluating Petitioner’s asserted grounds of unpatentability, we only consider substantive arguments made by the parties in

Appendix C

their papers during trial (i.e., the Petition, Response, Reply, and Sur- Reply). To the extent a document is filed in the record but never discussed in a paper, there is no substantive argument pertaining to that document that can be considered. Accordingly, Patent Owner's Motion with respect to uncited portions of Dr. Guccione's testimony is dismissed.

B. Level of Ordinary Skill in the Art

In determining the level of ordinary skill in the art for a challenged patent, we look to "1) the types of problems encountered in the art; 2) the prior art solutions to those problems; 3) the rapidity with which innovations are made; 4) the sophistication of the technology; and 5) the educational level of active workers in the field." *Ruiz v. A.B. Chance Co.*, 234 F.3d 654, 666–667 (Fed. Cir. 2000); *see also Custom Accessories, Inc. v. Jeffrey-Allan Indus., Inc.*, 807 F.2d 955, 962–963 (Fed. Cir. 1986). "Not all such factors may be present in every case, and one or more of them may predominate." *Ruiz*, 234 F.3d at 666–667; *see also Custom Accessories*, 807 F.2d at 963.

Petitioner asserts that a person of ordinary skill in the art, at the time of the '524 patent, would have had "a Bachelor's degree in Electrical Engineering, Computer Engineering, Computer Science, or in a related field, and four years of experience with the design or use of field programmable gate array based systems or network adapters." Pet. 5 (citing Ex. 1003 ¶ 41). Petitioner alternatively argues that a person of ordinary skill in the art "would have had an advanced degree in one of those

Appendix C

fields and two years of related experience.” *Id.* Moreover, such a person, according to Petitioner, would have been knowledgeable about computer architectures and how FPGAs could be included in them. *Id.* at 5. Such a person also would have been knowledgeable about software algorithms that could be implemented on FPGAs and how to configure FPGAs to carry out such implementation. *Id.* (citing Ex. 1003 ¶ 41).

Patent Owner disputes Petitioner’s assessment of the level of ordinary skill in the art, but does not provide a proposed level of ordinary skill in the art that we should apply in this proceeding.¹² PO Resp. 30–33. Patent Owner contends Petitioner’s assessment is incorrect because a person of ordinary skill in the art would not have had the detailed FPGA knowledge assumed by Petitioner’s definition. *Id.* Patent Owner further contends the technical problem the ’524 patent sought to address is in the field of High-Performance Computing (“HPC”). *Id.* at 3. Concatenating the experience Petitioner argues an artisan would have had with Patent Owner’s assessment of the technical field, Patent Owner argues the result would have been a “mythical person” knowledgeable about both FPGAs and HPC. *Id.* at 30–33. According to Patent Owner, it would have been rare to find a person knowledgeable in both disciplines, and that any such person would have had more education and experience than that proposed by Petitioner. *Id.*

12. Patent Owner’s declarant, Dr. Homayoun, agrees with “the level of education and skill” that a person of ordinary skill in the art would have according to Petitioner. Ex. 2102 ¶ 123.

Appendix C

Patent Owner's assertions are unavailing. Patent Owner relies on the Declaration of Dr. Guccione to support its assertion that there were few engineers at the time of the '524 patent who knew how to program FPGAs to run software algorithms. PO Resp. 31–32 (citing Ex. 2146 ¶¶ 87, 193–198). Dr. Guccione testifies, for example, that implementing algorithms in FPGAs (i.e., hardware) is more complex than implementing algorithms in software, and therefore requires a different skillset. Ex. 2146 ¶ 87. Dr. Guccione's testimony, however, is at odds with what is reflected in the prior art of the period. *See In re GPAC Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995); *In re Oelrich*, 579 F.2d 86, 91 (CCPA 1978). As Petitioner points out, "numerous prior art references disclose computer architectures with multiple processing elements and FPGAs used for network interface," including the main references on which Petitioner relies, Tsutsui and Tsutsui II. Reply 4 (citing Exs. 1007, 1009, 1040, 1041). The art of record therefore indicates that Petitioner's assessment of the education and experience required in the field of programming FPGAs is accurate.

Also, we find insufficient evidence that the level of ordinary skill in the art would have required specialized knowledge in the field of High Performance Computing. Patent Owner cites to portions of the '524 patent that describe enhancing overall processing speed in a multiprocessor computer architecture incorporating a number of memory algorithm processors, but has not shown why these systems are HPC architectures, much less limited to being HPC architectures. PO Resp. 4–5. The '524 patent does not limit the invention to HPC, but rather more broadly states

Appendix C

The present invention relates, in general, to the field of computer architectures incorporating multiple processing elements. More particularly, the present invention relates to a switch/network adapter port (“SNAP”) for clustered computers employing a chain of multi-adaptive processors (“MAP[.]”) . . . in a dual in-line memory module (“DIMM”) format to significantly enhance data transfer rates over that otherwise available from the peripheral component interconnect (“PCI”) bus.

Ex. 1001, 1:29–37. This description, however, is non-limiting, and indicates a shortcoming of systems with “multiple processing elements,” namely very high latency due to the location of chips below the processor bus. Ex. 1001, 1:51–56. This evidence does not suggest the ’524 patent is directed only to HPC systems.

We also find unavailing Patent Owner’s argument that in assessing the level of skill in the art, Petitioner erred by focusing on the technical solution—using programmable computing elements—rather than on the technical problem to be solved. PO Resp. 32; Sur-Reply 2–3. Patent Owner provides no legal basis to support its contention that the only appropriate factor to consider in determining the hypothetical person of ordinary skill in the art is the technical problem to be solved. Contrary to Patent Owner’s assertion, several factors may be considered, including the types of problems encountered in the art, prior art solutions to those problems, rapidity with which innovations are made, sophistication of the

Appendix C

technology, and educational level of active workers in the field. *Ruiz*, 234 F.3d at 666–667; *Custom Accessories*, 807 F.2d at 963. Moreover, such factors may or may not be present in every case, and one or more of the factors may predominate. *Ruiz*, 234 F.3d at 666–667; *see also Custom Accessories*, 807 F.2d at 963. Here, we have considered the entirety of the '524 entirety disclosure, and the prior art of record. Further, to the extent Patent Owner disputes the number of *actual* persons with the proposed education and technical experience at the time of the '524 patent, we note that “[t]he person of ordinary skill is a *hypothetical* person who is presumed to be aware of all the pertinent prior art.” *Custom Accessories*, 807 F.2d at 962 (emphasis added).

For the foregoing reasons, we adopt Petitioner’s proposal as to the level of ordinary skill in the art. We note that had we not adopted Petitioner’s proposal that the relevant experience be specific to FPGAs—in particular, had we not articulated a specialty within the field of computer architecture—it would not have altered the outcome of this Decision. As we noted above, the level of skill in the art also may be reflected in the prior art. *See GPAC*, 57 F.3d at 1579; *Oelrich*, 579 F.2d at 91. Here, *Tsutsui* and *Tsutsui II*, for example, reflect the appropriate level of skill in the art. We also note that to the extent we determine Petitioner has demonstrated certain claims are unpatentable as obvious, below, we make such findings under an even lower standard than that suggested by Patent Owner’s arguments—namely, we determine the claims would have been obvious even without assuming the level of skill in the art requiring specialized knowledge in the field of High-Performance Computing. *See Kinetic*

Appendix C

Concepts, Inc. v. Smith & Nephew, Inc., 688 F.3d 1342, 1366 (Fed. Cir. 2012) (“[I]t is generally easier to establish obviousness under a higher level of ordinary skill in the art.”).

C. Claim Construction

The parties agree that the ’524 patent has expired. *See* Pet. 12; Prelim. Resp. 18. Accordingly, we apply the district court claim construction standard. *See* 37 C.F.R. § 42.100(b). In district court, claim terms are given their plain and ordinary meaning as would be understood by a person of ordinary skill in the art at the time of the invention and in the context of the entire patent disclosure. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1313 (Fed. Cir. 2005) (en banc). “There are only two exceptions to this general rule: 1) when a patentee sets out a definition and acts as his own lexicographer, or 2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). We determine that the following terms and limitations require construction.

1. “memory module bus”

Claim 1 recites a “memory module bus.” Ex. 1001, 9:42. Based on the arguments presented in the Petition and Preliminary Response, we were persuaded by Petitioner and preliminarily construed “memory module bus” in the Decision on Institution to mean “a bus used to communicate with a memory module.” Dec. 10–11; *see* Pet. 13.

Appendix C

Patent Owner argues that “memory module bus” should be construed to mean “[a] bus designed to comply with applicable standards for connecting the main system processor(s) to a standardized memory device module according to the physical and logical connection protocols that are required by the memory module’s specifications, which does not include a PCI bus or other I/O bus that is used to connect the system processor(s) to peripheral or external devices.” PO Resp. 33–34; *see id.* at 39–40 (citing Ex. 2146 ¶ 216). Patent Owner asserts that the plain and ordinary meaning, and the meaning in the ’524 patent specification, of “memory module bus” only includes “a bus that is intended to communicate with a memory module *according to the adopted standards* for interfacing with the memory module.” *Id.* at 35; *see id.* at 39–40 (citing Ex. 2146 ¶ 216); *see* PO Sur-Reply 16–18. Patent Owner argues that the ’524 patent specification explains that “the MAP element [is] inserted into a DIMM-standard memory slot.” PO Resp. 40 (citing Ex. 1001, Fig. 5, 8:4–26). Patent Owner further argues that Petitioner’s expert, Dr. Hauck, testifies that the ’524 patent disclaims a PCI bus as a “memory module bus.” *Id.* at 35–36 (citing Ex. 1003 ¶¶ 76–82).¹³ Patent Owner argues that because Dr. Hauck’s construction excludes a PCI bus, similar peripheral buses

13. Patent Owner also requests sanctions on the basis that Petitioner exhibited a lack of good faith and violated its duty of candor by failing to alert the Board to the testimony of Dr. Hauck which Patent Owner alleges is inconsistent with Petitioner’s proposed construction of “memory module bus.” PO Resp. 35–36 (citing Ex. 1003 ¶¶ 76–81); PO Sur-Reply 27–28. We dismiss Patent Owner’s request for sanctions for the reasons discussed below. *See infra* Section G. Moreover, *all of* the cited testimony in Dr. Hauck’s declaration (Ex. 1003) was in the record and available to the Board at the time of institution.

Appendix C

with similar characteristics should also be excluded. *Id.* at 36–37.

In response, Petitioner argues that Patent Owner’s proposed construction is substantially different than what was proposed by Patent Owner in its Preliminary Response and in parallel litigation between the parties. Pet. Reply 5. Petitioner argues that nothing in the intrinsic or extrinsic record supports Patent Owner’s contention that the “memory module bus” must be “designed to comply with applicable standards.” *Id.* at 6–7. Petitioner also asserts that Patent Owner’s proposed construction improperly extends the disclaimer of a PCI bus to other types of peripheral buses. *Id.* at 8–9. Petitioner further asserts that the PCI bus disclaimer is irrelevant, as the relied-upon prior art, Tsutsui, does not disclose a PCI bus. *Id.* at 8; *see also* Tr. 19:12–20:3, 21:1–11, 22:18–23:10 (explaining the same).

We are persuaded that “memory module bus” should be construed to mean “a non-PCI bus used to communicate with a memory module.” This construction is consistent with the plain and ordinary meaning of “memory module bus,” is consistent with the ’524 patent specification, and incorporates the ’524 patent specification’s disclaimer of a standard PCI bus. *See* Ex. 1001, 7:56–59; Ex. 1003 ¶¶ 77–78 (citing Ex. 1001, Abstract, 1:29–67, 2:1–13, 4:7–13, 7:34–41, 8:4–29). We limit the scope of “memory module bus” to exclude a PCI bus given the specific statements in the ’524 patent specification and the parties’ agreement that the ’524 patent disclaimed a PCI bus. PO Resp. 55 (citing Ex. 1003 ¶¶ 76–82; Ex. 2146 ¶ 210); *see also* Tr. 19:13–20:3 (“[I]f the Board wants to find disclaimer of a

Appendix C

PCI bus, we made that argument in District Court. We believe it's disclaimed . . ."). Specifically, the '524 patent specification explains how a PCI bus is different from the bus used in the disclosed invention and why the disclosed bus solves the problems that a PCI bus had at the time. *See, e.g.*, Ex. 1001, 1:29–67, 2:1–38, 4:7–13, 7:34–41, 8:4–29, code (57).

We are unpersuaded by Patent Owner's attempt to limit the scope of a "memory module bus" to require a connection "according to the adopted standards for interfacing with the memory module," or to extend the disclaimer beyond PCI buses. PO Resp. 35–37 (emphasis omitted). We are not persuaded that the claims of the '524 patent and the '524 patent specification support such a construction. The specification states that the disclosed arrangement has improved data transfer rates "over that otherwise available" from a PCI bus, but does not distinguish any other type of bus or generalize beyond PCI buses. *See, e.g.*, Ex. 1001, 1:31–37, 2:8–12, 4:7–13, 7:34–41, code (57). The specification also explains that controller 204 is "*conventionally coupled* to a number of DIMM slots 214." Ex. 1001, 7:56–57 (emphasis added). Thus, Patent Owner's proposed construction seeking to extend the explicit disclaimer to "other types of peripheral buses" with the "same characteristics" as a PCI bus does not find support in the intrinsic record. *See* PO Resp. 36–37. In addition, we are not persuaded that a construction using the term "applicable standards," as Patent Owner proposes, would be appropriate, as it is unclear what standards would or would not be encompassed within such a construction.

Appendix C

Accordingly, we construe “memory module bus” to mean “a non-PCI bus used to communicate with a memory module.”

2. *“a direct data connection coupled to said field programmable gate array for providing said altered data directly from said memory module bus to an external device coupled thereto”*

Claim 1 recites “a direct data connection coupled to said field programmable gate array for providing said altered data directly from said memory module bus to an external device coupled thereto.” Ex. 1001, 10:1–4. In the Decision on Institution, we preliminarily construed this limitation to mean that the field programmable gate array receives and operates on data directly from the memory module bus, producing altered data and providing that altered data directly to the external device via a direct data connection. Dec. 14.

The parties do not further argue the meaning of this limitation. We see no reason to depart from our construction in our Decision on Institution. Accordingly, we construe this limitation to mean that the field programmable gate array receives and operates on data directly from the memory module bus, producing altered data and providing that altered data directly to the external device via a direct data connection.

3. *“control connection”*

Claim 2 recites a “control connection.” Ex. 1001, 10:6. We construe this limitation in order to address Patent

Appendix C

Owner’s argument that Tsutsui fails to disclose it. *See infra* Section II.D.2.b.iv.; PO Resp. 63. Patent Owner contends that “control connection” should be construed to mean “a bus or data communication channel separate from the memory module bus.” PO Resp. 43 (emphasis omitted). Patent Owner contends that the ’524 patent discloses that a “control connection” is “either a PCI bus, graphics bus, accelerated graphics port (AGP) bus, or a system maintenance bus.” *Id.* at 41–42 (citing Ex. 1001, 7:34–41, 10:10–24, Fig. 5). Patent Owner argues that a person with ordinary skill in the art would have understood that the ’524 patent specification “describes the ‘memory module bus’ and ‘control connection’ as communication paths that are being used for fundamentally different things.” *Id.* at 42 (citing Ex. 2146 ¶ 229). Petitioner does not proffer a definition for this term. *See generally* Pet. Reply 4–11. We are persuaded by Patent Owner, and construe “control connection,” consistent with the ’524 patent specification, to mean “a bus or data communication channel separate from the memory module bus.”

4. “processor of said computer system” and
“plurality of processors”

Claim 2 recites “a processor of said computer system” and claim 14 recites “a plurality of processors.” Ex. 1001, 10:7, 10:39–40. We construe these limitations in order to address Patent Owner’s argument that Tsutsui fails to disclose them. *See infra* Section II.D.2.b.v.; PO Resp. 63–64. Patent Owner argues that “processor of said computer system” should be construed to mean “the main system processor consisting of a microprocessor or an FPGA programmed to function as a von Neumann-

Appendix C

style stored memory instruction processor.” PO Resp. 43 (citing Ex. 2146 ¶¶ 231–233) (emphasis omitted). Patent Owner argues that the claim language differentiates the “processor of said computer system” from the “processor element” of claim 1. *Id.* Patent Owner further argues that a person of ordinary skill in the art in 2001 would understand “processor of said computer system” to refer to “main system processors coupled to the main memory modules via a memory module bus.” *Id.* (citing Ex. 2146 ¶¶ 231–233).

Petitioner argues that Patent Owner has provided no intrinsic evidence to support its proposed construction. Pet. Reply 10–11. Petitioner contends that the ordinary meaning of “processor of said computer system,” and as discussed in the ’524 patent specification, would encompass “any processor in the computer system.” *Id.* (emphasis omitted); Pet. 21–22 (citing Ex. 1003 ¶¶ 344, 346–349; Ex. 1001, Title, 1:29–37, 3:40–42, 5:17–24, Fig. 3, 3:66–4:2, 6:19–38).

We agree with Petitioner that the plain and ordinary meaning of “a processor in the computer system” is “any processor in the computer system.” Pet. Reply 10–11. Although we agree with Patent Owner that the claim language demonstrates that the “processor of said computer system” is distinct from the “processor element” of claim 1, we are not persuaded that either the intrinsic or extrinsic record supports limiting the scope of this limitation to be “the main system processor” or a “processor consisting of a microprocessor or an FPGA programmed to function as a von Neumann- style stored memory instruction processor.” Patent Owner’s sole

Appendix C

support for its proposed construction is the testimony of Dr. Guccione, and we do not credit that testimony given the lack of any support in the '524 patent claims, written description, or prosecution history. Accordingly, we construe “a processor in the computer system” to mean “any processor in the computer system that is not the processor element of claim 1.”

Patent Owner further argues that “plurality of processors,” as recited in claim 14, should be construed to mean “more than one processor.” PO Resp. 43–44. Petitioner does not dispute this proposed construction. *See generally* Pet. Reply 4–11. We agree with Patent Owner that the plain and ordinary meaning of a “plurality of processors” is “more than one processor.” Accordingly, we construe “a plurality of processors” to mean “more than one processor.”

5. Remaining Terms and Limitations

We determine that no other express claim construction analysis of any claim term is necessary. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (holding that only terms in controversy must be construed and only to the extent necessary to resolve the controversy) (citing *Vivid Techs., Inc. v. Am. Sci. & Eng'g*, 200 F.3d 795, 803 (Fed. Cir. 1999)).

D. Anticipation of Claims 1, 2, and 13–15 by Tsutsui

Petitioner contends that claims 1, 2, and 13–15 of the '524 patent are unpatentable under 35 U.S.C.

Appendix C

§ 102 as anticipated by Tsutsui. Pet. 30–53. For the reasons discussed below, we determine Petitioner has demonstrated by a preponderance of the evidence that claims 1, 2, and 13–15 of the '524 patent are unpatentable under 35 U.S.C. § 102 as anticipated by Tsutsui.

1. *Tsutsui (Ex. 1007)*¹⁴

Tsutsui is directed to a “system architecture applicable to high-performance and flexible transport data processing which includes complex protocol operation and a network control algorithm.” Ex. 1007, 1. Tsutsui discloses an FPGA and micro-processing unit (“MPU”)

14. Petitioner asserts that Tsutsui qualifies as a prior art printed publication under 35 U.S.C. § 102(a). Pet. 22; Reply 13–16. Patent Owner contends that “Petitioner did not Authenticate EX1007 (Tsutsui []),” and that “[a]s a result, Petitioner has failed to prove that EX1007 is a prior art publication under 35 U.S.C. § 102(b).” *Id.*; PO Sur-Reply 13–14 (arguing that Petitioner “failed to authenticate EX1007”). Thus, Patent Owner’s sole contention that Tsutsui (i.e., the document filed as Exhibit 1007) is not a prior art printed publication is that Petitioner failed to authenticate it. Patent Owner confirmed that to be the case during the hearing. Tr. 39:25–40:12 (in response to a question asking whether Patent Owner is contesting authentication or qualification as a printed publication, stating: “It’s really just authentication In this case we’re not conceding that it was publicly available but we haven’t contested it but the authentication, yes, we’re pressing that argument.”). We address Patent Owner’s authentication arguments above in deciding Patent Owner’s Motion to Exclude Tsutsui as unauthenticated under Federal Rule of Evidence 901. *See supra* Section II.A.2; PO Mot. 5–7. Petitioner has proven, for the reasons stated in its papers, which are supported by the evidence of record cited in those papers, that Tsutsui is a prior art printed publication under 35 U.S.C. § 102(a). *See* Pet. 22; Reply 13–16.

Appendix C

system named Yet Another Re-Definable System (“YARDS”). *Id.* YARDS comprises three programmable devices, thereby providing high flexibility. *Id.* The system “supports various styles of coupling between the FPGAs and the MPU.” *Id.* The system architecture is suitable for implementing flexible and real-time transport data processing operations. *Id.*

*2. Analysis**a. Petitioner’s Contentions*

Independent claim 1 recites “[a] processor element for a memory module bus of a computer system.” Ex. 1001, 9:42–43. Petitioner contends that Tsutsui discloses “multiple FPGAs, the VME-I/F, the Interrupt Signals from the FPGAs to the MPU, and the Direct-I/O connection,” which comprise a “processing element.” Pet. 30–31 (citing Ex. 1007, Fig. 1). Petitioner argues that the “‘FPGAs’ within YARDS can be ‘treated as co-processors or special peripheral devices of the MPU.’” *Id.* at 31 (quoting Ex. 1007, 2; citing Ex. 1003 ¶¶ 188–189). Petitioner further argues that the “processing element” is “coupled to the Local Bus of YARDS,” and, therefore, is “for a memory module bus.” *Id.* (citing Ex. 1007, 3; Ex. 1003 ¶¶ 190, 197) (emphasis omitted).

Claim 1 further recites “a field programmable gate array configurable to perform an identified algorithm on an operand provided thereto and operative to alter data provided directly thereto on said memory module bus.” Ex. 1001, 9:44–47. Petitioner contends that Tsutsui discloses

Appendix C

FPGAs that are configured to perform processing functions. Pet. 32 (citing Ex. 1007, 2–3). Petitioner further argues that Tsutsui discloses that FPGAs, during operations and maintenance processing (“OAM”) cell operations, determine the appropriate action for extracted cell data, including determining the control method using a control algorithm and accessing a database. *Id.* at 34 (citing Ex. 1007, 3, Fig. 16; Ex. 1003 ¶ 208). Petitioner also argues that Tsutsui discloses FPGAs that perform an algorithm on a received operand, including performing hash functions on received data, processing data provided directly from memory, and “re-shaping” data before transmission, which “would necessarily include ‘*identified algorithms*’ that operate on data.” *Id.* at 32–33, 35–36 (citing Ex. 1007, 1, 3; Ex. 1003 ¶¶ 205–207, 217–220). According to Petitioner, Tsutsui’s “re-shaping” of data includes “reordering and changing the spacing of the ATM cells within a stream.” *Id.* at 35 (citing Ex. 1007, 1, 3; Ex. 1021, 2:36–55; Ex. 1003 ¶¶ 211–214). Petitioner argues that the “re-shaping” of data constitutes altering data. *Id.* (citing Ex. 1021, 2:36–56; Ex. 1003 ¶¶ 211–214). Petitioner also argues that there is “a direct connection between the Local Bus and the FPGA.” *Id.* (citing Ex. 1007, 2).

Claim 1 also recites “a direct data connection coupled to said field programmable gate array for providing said altered data directly from said memory module bus to an external device coupled thereto.” Ex. 1001, 10:1–4. As discussed above, we construe this limitation to mean that the field programmable gate array receives and operates on data directly from the memory module bus, producing altered data and providing that altered data directly

Appendix C

to the external device via a direct data connection. *See supra* Section II.B.2. Petitioner argues that, under this construction, Tsutsui discloses that the “direct I/O channel enables a direct exchange of signals between other devices and the FPGA card,” where the Direct I/O channel is directly connected to an FPGA. Pet. 41 (quoting Ex. 1007, 2; citing Ex. 1007, Fig. 1).

Notwithstanding Patent Owner’s arguments, which we address below, we are persuaded by Petitioner that Tsutsui discloses the limitations of claim 1. As discussed above, Tsutsui discloses an FPGA and MPU system named YARDS. Ex. 1007, 1; *see supra* Section II.D.1. The main parts of the system include programmable switching devices and 2-port SRAMs. Ex. 1007, 2. YARDS supports three different styles of connection between the FPGAs and MPUs: a bus, a direct interrupt, and a 2-port SRAM channel. *Id.* Using the local bus, the transport data stream is directly input into the FPGA, which executes operations and transfers the results to the main memory. *Id.* at 3. After the MPU completes high level protocol operations, the FPGA “re-shapes” the data as the output transport data stream. *Id.* The system further includes external interfaces via a VME-Bus and Direct-I/O. *Id.* at 2. Accordingly, we are persuaded by Petitioner that Tsutsui discloses the claimed FPGA and direct data connection of the processor element for a memory module bus of a computer system.

Petitioner provides a similar analysis for dependent claims 2 and 13–15, explaining exactly how Tsutsui discloses the limitations of the claims. Pet. 30–53. We

Appendix C

find that Petitioner supports its arguments with credible evidence, and we are persuaded by Petitioner that Tsutsui discloses the limitations of claims 2 and 13–15. *See id.* Accordingly, notwithstanding Patent Owner’s arguments, which we address below, we are persuaded that Tsutsui anticipates claims 1, 2, and 13–15.

b. Patent Owner’s Arguments

Patent Owner argues that Tsutsui fails to disclose the following limitations: (i) “memory module bus” (claim 1), (ii) “data provided directly thereto on said memory module bus” (claim 1), (iii) “direct data connection” (claim 1), (iv) “control connection for indicating to a processor of said computer system” (claim 2), and (v) “a plurality of processors” (claim 14). PO Resp. 55–64. We address these arguments in turn.

i. “memory module bus”

Patent Owner argues that Tsutsui fails to disclose the claimed “memory module bus,” as Patent Owner construes the term. PO Resp. 55–57; PO Sur-Reply 20–23. Specifically, Patent Owner argues that the proper claim construction for “memory module bus” excludes a PCI bus (*see supra* Section II.C.1), and argues that Tsutsui’s Local Bus is “functionally equivalent” to a PCI bus. PO Resp. 55–56 (citing Ex. 1003 ¶¶ 76–82; Ex. 2146 ¶¶ 210, 264–265; Ex. 1007, 3). Patent Owner further argues that Tsutsui’s Local Bus is not a “memory module bus” because there is no evidence that it is used to communicate with a memory module. *Id.* at 56–57 (citing Ex. 1007, Fig.

Appendix C

1); PO Sur-Reply 21–22. Patent Owner argues that the Local Bus “has . . . many components attached to it” and “supports only one bus master at a time,” requiring “bus contention,” where local bus congestion “is the same technical problem addressed by the patent.” PO Resp. 56 (citing Ex. 2146 ¶ 265, Ex. 1007, 3) (emphasis omitted); PO Sur-Reply 20–21.

In response, Petitioner asserts that Patent Owner’s arguments are both unsupported and premised on an improper claim construction. Pet. Reply 17–18. Petitioner further argues that Patent Owner fails to explain the alleged equivalence between Tsutsui’s Local Bus and a PCI bus. *Id.* at 17 (citing PO Resp. 55). Petitioner argues that Patent Owner’s arguments relating to bus contention and the Local Bus “suffer[ing] from congestion” are not even requirements of Patent Owner’s proposed claim construction. *Id.* Finally, Petitioner argues that Tsutsui satisfies Patent Owner’s proposed construction because Tsutsui’s Local Bus is not a PCI or I/O bus that connects to peripheral/external devices. *Id.* at 17–18 (citing Pet. 30–31; Pet. Reply 17; Ex. 2156, 447; Ex. 1003 ¶¶ 130, 197; Ex. 1007, Fig. 1).

We are not persuaded by Patent Owner’s arguments. We construe “memory module bus” to mean “a non-PCI bus used to communicate with a memory module.” *See supra* Section II.C.1. Tsutsui does not describe the Local Bus as a PCI bus. Tsutsui discloses a “processing element” that includes multiple FPGAs, the VME-I/F, Interrupt Signals from the FPGAs to the MPU, and the Direct-I/O connection, and discloses that the “processing element”

Appendix C

is “coupled to the Local Bus of YARDS.” Ex. 1007, 1–3, Figs. 1, 13; Pet. 31 (citing Ex. 1003 ¶¶ 190, 197). Tsutsui discloses FPGAs and DRAM (SIMM) both coupled to the Local Bus. Ex. 1007, 1–3. Further, we are not persuaded by Patent Owner’s argument that Tsutsui’s Local Bus is a PCI bus because Patent Owner merely alleges that the Local Bus “shares characteristics that a [person of ordinary skill in the art] would recognize are indicative of a PCI bus,” but fails to set forth persuasive evidence and argument to establish that Tsutsui’s Local Bus is functionally the same as a PCI Bus. PO Resp. 55–56 (citing Ex. 2146 ¶ 264). That is, even if the Local Bus shared certain characteristics with a PCI bus, that does not mean that the Local Bus is the same thing as a PCI bus. As such, we are not persuaded that Tsutsui’s Local Bus is a PCI Bus. We determine that Tsutsui’s DRAM is memory and is connected to the Local Bus, and, therefore, the Local Bus is a “memory module bus.” Accordingly, we are not persuaded by Patent Owner’s argument that Tsutsui fails to disclose a “memory module bus.”

ii. “data provided directly thereto on said memory module bus”

Patent Owner further argues that Tsutsui fails to disclose “data provided directly thereto on said memory module bus” because “the FPGAs in Tsutsui never receive data directly from the memory module bus.” PO Resp. 57. Rather, Patent Owner argues that “the MPU and FPGAs communicate through shared memory to avoid bus contention of the ‘Local Bus.’” *Id.* at 57–59 (citing Ex. 1007, 2–4, Figs. 2, 11, 17; Ex. 2146 ¶¶ 242–244, 267,

Appendix C

269–270). Patent Owner further argues that Tsutsui did not implement, and teaches away from, the “Typical Implementation Style Using Local Bus” depicted in Figure 11 of Tsutsui because data communication between the MPU and FPGAs occurs frequently, and therefore the local bus “would be blocked by the ‘repetitive data transformations among the MPU, the FPGAs, and the memories.’” *Id.* at 58–59 (quoting Ex. 1007, 2–4). Patent Owner argues that “[t]he FPGAs do not receive any data directly from the Local Bus,” as Tsutsui uses a 2-port SRAM channel “via the I-Cube switching device” between the MPU and FPGAs to avoid problems associated with the local bus architecture. *Id.* (citing Ex. 1007, 2; Ex. 2146 ¶¶ 267, 269–270). Finally, according to Patent Owner, “any data path that includes the I-Cube switch cannot meet the ‘direct’ data connection requirements of Claim 1 because it is unknown whether the data would be stored in the SRAM memory contained in the I-Cube switching device,” and an ambiguous disclosure is insufficient to show anticipation. *Id.* at 60–61.

Petitioner responds that Tsutsui teaches the use of the bus style of Figure 11 in which data is directly communicated between the main memory and FPGAs. Pet. Reply 20 (citing Ex. 1007, 2–3; Ex. 1003 ¶ 225). Further, Petitioner asserts that Patent Owner’s arguments regarding teaching away are not relevant to the issue of anticipation. *Id.* (citing *ClearValue, Inc. v. Pearl River Polymers, Inc.*, 668 F.3d 1340, 1344 (Fed. Cir. 2012)).

We agree with Petitioner. Tsutsui discloses that “YARDS supports three different styles of connection

Appendix C

between FPGAs and [the] MPU: a bus, a direct interrupt, and a 2-port SRAM channel.” Ex. 1007, 2. Although Tsutsui discloses expected “frequent[.]” data communication between the MPU and FPGAs, Tsutsui expressly discloses that YARDS supports a local bus connection between the MPU and FPGAs. *Id.* at 2–3 (“The bus style is the same as the conventional one. . . . Using only the bus architecture, an implementation style of our target system should be similar to Figure 11. The transport data stream is input into the FPGA directly.”), Fig. 1 (depicting the connection between each FPGA and the Local Bus, which Tsutsui states on page 2 are “direct[.]” connections), Fig. 9, (depicting the conventional Local Bus style), Fig. 11 (depicting a direct connection to the FPGA via the Local Bus). We credit Dr. Hauck’s testimony as to how a person of ordinary skill in the art would have understood Tsutsui in that regard, as it is consistent with the figures and language used in the reference. Ex. 1003 ¶¶ 217–224. Further, even with respect to the I-Cube switching device implementation in particular, based on the disclosures of Tsutsui cited above that the system can implement local bus interconnection as shown in Figures 9 and 11, as well as Tsutsui’s statement that the I-Cube switching device “supports various types of connections among its pins such as . . . a bus,” a person of ordinary skill in the art would have understood the I-Cube switching device implementation to use a local bus architecture with direct connection to the FPGA. *Id.* ¶ 225 (quoting Ex. 1007, 2). Given the express disclosure in Tsutsui of direct connection to the FPGA, we agree.

Accordingly, we determine that Tsutsui discloses a local bus connection to the FPGAs that provides data

Appendix C

that the FPGAs then “re-shap[e]” and output (Pet. 35–36 (citing Ex. 1007, 1–3; Ex. 1003 ¶¶ 217–220)), and we are not persuaded by Patent Owner’s argument that Tsutsui fails to disclose that the FPGA alters data “provided directly thereto on said memory module bus.”

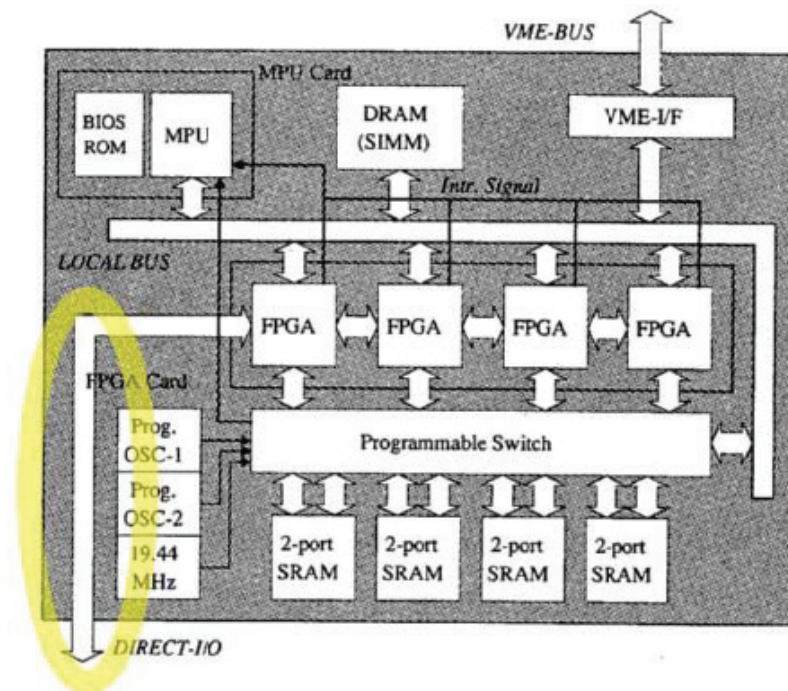
iii. “direct data connection”

Patent Owner further argues that Tsutsui fails to disclose a “direct data connection” between the FPGAs and an external device. PO Resp. 61–63. Specifically, Patent Owner argues that in Tsutsui, “all data going in and coming out of the FPGAs gets stored in either the 2-Port SRAM on the YARDs card, the memory contained in the I-Cube, or both.” *Id.* at 61. Patent Owner further asserts that Petitioner has failed to show support for Tsutsui’s VME-BUS I/F acting as a direct data connection to external devices, and that Tsutsui instead uses the VME bus for controlling and monitoring the system. *Id.* at 61–62 (citing Ex. 2146 ¶¶ 277–278; Ex. 1003 ¶¶ 281–282).

Petitioner responds that Tsutsui discloses that the “Direct I/O connection is connected directly to the FPGAs, *i.e.*, without intervening memory,” as the claimed direct data connection. Pet. Reply 21–22 (citing Pet. 40–43; Ex. 1007, Figs. 1, 13) (emphasis omitted). Petitioner argues that Patent Owner’s expert, Dr. Guccione, explains that the Direct I/O connection of Tsutsui “is a fast, dedicated set of signals on the FPGA that connect directly to the data (network) interface.” *Id.* (quoting Ex. 2146 ¶ 274) (emphasis omitted).

Appendix C

We agree with Petitioner. As explained above, we construe “direct data connection” to mean that the field programmable gate array receives and operates on data directly from the memory module bus, producing altered data and providing that altered data directly to the external device via a direct data connection. *See supra* Section II.B.2. Tsutsui discloses the Direct I/O provides a connection to the FPGAs without intervening memory. Ex. 1007, Figs. 1, 13; *see Pet. Reply* 21–22 (citing Pet. 40). Petitioner provides the following annotated version of Figure 1 of Tsutsui (Pet. 40):



Appendix C

The annotated Figure 1 depicts the “DIRECT-I/O” connection to a FPGA, which Tsutsui describes as “enabl[ing] a direct exchange of signals between other devices and the FPGA card.” Ex. 1007, 2; *see id.*, Figs. 1, 13 (both showing the “DIRECT-I/O” connection); Pet. 40–43. Tsutsui further states that “FPGAs have many direct I/O ports and are suited to handle real-time operations for continuous data streams.” *Id.* at 3. Dr. Guccione also agrees that the “DIRECT-I/O” connection “is a fast, dedicated set of signals on the FPGAs that connect *directly* to the data (network) interface.” Ex. 2146 ¶ 274 (emphasis added).

Accordingly, we determine that Tsutsui’s description of the direct I/O channel discloses the claimed “direct data connection.”

iv. “control connection”

Patent Owner further argues Tsutsui fails to disclose a “control connection . . . for indicating to a processor of said computer system,” as recited in claim 2. PO Resp. 63 (citing Ex. 2146 ¶¶ 282–283). Patent Owner argues that a “control connection implies some sophisticated method of communicating status information back to the processor, and more critically, a way to receive information back, for the purposes of control.” *Id.* Patent Owner argues that the ’524 patent specification discloses the use of a PCI bus that permits bi-directional data flow. *Id.*

We are not persuaded by Patent Owner’s argument. Claim 2 recites “a control connection coupled to said

Appendix C

processor element for indicating to a processor of said computer system an arrival of data on said data connection from said external device.” As discussed above, we adopt Patent Owner’s proposed construction of “control connection” to mean “a bus or data communication channel separate from the memory module bus.” *See supra* Section II.C.3. As argued by Patent Owner, a person with ordinary skill in the art would have understood that the “memory module bus” and “control connection” are separate communication paths. *See id.*

Claim 2 only requires indicating to the processor the *arrival* of data from an external device. Petitioner contends that the YARDS processor element includes interrupt signals that enable the FPGA to control the behavior of the processor element. Pet. 43 (citing Ex. 1007, 2–3; Ex. 1003 ¶ 308). That is, the interrupt signal indicates the *arrival* of data to the processor element. We see no reason to limit the scope of claim 2 to require a PCI Bus or a bi-directional flow of data. As such, we are not persuaded by Patent Owner’s argument.

v. “plurality of processors”

Patent Owner further argues that Tsutsui fails to disclose a “plurality of processors,” as recited in claim 14. PO Resp. 64 (citing Ex. 2146 ¶ 284). Patent Owner argues that claim 14 requires the main system host processor to be a plurality of instruction-based processors, and Tsutsui’s YARDS card only discloses a single MPU microprocessor. *Id.* (citing Ex. 2146 ¶ 274–280, 285).

Appendix C

Petitioner responds that Patent Owner’s argument is based on an erroneous claim construction. Pet. Reply 24. We agree. As noted above, we construe “a plurality of processors” to mean “more than one processor,” and “processor” to mean “any processor in the computer system that is not the processor element of claim 1.” *See supra* Section II.C.4. As such, we are not persuaded by Patent Owner’s argument that claim 14 requires a main system host processor to be a plurality of instruction-based processors. PO Resp. 64. Tsutsui discloses a plurality of processors by disclosing a processing element that includes multiple FPGAs that process data. Ex. 1007, 3; Pet. 48–49 (citing Ex. 1003 ¶¶ 351–352, 353–354). Accordingly, we are not persuaded by Patent Owner’s argument.

3. Conclusion

We are persuaded by Petitioner’s arguments, as they are supported by the cited evidence, including the testimony of Dr. Hauck, which we credit, explaining how a person of ordinary skill in the art would have understood the disclosure of Tsutsui, and notwithstanding Patent Owner’s arguments addressed above. We determine that Petitioner has demonstrated by a preponderance of the evidence that claims 1, 2, and 13–15 of the ’524 patent are anticipated by Tsutsui.

E. Obviousness of Claims 1, 2, and 13–15 over Tsutsui and Tsutsui II

Petitioner contends that claims 1, 2, and 13–15 of the ’524 patent are unpatentable under 35 U.S.C. § 103 as

Appendix C

obvious over Tsutsui and Tsutsui II.¹⁵ Pet. 65–69. For the reasons discussed below, we determine that Petitioner has demonstrated by a preponderance of the evidence that claims 1, 2, and 13–15 of the '524 patent are unpatentable under 35 U.S.C. § 103 as obvious over the combination of Tsutsui and Tsutsui II.

1. Tsutsui II (Ex. 1009)

Tsutsui II is cited by Tsutsui as a reference, and shares two of the same authors as Tsutsui. Ex. 1007, 4. Tsutsui II discloses a FPGA for high-speed digital telecommunication systems. Ex. 1009, 486. Tsutsui II states that the FPGA can realize high-speed transport data processing and its inter-chip connection mechanism enables flexible multi-FPGA modules. *Id.* Tsutsui II further discloses the use of the FPGA in a transport processing system that executes operations due to extraction of data structures from a bit-serial main data stream. *Id.* The data structures consist of two types of data: one being for transportation control (i.e., the “Header”) and the other for the transport of data (i.e., the “Payload”). *Id.* Tsutsui II discloses that transport processing mainly treats the Header part to determine the operation appropriate for a corresponding data block, although other essential operations for transport data, such as bit-error detection/correction and scrambling, are included. *Id.* at 486–87.

15. Petitioner argues, and we agree, that Tsutsui II is a prior art printed publication under 35 U.S.C. §§ 102(a) and (b). *See* Pet. 29.

*Appendix C**2. Analysis**a. Petitioner's Contentions*

As discussed above, Petitioner contends that Tsutsui discloses the limitations of claim 1. *See supra* Section II.D.2.a. Petitioner further argues that Tsutsui II discloses an FPGA that performs data processing functions, including scrambling data. Pet. 65 (citing Ex. 1009, 486). As such, Petitioner argues that to the extent the FPGA of Tsutsui is not “operative to alter data,” as recited in claim 1, the combination of Tsutsui and Tsutsui II discloses “a field programmable gate array configurable to perform an identified algorithm on an operand provided thereto and operative to alter data provided directly thereto on said memory module bus.” *Id.* at 65–66. Specifically, according to Petitioner, it would have been obvious to include Tsutsui II’s data scrambling in the system of Tsutsui, such that the Direct-I/O connection in Tsutsui (“direct data connection”) would provide scrambled payload data (“altered data”) to the ANT card (“external device”) connected to the Direct-I/O connection. *Id.* at 65–66, 69. That is, the only modification being made to Tsutsui is the alteration of data by the FPGA (i.e., scrambling data rather than re-shaping data).

Petitioner argues that a person with ordinary skill in the art would have been motivated to combine Tsutsui and Tsutsui II because Tsutsui cites to Tsutsui II as disclosing “an original FPGA especially designed for high- speed telecommunication data processing.” *Id.* at 67 (quoting Ex. 1007, 1). Petitioner further argues that a

Appendix C

person with ordinary skill in the art would have looked to Tsutsui II to combine its scrambling processing with the ATM protocol system of Tsutsui. *Id.* (citing Ex. 1009, 486; Ex. 1003 ¶ 240). Petitioner further contends that Tsutsui and Tsutsui II are analogous art to the '524 patent and the combination of Tsutsui and Tsutsui II “would have been only the arrangement of old elements . . . with each performing the function it was known to perform and yielding no more than what one would expect from such an arrangement (low level transport processing on a multi-FPGA system).” *Id.* (citing Ex. 1003 ¶ 238).

Notwithstanding Patent Owner’s arguments and alleged objective indicia of nonobviousness, which we address below, we are persuaded by Petitioner that the combination of Tsutsui and Tsutsui II discloses the limitations of claim 1, and a person with ordinary skill in the art would have combined Tsutsui with Tsutsui II. Pet. 65–69. As discussed above, we find that Tsutsui discloses the limitations of claim 1. Tsutsui II discloses features for high-speed transport data processing, including “bit-error detection/correction and scrambling.” Ex. 1009, 486. As such, we agree with Petitioner that the combination of Tsutsui and Tsutsui II discloses “a field programmable gate array configurable to perform an identified algorithm on an operand provided thereto and operative to alter data provided directly thereto on said memory module bus.” We specifically agree with Petitioner that a person with ordinary skill in the art would have looked to improve synchronization, and would have determined that application of a scrambling technique, as taught by Tsutsui II, would have improved synchronization. *See* Pet. 68–69.

Appendix C

For the reasons discussed above, Petitioner’s arguments are supported by the record, including the testimony of Dr. Hauck, and we find them persuasive. Accordingly, we are persuaded that a person with ordinary skill in the art would have been motivated to combine Tsutsui and Tsutsui II. *See* Pet. 66–69; Ex. 1003 ¶¶ 233–244.

Petitioner provides an analysis for dependent claims 2 and 13–15, explaining exactly how the combination of Tsutsui and Tsutsui II discloses the limitations of the claims, and articulates a rationale to combine Tsutsui and Tsutsui II. Pet. 43–53, 65–69. We find that Petitioner supports its arguments with credible evidence, and we are persuaded by Petitioner that Tsutsui and Tsutsui II disclose the limitations of claims 2 and 13–15 and there would have been reason for a person with ordinary skill in the art to combine Tsutsui and Tsutsui II. *See id.*

b. Patent Owner’s Arguments

Patent Owner argues that the combination of Tsutsui and Tsutsui II fails to disclose certain elements of independent claim 1 for the same reasons asserted in opposing Petitioner’s anticipation challenge of claims 1, 2, and 13–15. *See* PO Resp. 67; *supra* Section II.D.2.b. We do not repeat those arguments here. We are not persuaded by Patent Owner’s arguments for the same reasons discussed above. *See id.*

Patent Owner does not separately argue that Tsutsui II in combination with Tsutsui fails to disclose “a field programmable gate array configurable to perform an

Appendix C

identified algorithm on an operand provided thereto and operative to alter data provided directly thereto on said memory module bus.” *See generally* PO Resp. 67–69; PO Sur-Reply 25–26. Rather, Patent Owner generally argues that a person with ordinary skill in the art “would not [have been] motivated to make the proposed combination of Tsutsui with the other prior art.” PO Resp. 71. Patent Owner further generally argues that “Petitioner and its expert resort to both hindsight reasoning and *ipse dixit* errors that are generally fatal to [*inter partes* review] petitions.” *Id.* at 72–73. Petitioner argues that Patent Owner does not address any of the Petition’s motivation arguments in detail and improperly incorporates material by reference. Pet. Reply 27–29.

We are not persuaded by Patent Owner’s arguments. Patent Owner generally alleges deficiencies in Petitioner’s arguments towards a rationale to combine Tsutsui and Tsutsui II, but fails to set forth persuasive argument and evidence to undermine Petitioner’s showing that a person with ordinary skill in the art would have been motivated to combine Tsutsui and Tsutsui II. That is, we are not persuaded by Patent Owner’s general allegations that a person with ordinary skill in the art would not have been motivated to modify Tsutsui. Rather, Petitioner has set forth multiple persuasive reasons to combine Tsutsui and Tsutsui II (two related references with very similar disclosures), and Patent Owner has not meaningfully contested Petitioner’s arguments.

As discussed above, as argued by Petitioner and supported by Dr. Hauck’s testimony, a person with

Appendix C

ordinary skill in the art would have been motivated to combine Tsutsui and Tsutsui II because Tsutsui cites to Tsutsui II as disclosing “an original FPGA especially designed for high-speed telecommunication data processing.” Pet. 67 (quoting Ex. 1007, 1). Furthermore, Tsutsui provides explicit motivation to look to Tsutsui II’s disclosure, stating the system of Tsutsui II “is useful in implementing lower-layer transport operations.” *Id.* (quoting Ex. 1007, 1). Petitioner persuasively argues that a person with ordinary skill in the art, moreover, would have looked to Tsutsui II to combine its scrambling processing with the ATM protocol system of Tsutsui because Tsutsui II refers to scrambling as an “essential operation[] for transport data.” *Id.* (citing Ex. 1009, 486; Ex. 1003 ¶ 240). Petitioner further contends, and we agree, that Tsutsui and Tsutsui II are analogous art to the ’524 patent and the combination of Tsutsui and Tsutsui II “would have been only the arrangement of old elements (the system of Tsutsui with the FPGA payload scrambling processing of Tsutsui II) with each performing the function it was known to perform and yielding no more than what one would expect from such an arrangement (low level transport processing on a multi-FPGA system).” *Id.* at 67 (citing Ex. 1003 ¶ 238).

Accordingly, we are not persuaded by Patent Owner’s general allegations of deficiencies in Petitioner’s arguments regarding a motivation to combine.

Patent Owner additionally argues, with respect to all of Petitioner’s asserted obviousness combinations, that Dr. Hauck’s testimony is conclusory, grounded in hindsight

Appendix C

bias, fails to disclose the underlying facts or data on which his opinions are based under 37 C.F.R. § 42.65(a), fails to consider whether the asserted combinations were “feasible,” and “assume[s] all the benefits [of the asserted combinations] and ignore[s] any drawbacks of cobbling together [the] prior art as proposed in the Petition.” PO Resp. 71–74; PO Sur-Reply 13. Patent Owner acknowledges Dr. Hauck’s experience with FPGAs, but argues that he “lacks the expertise in memory system design and computer architecture that are necessary in this case.” PO Sur-Reply 13. We do not see how any such lack of experience undermines Dr. Hauck’s testimony regarding how a person of ordinary skill in the art would have understood the FPGA-related disclosures of Tsutsui and Tsutsui II and why a person of ordinary skill in the art would have been motivated to combine their teachings. To the extent Patent Owner argues that we should not afford Dr. Hauck’s declaration any weight, we appropriately weigh Dr. Hauck’s testimony pertaining to Petitioner’s anticipation ground and obviousness ground premised on the combination of Tsutsui and Tsutsui II based on the specific testimony provided to determine whether Petitioner has established, by a preponderance of the evidence, the unpatentability of the challenged claims.

3. Objective Indicia of Nonobviousness

Patent Owner also argues that objective indicia of nonobviousness demonstrate that claims 1, 2, and 13–15 would not have been obvious to a person of ordinary skill in the art. PO Resp. 74–75. Patent Owner lists various alleged types of indicia of nonobviousness,

Appendix C

including long-felt need, teaching away, well-known limitations, commercial success, and unexpected results, and provides citations to various documents. *Id.* Patent Owner's citations to documents appear to indicate that the documents support Patent Owner's listed types of objective indicia of nonobviousness.

Patent Owner, however, fails to provide any argument or analysis demonstrating the nonobviousness of the claims. At best, Patent Owner's citations are an inappropriate incorporation by reference of documents. 37 C.F.R. § 42.6(a)(3). Patent Owner neither argues that there is a presumption of a nexus between the claims and the objective indicia of nonobviousness nor provides any explanation establishing a nexus between any objective evidence of nonobviousness and the challenged claims. *See generally* PO. Resp. 74–75. For example, Patent Owner does not explain how the cited testimony shows skepticism in the industry regarding the claimed processor element. Nor does Patent Owner provide any explanation or analysis demonstrating that any products embody the claimed processor element to establish nexus or show commercial success. Accordingly, we conclude that Patent Owner's evidence purportedly showing long felt need, skepticism in the industry, teaching away, well-known limitations, commercial success, and unexpected results does not weigh in favor of nonobviousness of the challenged claims.

4. Conclusion

For the foregoing reasons, we are persuaded that Petitioner has shown by a preponderance of the evidence

Appendix C

that claims 1, 2, and 13–15 of the '524 patent are unpatentable as obvious over the combination of Tsutsui and Tsutsui II.

F. Additional Grounds

Petitioner also challenges claims 1, 2, and 13–15 as obvious over Tsutsui alone, claim 2 as obvious over Tsutsui and Stone, with or without Tsutsui II, claim 14 as obvious over Tsutsui and Collins, with or without Tsutsui II, and claim 15 as obvious over Tsutsui and Hayashi, with or without Tsutsui II. Pet. 53–65, 69–80. Petitioner has proven by a preponderance of the evidence that these claims are unpatentable on other grounds. *See supra* Sections II.D–E. As such, we need not address Petitioner's alternative grounds of unpatentability as to claims 1, 2, and 13–15.

G. Request for Sanctions

Patent Owner “requests a finding that Petitioner violated its duty of candor in this proceeding, and requests appropriate sanctions.” PO Resp. 35–36; PO Sur-Reply 27–28. Specifically, Patent Owner argues that Petitioner failed to “cite its own expert’s testimony” that contradicted its position. *Id.* Patent Owner argues that the failure to “bring this testimony to the attention of the Board” was a violation of Petitioner’s duty of candor to the Board. *Id.* (citing 37 C.F.R. §§ 42.11, 42.12).

“Relief, other than a petition requesting the institution of a trial, must be requested in the form of a motion.”

Appendix C

37 C.F.R. § 42.20(a). 37 C.F.R. § 42.11(d)(2) permits a party, under certain circumstances, to file a motion for sanctions based on conduct that violates the duty of candor in representations made to the Board. *See* 37 C.F.R. §§ 42.11(a), (c). It requires, however, that such a motion for sanctions “must be made separately from any other motion” and “must be authorized by the Board under § 42.20 prior to filing the motion.” 37 C.F.R. § 42.11(d)(2). It also requires that such a motion be served on opposing counsel “[a]t least 21 days prior to seeking authorization to file a motion for sanctions” and must not be filed or presented to the Board if withdrawn or appropriately corrected within twenty-one days of service of the proposed motion. *Id.*

Here, Patent Owner has requested a finding of unspecified sanctions in Patent Owner’s Sur-Reply. PO Sur-Reply 27–28. Patent Owner never requested or obtained authorization from the Board before filing a motion for sanctions, and did not file a motion for sanctions separate from “any other motion.” Furthermore, Patent Owner does not present any evidence or persuasive argument that it had served Petitioner with a proposed motion for sanctions, thereby allowing Petitioner to appropriately correct its alleged violation. As such, Patent Owner has not followed the proper procedure set forth in 37 C.F.R. § 42.11(d)(2), and, accordingly, Patent Owner’s request for sanctions is dismissed.

*Appendix C**H. Constitutionality of the Proceedings*

Patent Owner “objects to the entirety of these proceedings based on the Federal Circuit’s recent opinion in *Arthrex, Inc. v. Smith & Nephew, Inc.*, 941 F.3d 1320 (Fed. Cir. 2019).” PO Sur-Reply 26. Patent Owner argues that “the current structure of the Board violates the Appointments Clause.” *Id.* (citing *Arthrex*, 941 F.3d at 1335). Patent Owner “requests that this proceeding be dismissed in its entirety on the grounds that the panel lacks the constitutional authority to enter a final decision in this case.” *Id.* at 27.

However, Patent Owner’s constitutional challenge as to this issue— whether the as-constituted panel is constitutional—has been addressed by the Federal Circuit’s decision. *See Arthrex*, 941 F.3d at 1337 (“This as-applied severance . . . cures the constitutional violation.”). Accordingly, we do not consider this issue any further.

III. CONCLUSION

Based on the information presented, we conclude that Petitioner has shown, by a preponderance of the evidence, that claims 1, 2, and 13–15 of the ’524 patent are unpatentable.¹⁶

16. As discussed above, we do not reach Petitioner’s challenges to claims 1, 2, and 13–15 as obvious over Tsutsui alone, claim 2 as obvious over Tsutsui and Stone, with or without Tsutsui II, claim 14 as obvious over Tsutsui and Collins, with or without Tsutsui II, and claim 15 as obvious over Tsutsui and Hayashi, with or without Tsutsui II. *See supra* Section II.F.

Appendix C

In summary:

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpat- entable	Claims Not shown Unpat- entable
1, 2, 13-15	102	Tsutsui	1, 2, 13-15	
1, 2, 13-15	103	Tsutsui		
1, 2, 13-15	103	Tsutsui and Tsutsui II	1, 2, 13-15	
2	103	Tsutsui and Stone, with or without Tsutsui II		
14	103	Tsutsui and Collins, with or without Tsutsui II		
15	103	Tsutsui and Hayashi, with or without Tsutsui II		
Overall Outcome			1, 2, 13-15	

63a

Appendix C

IV. ORDER

After due consideration of the record before us, and for the foregoing reasons, it is:

ORDERED that claims 1, 2, and 13–15 of the '524 patent are held unpatentable;

FURTHER ORDERED that Patent Owner's request for sanctions is *dismissed*;

FURTHER ORDERED that Petitioner's Motion to Exclude (Paper 63) is *denied-in-part* and *dismissed-in-part*;

FURTHER ORDERED that Patent Owner's Motion to Exclude (Paper 64) is *denied-in-part* and *dismissed-in-part*; and

FURTHER ORDERED that because this is a final written decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2018-01604
Patent 7,421,524 B2

For PETITIONER:

Joseph Micallef
Jason Greenhut
Scott Border

64a

Appendix C

SIDLEY AUSTIN LLP
jmicallef@sidley.com
jgreenhut@sidley.com
sborder@sidley.com

For PATENT OWNER:

Alfonso Chan
Joseph DePumpo
SHORE CHAN DEPUMPO LLP
achan@shorechan.com
jdepumpo@shorechan.com

65a

APPENDIX D — JUDGMENT OF THE UNITED STATES PATENT AND TRADEMARK OFFICE, BEFORE THE PATENT TRIAL AND APPEAL BOARD, DATED APRIL 9, 2020

UNITED STATES PATENT
AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL
AND APPEAL BOARD

MICROSOFT CORPORATION,

Petitioner,

v.

FG SRC LLC,¹

Patent Owner.

IPR2018-01601²
Patent 7,225,324 B2

Before KALYAN K. DESHPANDE, JUSTIN T. ARBES,
and CHRISTA P. ZADO, *Administrative Patent Judges.*

ARBES, *Administrative Patent Judge.*

1. Patent Owner filed updated mandatory notice information indicating that DirectStream, LLC (“DirectStream”) assigned the challenged patent to FG SRC LLC. Paper 69, 1. Accordingly, the caption for this proceeding has been changed.

2. Cases IPR2018-01602 and IPR2018-01603 have been consolidated with this proceeding.

*Appendix D***JUDGMENT**

Final Written Decision
 Determining All Challenged Claims Unpatentable
 35 U.S.C. § 318(a)

I. INTRODUCTION**A. Background and Summary**

Petitioner Microsoft Corporation filed three Petitions, collectively requesting *inter partes* review of claims 1–5, 7–9, 15, 17, 18, and 20–24 of U.S. Patent No. 7,225,324 B2 (Ex. 1001, “the ’324 patent”) pursuant to 35 U.S.C. § 311(a), as listed in the following chart.³

Case Number	Challenged Claims	Petition
IPR2018-01601	1, 8, 9, and 20	Paper 1 (“Pet.”)
IPR2018-01602	1, 7, 15, 17, and 24	Paper 1 (“-1602 Pet.”)
IPR2018-01603	1–5, 18, and 21–23	Paper 1 (“-1603 Pet.”)

On April 12, 2019, we instituted an *inter partes* review as to all challenged claims on all grounds of unpatentability asserted in the Petitions, and exercised our authority under 35 U.S.C. § 315(d) to consolidate the three proceedings and conduct the proceedings as one trial.

3. Unless otherwise noted, references herein are to the exhibits filed in Case IPR2018-01601.

Appendix D

Paper 21 (“Decision on Institution” or “Dec. on Inst.”). Patent Owner FG SRC LLC subsequently filed a Patent Owner Response (Paper 36, “PO Resp.”), Petitioner filed a Reply (Paper 49, “Reply”), and Patent Owner filed a Sur-Reply (Paper 59, “Sur-Reply”). Petitioner filed a Motion to Exclude (Paper 60, “Pet. Mot.”) certain evidence submitted by Patent Owner, to which Patent Owner filed an Opposition (Paper 63, “PO Opp.”) and Petitioner filed a Reply (Paper 66, “Pet. Mot. Reply”). Patent Owner filed a Motion to Exclude (Paper 61, “PO Mot.”) certain evidence submitted by Petitioner, to which Petitioner filed an Opposition (Paper 62, “Pet. Opp.”) and Patent Owner filed a Reply (Paper 65, “PO Mot. Reply”). An oral hearing was held on February 4, 2020, and a transcript of the hearing is included in the record (Paper 71, “Tr.”).

We have jurisdiction under 35 U.S.C. § 6. This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a). For the reasons that follow, we determine that Petitioner has shown by a preponderance of the evidence that claims 1–5, 7–9, 15, 17, 18, and 20–24 are unpatentable.

B. Related Matters

The parties indicate that the ’324 patent is the subject of the following district court cases: *SRC Labs, LLC v. Microsoft Corp.*, No. 2:18-cv-00321 (W.D. Wash.), and *SRC Labs, LLC v. Amazon Web Servs., Inc.*, No. 2:18-cv-00317 (W.D. Wash.). *See* Pet. 3–4; Paper 69, 1.

*Appendix D***C. The '324 Patent**

The '324 patent⁴ discloses “multi-adaptive processing systems and techniques for enhancing parallelism and performance of computational functions.” Ex. 1001, col. 1, ll. 38–41. Parallel processing “allows multiple processors to work simultaneously on the same problem to achieve a solution” in less time than it would take a single processor. *Id.* at col. 1, ll. 42–47. “[A]s more and more performance is required, so is more parallelism, resulting in ever larger systems” and associated difficulties, including “facility requirements, power, heat generation and reliability.” *Id.* at col. 1, ll. 51–59. The '324 patent discloses that

if a processor technology could be employed that offers orders of magnitude more parallelism per processor, these systems could be reduced in size by a comparable factor. Such a processor or processing element is possible through the use of a reconfigurable processor. Reconfigurable processors instantiate only the functional units needed to solve a particular application, and as a result, have available space to instantiate as many functional units as may be required to solve the problem up to the total capacity of the integrated circuit chips they employ.


Id. at col. 1, l. 63–col. 2, l. 5. The '324 patent describes a known issue where each processor in a multi-processor

4. U.S. Patent No. 7,620,800 B2 (Ex. 1005), a continuation of the '324 patent, is challenged by Petitioner in Case IPR2018-01605. We enter a Final Written Decision in Case IPR2018-01605 concurrently with this Decision.

Appendix D

system is allocated a portion of a problem called a “cell” and “to solve the total problem, results of one processor are often required by many adjacent cells because their cells interact at the boundary.” *Id.* at col. 2, ll. 25–31. Passing intermediate results around the system to complete the problem requires using “numerous other chips and busses that run at much slower speeds than the microprocessor,” diminishing performance. *Id.* at col. 2, ll. 31–37, col. 5, ll. 18–30, Fig. 1 (depicting a conventional multi-processor arrangement). In an adaptive processor-based system, however, “any boundary data that is shared between . . . functional units need never leave a single integrated circuit chip,” reducing “data moving around the system” and improving performance. *Id.* at col. 2, ll. 38–48.

Figure 2 of the '324 patent is reproduced below.



200 **Fig. 2**

Appendix D

Figure 2 is “a functional block diagram of an adaptive processor 200 communications path for implementing the technique of the present invention.” *Id.* at col. 5, ll. 32–35. Adaptive processor 200 includes adaptive processor chip 202, which is coupled to memory element 206, interconnect 208, and additional adaptive processor chips 210. *Id.* at col. 5, ll. 35–40. Adaptive processor chip 202 includes thousands of functional units (“FU”) 204 interconnected by “reconfigurable routing resources” inside adaptive processor chip 202, allowing functional units 204 to “exchange data at much higher data rates and lower latencies than a standard microprocessor.” *Id.* at col. 5, ll. 41–47.

Figures 4A and 4B of the '324 patent are reproduced below.




Figure 4A depicts conventional sequential processing operation 400 where “nested Loops A (first loop 402) and B (second loop 404) are alternately active on different

Appendix D

phases of the process.” *Id.* at col. 6, ll. 1–5. Because first loop 402 must be completed before beginning second loop 404, “all of the logic that has been instantiated is not being completely utilized.” *Id.* at col. 6, ll. 6–12. Figure 4B depicts “multi-dimensional process 410 in accordance with the technique of the present invention.” *Id.* at col. 6, ll. 13–16. “[M]ulti-dimensional process 410 is effectuated such that multiple dimensions of data are processed by both Loops A (first loop 412) and B (second loop 414) such that the computing system logic is operative on every clock cycle.” *Id.* at col. 6, ll. 16–20. A “dimension” of data can be “multiple vectors of a problem, multiple plans of a problem, multiple time steps in a problem and so forth.” *Id.* at col. 6, ll. 27–30. The ’324 patent discloses that available resources are utilized more effectively in the multi-dimensional process by “hav[ing] an application evaluate a problem in a data flow sense. That is, it will ‘pass’ a subsequent dimension of a given problem through the first loop 412 of logic concurrently with the previous dimension of data being processed through the second loop.” *Id.* at col. 6, ll. 21–27.

The ’324 patent states that the disclosed process can be utilized for a variety of applications. *Id.* at col. 9, ll. 18–29. For example, seismic imaging applications, which “process echo data to produce detailed analysis of subsurface features” for oil and gas exploration, would “particularly benefit from the tight parallelism that can be found in the use of adaptive or reconfigurable processors” because they “use data collected at numerous points and consisting of many repeated parameters” and “the results of the computation on one data point are used

Appendix D

in the computation of the next.” *Id.* at col. 9, ll. 35–44; *see id.* at col. 6, l. 31–col. 7, l. 41, Figs. 5A–5B, 6A–6B (describing a seismic imaging function that can be adapted to utilize the disclosed parallelism, where computational process 610 “loops over the depth slices as indicated by reference number 622 and loops over the shots as indicated by reference number 624”). Also, reservoir simulation applications, which “process fluid flow data in . . . oil and gas subsurface reservoirs to produce extraction models,” would benefit from the disclosed process because they define a three dimensional set of cells for the reservoir, utilize repeated operations on each cell, and “information computed for each cell is then passed to neighboring cells.” *Id.* at col. 10, ll. 2–13; *see id.* at col. 7, l. 42–col. 8, l. 26, Figs. 7A–7D (describing “process 700 for performing a representative systolic wavefront operation in the form of a reservoir simulation function” in which “the computation of fluid flow properties are communicated to neighboring cells 710” without storing data in memory, “a set of cells can reside in an adaptive processor,” and “the pipeline of computation can extend across multiple adaptive processors,” where the process involves nested loops and systolic walls 712 and 714 of computation at different time sets), col. 8, l. 27–col. 9, l. 17, Figs. 8A–8C, 9A–9C (describing two other processes for performing “a representative systolic wavefront operation”). Finally, the disclosed process may be used for genetic pattern matching applications, which “look[] for matches of a particular genetic sequence (or model) to a database of genetic records,” performing repeated operations to “compare[] each character in the model to the characters in [a particular] genetic record.” *Id.* at col. 11, ll. 55–64.

*Appendix D***D. Illustrative Claims**

Claim 1 of the '324 patent is independent. Claims 2–5, 7–9, 15, 17, 18, and 20–24 each depend from claim 1. Claims 1 and 15 recite:

1. A method for data processing in a reconfigurable computing system, the reconfigurable computing system comprising at least one reconfigurable processor, the reconfigurable processor comprising a plurality of functional units, said method comprising:

transforming an algorithm into a calculation that is systolically implemented by said reconfigurable computing system at the at least one reconfigurable processor;

instantiating at least two of said functional units at the at least one reconfigurable processor to perform said calculation wherein only functional units needed to solve the calculation are instantiated and wherein each instantiated functional unit at the at least one reconfigurable processor interconnects with each other instantiated functional unit at the at least one reconfigurable processor based on reconfigurable routing resources within the at least one reconfigurable processor as established at instantiation, and wherein systolically linked lines of code of said calculation are instantiated as clusters of functional units within the at least one reconfigurable processor;

Appendix D

utilizing a first of said instantiated functional units to operate upon a subsequent data dimension of said calculation forming a first computational loop; and

substantially concurrently utilizing a second of said instantiated functional units to operate upon a previous data dimension of said calculation forming a second computational loop wherein said systolic implementation of said calculation enables said first computational loop and said second computational loop execute concurrently and pass computed data seamlessly between said computational loops.

15. The method of claim 1 wherein instantiating includes establishing a stream communication connection between functional units.

E. Evidence

The pending grounds of unpatentability in the instant *inter partes* review are based on the following prior art:

Jean-Luc Gaudiot, "Data-Driven Multicomputers in Digital Signal Processing," *Proceedings of the IEEE, Special Issue on Hardware and Software for Digital Signal Processing*, vol. 75, no. 9, Sept. 1987, pp. 1220–1234 (Ex. 1010, "Gaudiot");

Appendix D

Duncan A. Buell, Jeffrey M. Arnold, & Walter J. Kleinfelder, SPLASH2: FPGAs IN A CUSTOM COMPUTING MACHINE (1996) (Ex. 1007, “Splash2”);

Carl Ebeling *et al.*, “Mapping Applications to the RaPiD Configurable Architecture,” *Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines*, Apr. 16–18, 1997, pp. 106–115 (Ex. 1009, “RaPiD”);

Michael Rencher & Brad L. Hutchings, “Automated Target Recognition on SPLASH 2,” *Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines*, Apr. 16–18, 1997, pp. 192–200 (Ex. 1011, “Chunky SLD”);

Yong-Jin Jeong & Wayne P. Burleson, “VLSI Array Algorithms and Architectures for RSA Modular Multiplication,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 2, June 1997, pp. 211–217 (Ex. 1061, “Jeong”); and

D. Roccatano *et al.*, “Development of a Parallel Molecular Dynamics Code on SIMD Computers: Algorithm for Use of Pair List Criterion,”

Appendix D

Journal of Computational Chemistry, vol. 19, no. 7, May 1998, pp. 685–694 (Ex. 1012, “Roccatano”)⁵

Petitioner filed a declaration from Harold Stone, Ph.D. (Ex. 1003) with its Petition and a reply declaration from Dr. Stone (Ex. 1076) with its Reply. Patent Owner filed declarations from Jon Huppenthal (Ex. 2100), Houman Homayoun, Ph.D. (Exs. 2029, 2111), and Tarek El-Ghazawi, Ph.D. (Ex. 2164).

F. Asserted Grounds

The instant *inter partes* review involves the following grounds of unpatentability:

Petition(s)	Claim(s) Challenged	35 U.S.C. §	Reference(s)/ Basis
IPR2018-01601, IPR2018-01602, IPR2018-01603	1, 15, 18, 21, 22	102(a), 102(b) ⁶	Splash2

5. When citing the prior art references and other exhibits, we refer to the page numbers in the bottom-right corner added by the filing party. *See* 37 C.F.R. § 42.63(d)(2).

6. The Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (“AIA”), amended 35 U.S.C. §§ 102 and 103. Because the challenged claims of the ’324 patent have an effective filing date before the effective date of the applicable AIA amendments, we refer to the pre-AIA versions of 35 U.S.C. §§ 102 and 103.

Appendix D

Petition(s)	Claim(s) Challenged	35 U.S.C. §	Reference(s)/ Basis
IPR2018-01601, IPR2018-01602, IPR2018-01603	1, 15, 18, 21, 22	103(a)	Splash2
IPR2018-01601, IPR2018-01602, IPR2018-01603	1, 15, 18, 21, 22	103(a)	Splash2, Gaudiot
IPR2018-01601	8, 9	103(a)	Splash2, RaPiD
IPR2018-01601	8, 9	103(a)	Splash2, RaPiD, Gaudiot
IPR2018-01601	20	103(a)	Splash2, Jeong
IPR2018-01601	20	103(a)	Splash2, Jeong, Gaudiot
IPR2018-01602	7, 17, 24	103(a)	Splash2, Chunky SLD
IPR2018-01602	7, 17, 24	103(a)	Splash2, Chunky SLD, Gaudiot
IPR2018-01603	2-5, 22, 23	103(a)	Splash2, Roccatano
IPR2018-01603	2-5, 22, 23	103(a)	Splash2, Roccatano, Gaudiot

*Appendix D***II. ANALYSIS****A. Motions to Exclude**

The party moving to exclude evidence bears the burden of proof to establish that it is entitled to the relief requested—namely, that the material sought to be excluded is inadmissible under the Federal Rules of Evidence. *See* 37 C.F.R. §§ 42.20(c), 42.62(a). For the reasons discussed below, Petitioner’s Motion is granted-in-part, denied-in-part, and dismissed-in-part, and Patent Owner’s Motion is denied-in-part and dismissed-in-part.

1. Petitioner’s Motion to Exclude

Exhibits 2100, 2066, 2076, and 2092: Petitioner moves to exclude the entirety of the declaration (Ex. 2100) of Mr. Huppenthal, one of the named inventors of the ’324 patent, “as not being relevant to any issue on which trial has been instituted, and for lacking foundation, containing hearsay, and/or causing undue prejudice.” Pet. Mot. 3–6. Petitioner argues that the declaration includes “irrelevant narrative discussion of [Mr. Huppenthal’s] participation in reconfigurable computing” and statements “either based on hearsay or lack of personal knowledge.” *Id.* at 3–4. Petitioner also moves to exclude paragraphs 80 and 82–86 of the declaration based on Mr. Huppenthal’s alleged “refusal to answer questions concerning those portions of the declaration” during cross-examination. *Id.* at 1–3.

Petitioner also moves to exclude three transcripts (Exs. 2066, 2076, 2092) of depositions of Petitioner’s

Appendix D

declarants from other *inter partes* reviews as “not being relevant to any issue on which trial has been instituted, for containing hearsay, and/or causing undue prejudice.” *Id.* at 6–7. Petitioner argues that allowing the transcripts in the record would be “highly prejudicial as they present themselves with the indicia of expert testimony while being totally devoid from the necessary context of the matters from which they originate.” *Id.* Patent Owner cites Exhibits 2066 and 2076 in its Response, but does not cite Exhibit 2092 in its Response or Sur-Reply.

Petitioner’s Motion is dismissed as moot, as we do not rely on the cited portions of the testimony in a manner adverse to Petitioner in this Decision. As explained below, even if the testimony is considered, we are not persuaded by Patent Owner’s arguments regarding the state of the art or alleged nonobviousness of the challenged claims, and Patent Owner has not shown proof of secondary considerations that would support a conclusion of nonobviousness. *See infra* Sections II.E–II.J.

Exhibit 2111: Petitioner moves to exclude paragraphs 41, 145, 180, 181, and 247 of the declaration of Dr. Hodayoun, which refer to Exhibits 2066 and 2100. Pet. Mot. 8–9. Because we do not exclude those exhibits, we also dismiss as moot Petitioner’s Motion with respect to Exhibit 2111.

Exhibits 2067–2075, 2077, 2079–2099, 2101–2103, 2105, 2106, 2109, 2110, 2112–2133, 2139–2151, 2155, 2161–2163, and 2168: Petitioner moves to exclude a number of exhibits as “not being relevant to any issues on which trial has

Appendix D

been instituted, lacking foundation, and/or causing undue prejudice” because the exhibits were not discussed or cited, or “only cited superficially,” in Patent Owner’s Response and Sur-Reply. Pet. Mot. 7–8. Petitioner’s Motion is dismissed as moot, as we do not rely on the exhibits in a manner adverse to Petitioner in this Decision. We note, however, that in evaluating Petitioner’s asserted grounds of unpatentability, we only consider substantive arguments made by the parties in their papers during trial (i.e., the Petitions, Response, Reply, and Sur-Reply). To the extent a document is filed in the record but never discussed in a paper, there is no substantive argument pertaining to that document that can be considered.

Exhibit 2168: Petitioner moves to exclude Exhibit 2168 under Federal Rules of Evidence 401–403 as cumulative of Exhibit 1007. Pet. Mot. 8. Both exhibits are copies of Splash2. Dr. El-Ghazawi refers to the document in his declaration with the numeral “1007.” Ex. 2164 ¶ 40. To ensure a clear record, we grant Petitioner’s Motion, expunge Exhibit 2168, and refer herein to Exhibit 1007. *See* 37 C.F.R. §§ 42.7(a), 42.6(d) (“A document already in the record of the proceeding must not be filed again, not even as an exhibit or an appendix, without express Board authorization.”).

Portions of Patent Owner’s Response: Petitioner moves to exclude portions of Patent Owner’s Response referring to the exhibits that Petitioner seeks to exclude. Pet. Mot. 9. Patent Owner’s Response is a paper with

Appendix D

attorney arguments, not evidence that may be excluded.⁷ Further, we do not exclude any of the exhibits referred to in the identified portions of the Response. Petitioner’s Motion is denied as to Patent Owner’s Response.

2. Patent Owner’s Motion to Exclude

Exhibits 1074, 1077, and 1079: Patent Owner moves to exclude three technical documents filed by Petitioner with its Reply.

First, Patent Owner moves to exclude Exhibits 1074 and 1079 as unauthenticated under Federal Rule of Evidence 901. PO Mot. 6–7. “To satisfy the requirement of authenticating or identifying an item of evidence, the proponent must produce evidence sufficient to support a finding that the item is what the proponent claims it is.” Fed. R. Evid. 901(a). Certain evidence, though, is “self-authenticating” and “require[s] no extrinsic evidence of authenticity in order to be admitted.” Fed. R. Evid. 902. Exhibit 1074 is an article by Maya Gokhale and Ron Minnich titled “FPGA Computing in a Data Parallel C,” and includes an IEEE trade inscription, copyright symbol, and International Standard Book Number (ISBN) on the first page (“0-8186-3890-7/93 \$03.00 © 1993 IEEE”). It is self-authenticating under at least Federal Rule of Evidence 902(6) (“Printed material purporting to be a . . .

7. Petitioner did not seek authorization to file a motion to strike Patent Owner’s Response. *See* Patent Trial and Appeal Board Consolidated Trial Practice Guide (Nov. 2019), 80–81, *available at* <https://www.uspto.gov/TrialPracticeGuideConsolidated> (“Trial Practice Guide”).

Appendix D

periodical.”) and 902(7) (“An inscription, sign, tag, or label purporting to have been affixed in the course of business and indicating origin, ownership, or control.”). Exhibit 1079 is an excerpt from a book by Kevin Skahill titled “VHDL for Programmable Logic.” It includes a copyright notice (“Copyright © 1996 by Addison-Wesley Publishing, Inc.”), Library of Congress catalogue information, and hand-written library catalogue information. Ex. 1079, Cover 3, v. Petitioner points out where the book is available from “several well-known book sellers.” Pet. Opp. 4 & n.2. It is authenticated at least as an ancient document under Federal Rule of Evidence 901(b)(8) because it “is in a condition that creates no suspicion about its authenticity,” “was in a place where, if authentic, it would likely be,” and “is at least 20 years old when offered.” Nothing about either exhibit suggests that it is not what it points to be, and Patent Owner does not point to anything in particular in the exhibits that would indicate otherwise. *See* PO Mot. 6–7; PO Mot. Reply 1–4.

Second, Patent Owner moves to exclude Exhibits 1074, 1077, and 1079 as containing inadmissible hearsay under Federal Rule of Evidence 802. PO Mot. 7. Patent Owner states that Petitioner in its Reply “cites each of these documents to prove the truth of technical matters allegedly asserted in such documents, *i.e.* to support Petitioner’s specific factual assertions regarding a technical issue.” *Id.* We are not persuaded. Patent Owner does not identify any particular “statement” in any of the exhibits that is being offered “to prove the truth of the matter asserted in the statement,” and thus fails to meet its burden to prove inadmissibility as hearsay. *See* Fed.

Appendix D

R. Evid. 801(c); 37 C.F.R. § 42.20(c). Even if Patent Owner had done so, Petitioner cites the exhibits to show what a person of ordinary skill in the art would have known at the time of the '324 patent and, with respect to Exhibit 1074 in particular, how a person of ordinary skill in the art would have understood another reference relied on by Patent Owner that cites the article. *See* Reply 7, 17, 33; Pet. Opp. 5–6; Ex. 2167, 37–38, 208. The exhibits are not being offered for the truth of any particular matter discussed in the references. Finally, Exhibits 1074 and 1079 also are admissible under Federal Rule of Evidence 803(16), which provides as an exception to the hearsay rule “[a] statement in a document that was prepared before January 1, 1998, and whose authenticity is established.” *See* 1074, 94 (“1993” date); Ex. 1079, Cover 3, v (“1996” date).

Third, Patent Owner moves to exclude Exhibit 1077 as “irrelevant to the patent and claim construction issues in dispute.” PO Mot. 7–8. Exhibit 1077 is an excerpt of certain pages of the Microsoft Computer Dictionary (5th ed. 2002). Patent Owner argues that Exhibit 1077 “is extrinsic evidence pertaining to Petitioner’s proffered definitions from the Microsoft Computer Dictionary of ‘data structure’ and ‘data path,’ neither of which are claim terms in the patent.” PO Mot. 7. Federal Rule of Evidence 401 provides that “[e]vidence is relevant if: (a) it has any tendency to make a fact more or less probable than it would be without the evidence; and (b) the fact is of consequence in determining the action.”

Petitioner submits Exhibit 1077 in support of its arguments regarding the interpretation of disputed claim

Appendix D

language, specifically the term “stream communication.” Reply 33. The meaning of this phrase is “of consequence in determining” whether challenged claim 15 is unpatentable over the asserted prior art, and Exhibit 1077 provides insight as to the meaning of words used in both parties’ proposed interpretations. *See infra* Section II.C.4. Exhibit 1077 has some “tendency to make a fact more or less probable than it would be without the evidence” and is relevant under Federal Rule of Evidence 401. Thus, there is no basis to exclude Exhibit 1077. We also note that Patent Owner appears to have filed a full copy of the dictionary as Exhibit 2065, and the same pages filed as Exhibit 1077 (pages 144–145) that Patent Owner seeks to exclude are in the exhibit that Patent Owner filed. Further, the record contains numerous other dictionary references filed by both parties, including Exhibits 1025, 1059, 2024–2026, and 2038. Patent Owner’s Motion is denied as to Exhibits 1074, 1077, and 1079.

Exhibit 1076: Patent Owner also moves to exclude paragraphs 15–17 of the reply declaration of Dr. Stone, which refer to Exhibit 1074 “for the first time,” under Federal Rules of Evidence 402 and 403. PO Mot. 8. Because we find no basis to exclude Exhibit 1074, we also deny Patent Owner’s Motion with respect to Exhibit 1076. To the extent Patent Owner’s position is that Petitioner’s Reply and Dr. Stone’s reply declaration exceed the proper scope of a reply, we address those arguments below. *See infra* Section II.E.2.a.2; Trial Practice Guide, 79 (“A motion to exclude is not a vehicle for addressing the weight to be given evidence—arguments regarding weight should appear only in the merits documents. Nor should a motion

Appendix D

to exclude address arguments or evidence that a party believes exceeds the proper scope of reply or sur-reply.”).

Exhibits 1075 and 1078: Patent Owner moves to exclude certain portions of the transcript of the deposition of Dr. Homayoun because the questions asked were “vague, ambiguous, call[] for a legal conclusion, and misleading.” PO Mot. 9–11 (citing Ex. 1075, 27:4–12, 65:5–17). Patent Owner also moves to exclude a portion of the transcript of the deposition of Dr. El-Ghazawi because the question asked was “vague, ambiguous, and calls for a speculative answer.” *Id.* at 11 (citing Ex. 1078, 65:12–17). Patent Owner’s Motion is dismissed as moot, as we do not rely on the disputed portions of the testimony in rendering our Decision.

B. Level of Ordinary Skill in the Art

In determining the level of ordinary skill in the art for a challenged patent, we look to “1) the types of problems encountered in the art; 2) the prior art solutions to those problems; 3) the rapidity with which innovations are made; 4) the sophistication of the technology; and 5) the educational level of active workers in the field.” *Ruiz v. A.B. Chance Co.*, 234 F.3d 654, 666–667 (Fed. Cir. 2000). “Not all such factors may be present in every case, and one or more of them may predominate.” *Id.*

Petitioner’s declarant, Dr. Stone, testifies that a person of ordinary skill in the art at the time of the ’324 patent would have had “an advanced degree in electrical or computer engineering, or computer science with substantial

Appendix D

study in computer architecture, hardware design, and computer algorithms,” and “at least three years’ experience working in the field,” or alternatively “a bachelor’s degree covering those disciplines and at least four years working [in] the field.” Ex. 1003 ¶ 45. According to Dr. Stone,

[s]uch a person would also have been knowledgeable about the programming, design and operation of computer systems based on reconfigurable components such as FPGAs (field programmable gate arrays) and CPLDs (complex programmable logic devices), including computer systems for performing systolic and data driven calculations. That person would also have been familiar with hardware description languages such as [Very High Speed Integrated Circuit Hardware Description Language (VHDL)] that could be used to configure FPGAs and CPLDS that serve as components of reconfigurable computer systems. Finally, as demonstrated by many of the references discussed [in Dr. Stone’s declaration], such a person would also have been familiar with various other areas of technology that by 2002 had relied on high performance and parallel computing systems, such as genetic sequence comparisons, image processing, data mining, and processing related to proteins and organic structures.

Id.

Patent Owner states that it “does not dispute the level of education and skill promoted by [Dr. Stone],” and

Appendix D

Patent Owner's declarant, Dr. Homayoun, "agree[d] with Dr. Stone's assessment of the level of ordinary skill in the art," noting that such an individual also would have been "experienced in developing with high-level languages (C and Fortran), hardware description languages, and the unique problems involved with programming FPGAs and FPGA based systems." *See* PO Resp. 29; Ex. 2029 ¶ 17; *see also* Ex. 2111 ¶ 133 ("In general, I would agree to the level of education and skill promoted by [Petitioner's] expert [for the '324 patent]."). Dr. Homayoun further expands on his understanding of what a person of ordinary skill in the art at the time of the '324 patent would have known and considered. Ex. 2111 ¶¶ 133–146. For example, Dr. Homayoun testifies that in addition to the technical background set forth in Dr. Stone's definition, a person of ordinary skill in the art would have "considered all of the state of the art [described in Dr. Homayoun's declaration] in the design of computer architecture, . . . [i]ncluding the issues of reconfigurable programming, processor speed, FPGA speed, and cost/benefit analysis of overhead introduction as applied to [high performance computing (HPC)] applications." *Id.* ¶ 133. Also, according to Dr. Homayoun, an ordinarily skilled artisan would have "consider[ed] the technical problems [the '324 patent was] attempting to solve" without using the patent "as a technical road-map to then conflate the technical problem with the solution," and "would have considered the drawbacks in HPC computing and the deficiencies in FPGA systems and computer architecture design considerations at the time of the disclosed inventions." *Id.* ¶ 140. Patent Owner similarly argues in its Response that a person of ordinary skill in the art would have "considered" all of these issues. PO Resp. 17–29.

Appendix D

We have evaluated all of Patent Owner’s arguments and supporting evidence regarding what a person of ordinary skill in the art allegedly would have considered when reading the asserted references. Patent Owner’s arguments pertain more to its criticism of Dr. Stone’s analysis as allegedly failing to understand the problems solved by the ’324 patent and being based on “hindsight bias” than a dispute over the “level” of ordinary skill in the art. *See id.* (also arguing that “Petitioner essentially uses the patent itself as a roadmap for stitching together various prior art references”); Sur-Reply 1–8; Tr. 64:4–65:15. It suffices at this point to conclude that a person of ordinary skill in the art would have had the *technical education and work experience* set forth in Dr. Stone’s declaration (and agreed to by Dr. Homayoun). *See* Tr. 64:13–14 (Patent Owner stating that there is no dispute as to “the level of education”). Among other things, such an individual would have had *knowledge* of “multi-adaptive processing systems and techniques,” “parallel processing,” and field-programmable gate arrays (FPGAs). *See* Ex. 1001, col. 1, ll. 35–59; Ex. 1003 ¶ 45; Ex. 2029 ¶ 17; Ex. 2111 ¶¶ 107–119, 133; Ex. 1007, 1–9; *Custom Accessories, Inc. v. Jeffrey-Allan Indus., Inc.*, 807 F.2d 955, 962 (Fed. Cir. 1986) (“The person of ordinary skill is a hypothetical person who is presumed to be aware of all the pertinent prior art.”). What that individual would have considered in evaluating particular prior art references and making potential combinations, though, is an issue we address below in evaluating Petitioner’s grounds of unpatentability.⁸

8. We note that for Petitioner’s anticipation ground based on Splash2, the level of ordinary skill in the art is relevant to understanding how a person of ordinary skill in the art would have

Appendix D

Based on the full record developed during trial, including our review of the '324 patent and the types of problems and prior art solutions described in the '324 patent, as well as the sophistication of the technology described in the '324 patent, we conclude that a person of ordinary skill in the art would have had (1) an advanced degree in electrical or computer engineering, or computer science with substantial study in computer architecture, hardware design, and computer algorithms, and at least three years of experience working in the field, or (2) a bachelor's degree covering those disciplines and at least four years working in the field. We apply that level of skill for purposes of this Decision.

C. Claim Interpretation

In this proceeding, we interpret the claims of the unexpired '324 patent using the “broadest reasonable construction in light of the specification of the patent.”

interpreted the claims and understood the reference, but not with respect to, for example, any motivations or problems a person of ordinary skill in the art would have had in combining teachings. *See Wasica Finance GmbH v. Continental Auto. Sys., Inc.*, 853 F.3d 1272, 1284 (Fed. Cir. 2017) (“Anticipation is an inquiry viewed from the perspective of one skilled in the art.”); *Vivint, Inc. v. Alarm.com Inc.*, 741 F. App'x 786, 791–792 (Fed. Cir. 2018) (unpublished) (“While anticipation is proven based on the express and inherent teachings of a single prior art reference, an obviousness analysis reaches beyond the prior art reference and takes into account other considerations such as the level of ordinary skill in the art and any objective indicia of nonobviousness.”).

Appendix D

37 C.F.R. § 42.100(b) (2017).⁹ Under this standard, we interpret claim terms using “the broadest reasonable meaning of the words in their ordinary usage as they would be understood by one of ordinary skill in the art, taking into account whatever enlightenment by way of definitions or otherwise that may be afforded by the written description contained in the applicant’s specification.” *In re Morris*, 127 F.3d 1048, 1054 (Fed. Cir. 1997); see *In re Smith Int’l, Inc.*, 871 F.3d 1375, 1382–83 (Fed. Cir. 2017) (“[The] broadest reasonable interpretation . . . is an interpretation that corresponds with what and how the inventor describes his invention in the specification.”). “Under a broadest reasonable interpretation, words of the claim must be given their plain meaning, unless such meaning is inconsistent with the specification and prosecution history.” *TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1062 (Fed. Cir. 2016). Our interpretation “‘cannot be divorced from the specification and the record evidence,’ and ‘must be consistent with the one that those skilled in the art would reach.’ A

9. The Petitions in this proceeding were filed on September 5, 2018, prior to the effective date of the rule change that replaces the broadest reasonable interpretation standard with the federal court claim interpretation standard. See *Changes to the Claim Construction Standard for Interpreting Claims in Trial Proceedings Before the Patent Trial and Appeal Board*, 83 Fed. Reg. 51,340, 51,340, 51,358 (Oct. 11, 2018) (amending 37 C.F.R. § 42.100(b) effective November 13, 2018) (now codified at 37 C.F.R. § 42.100(b) (2019)); *Game & Tech. Co. v. Wargaming Grp. Ltd.*, 942 F.3d 1343, 1351 (Fed. Cir. 2019) (“If, as here, the [*inter partes* review] stems from a petition filed before November 13, 2018, the claims are given the ‘broadest reasonable interpretation’ consistent with the specification.” (citation omitted)).

Appendix D

construction that is ‘unreasonably broad’ and which does not ‘reasonably reflect the plain language and disclosure’ will not pass muster.” *Microsoft Corp. v. Proxyconn, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015) (citations omitted), *overruled on other grounds by Aqua Prods., Inc. v. Matal*, 872 F.3d 1290 (Fed. Cir. 2017).

Upon review of the parties’ arguments during trial and the evidence as a whole, we conclude that three terms in claim 1 (“systolic,” “computational loop,” and “pass computed data seamlessly between said computational loops”) and one term in claim 15 (“establishing a stream communication connection between functional units”) require interpretation.

1. “Systolic”

Claim 1 recites “transforming an algorithm into a calculation that is *systolically* implemented by said reconfigurable computing system at the at least one reconfigurable processor,” wherein “*systolically* linked lines of code of said calculation are instantiated as clusters of functional units within the at least one reconfigurable processor” and “said *systolic* implementation of said calculation enables said first computational loop and said second computational loop execute concurrently and pass computed data seamlessly between said computational loops” (emphases added).

In the Decision on Institution, based on the parties’ arguments and record at the time, we preliminarily interpreted the term “systolic” recited in claim 1 to mean

Appendix D

“the characteristic of rhythmically computing and passing data directly between processing elements in a manner that is transport triggered (i.e., by the arrival of a data object) rather than a program counter or clock driving movement of the data.” Dec. on Inst. 17–21. Petitioner agrees with that interpretation. Reply 26. Patent Owner argues that “systolic” has a “plain and ordinary meaning and need not be construed,” but if the term is interpreted, it means “[a]n array of many interconnected functional units that operates in a data flow sense and allows different data to flow in different directions.” PO Resp. 42, 50.

Based on our review of the full trial record, we maintain our preliminary conclusion in the Decision on Institution that the applicants described the plain meaning of “systolic” during prosecution of the ’324 patent. *See* Dec. on Inst. 17–21. Specifically, following an Office Action rejecting claim 1 under 35 U.S.C. § 112, first paragraph, and 35 U.S.C. § 103, the applicants conducted an interview with the examiner and proposed, among other amendments, adding the “systolic” language above to the claim. Ex. 1002, 186–212. The applicants then filed an Office Action response amending the claim, noting that the examiner had requested during the interview that the applicants “further define the term[s] instantiated and systolic” and arguing the following in response to the 35 U.S.C. § 112, first paragraph, written description rejection:

[T]he term systolic computation is derived from continual and pulsating pumping of the human heart. In computer architecture

Appendix D

a systolic array is an arrangement of data processing units similar to a central processing unit but without a program counter or clock that drives the movement of data. That is because the operation of the systolic array is transport triggered, i.e. by the arrival of a data object. Data flows across the array between functional units, usually with different data flowing in different directions. David J. Evans in his work, Systolic algorithms. Systolic algorithms, number 3 in Topics in Computer Mathematics. Gordon and Breach, 1991 define a Systolic system as a “network of processors which rhythmically compute an[d] pass data through the system[.]” Thus in the Applicant’s invention Systolic implementation will connect computational loops such that data from one compute loop will be passed as input data to a concurrently executing compute loop. In the Applicant’s invention data computed by computation units or groups of functional units flows seamlessly and concurrently with data being computed by other groups of functional units. Thus, the process claimed by the Applicant therefore significantly increases the computing processes taking place in a reconfigurable processor.

Ex. 1002, 224–226 (emphases added); *see also id.* at 119 (stating that the term “systolic,” coined by H.T. Kung of Carnegie-Mellon, refers to the rhythmic transfer of data through the pipeline, like blood flowing through

Appendix D

the vascular system”). Importantly, after the applicants filed the Office Action response amending claim 1 to include the “systolic” language and making the above argument, the examiner allowed the claims. *Id.* at 234–238. During prosecution, “an applicant’s amendment accompanied by explanatory remarks can define a claim term by demonstrating what the applicant meant by the amendment.” *Personalized Media Comm’cns, LLC v. Apple Inc.*, 952 F.3d 1336, 1340 (Fed. Cir. 2020). Patent Owner acknowledges that the prosecution history argument quoted above was “an explanation of the plain and ordinary meaning of the term ‘systolic.’” PO Resp. 47.

Our interpretation also is consistent with the Specification, in particular the description of Figures 8A and 8B showing an exemplary “systolic” operation where “data is continually passed directly from one computational loop to another for processing without intervening structures between the loops” and without a program counter or clock driving movement of the data. Ex. 1003 ¶ 70; *see* Ex. 1001, col. 8, ll. 27–45 (disclosing a “representative systolic wavefront operation” in which “systolic processing in the process 800 can pass previously computed data down within a column” and “to subsequent columns as well,” with the computational loops operating concurrently); Ex. 1002, 227 (applicants citing Figures 8A and 8B and the accompanying text as a “a systolic wave front operation” during prosecution of the ’324 patent).

Other documentation from the time also shows that the applicants’ description of the term “systolic” is how a person of ordinary skill in the art would have understood

Appendix D

the term. H.T. Kung, *Why Systolic Architectures?*, IEEE Jan. 1982 (Ex. 1016, “Kung”), discloses:

A systolic system consists of a set of interconnected cells, each capable of performing some simple operation. Because simple, regular communication and control structures have substantial advantages over complicated ones in design and implementation, cells in a systolic system are typically interconnected to form a systolic array or a systolic tree. *Information in a systolic system flows between cells in a pipelined fashion, and communication with the outside world occurs only at the “boundary cells.”* For example, in a systolic array, only those cells on the array boundaries may be I/O ports for the system.

...

The basic principle of a systolic architecture, a systolic array in particular, is illustrated in Figure 1. By replacing a single processing element with an array of [processing elements], or cells in the terminology of this article, a higher computation throughput can be achieved without increasing memory bandwidth. *The function of the memory in the diagram is analogous to that of the heart; it “pulses” data (instead of blood) through the array of cells. The crux of this approach is to ensure that once a data item is brought out from the memory*

Appendix D

it can be used effectively at each cell it passes while being “pumped” from cell to cell along the array. This is possible for a wide class of compute-bound computations where multiple operations are performed on each data item in a repetitive manner.

Id. at 39 (emphases added).

Figure 1 of Kung is reproduced below.

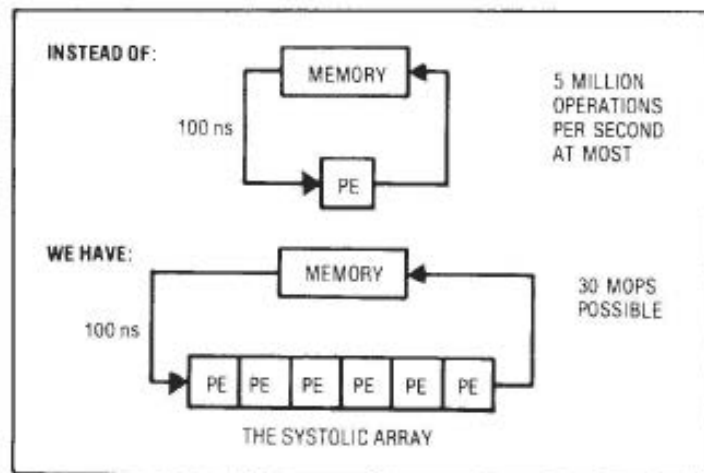


Figure 1. Basic principle of a systolic system.

Figure 1 depicts how data travels and interacts with memory using a systolic array of processing elements versus a single processing element. *Id.* at 38. Another reference by the same author states:

A systolic system is a network of processors which rhythmically compute and pass data

Appendix D

through the system. . . . In a systolic computing system, the function of a processor is analogous to that of the heart. Every processor regularly pumps data in and out, each time performing some short computation, so that a regular flow of data is kept up in the network.

Ex. 1015, 5. Including in the interpretation “rhythmically computing and passing data directly between processing elements” and operating in a “transport triggered” manner is consistent with Kung’s description of each processing element processing data and “puls[ing]” or “pump[ing]” it to the next processing element in the array. *See* Ex. 1016, 39; Ex. 1003 ¶¶ 76–78.

Patent Owner in its Response relies on the same portions of the prosecution history and Kung discussed above. *See* PO Resp. 42–50. That evidence, however, supports our preliminary interpretation, not Patent Owner’s proposed interpretation, for the reasons explained above. Patent Owner also argues that “systolic means an array of interconnected processing elements that only interact with memory at the array boundaries so that the data is processed by multiple processing elements before returning to memory.” PO Resp. 44–48 (citing Ex. 1001, col. 2, ll. 38–42, Fig. 2; Ex. 2046 ¶ 16; Ex. 2111 ¶¶ 125–131, 177).¹⁰ We do not see how that undermines the preliminary

10. Dr. Homayoun does not provide an opinion as to the proper interpretation of “systolic,” but testifies that “[a] person or ordinary skill in the art would *not* consider the terms systolic, data driven, and seamless to have distinct meanings.” Ex. 2111 ¶ 177 (emphasis added). We understand this to be a typographical

Appendix D

interpretation, however. Logically, if processing elements interact with memory *only* at array boundaries, there can be no memory receiving data between the processing elements; data would be passed “directly” between them. *See* Reply 29. We further address the parties’ arguments regarding the use of the word “directly” in interpreting “pass computed data seamlessly between said computational loops” below. *See infra* Section II.C.3.

Applying the broadest reasonable interpretation of the claims in light of the Specification, we interpret “systolic” to mean “the characteristic of rhythmically computing and passing data directly between processing elements in a manner that is transport triggered (i.e., by the arrival of a data object) rather than a program counter or clock driving movement of the data.”

2. “Computational Loop”

Claim 1 recites “utilizing a first of said instantiated functional units to operate upon a subsequent data dimension of said calculation forming a *first computational*

error, given the preceding heading stating the opposite and Patent Owner’s arguments in its Response. *See id.*; PO Resp. 34. We agree that “systolic” and “pass computed data seamlessly between said computational loops” in claim 1 have different meanings. Although both of our interpretations use the word “directly,” we interpret the terms to mean different things. *See infra* Sections II.C.1, II.C.3. Nevertheless, the terms plainly are related, as claim 1 recites that the “systolic” implementation of the calculation is what “enables” the first and second computational loops to “pass computed data seamlessly between said computational loops.”

Appendix D

loop” and “substantially concurrently utilizing a second of said instantiated functional units to operate upon a previous data dimension of said calculation forming a *second computational loop*” (emphases added).

In the Decision on Institution, we preliminarily interpreted “computational loop” in claim 1 to mean “a set of computations that is executed repeatedly, either a fixed number of times or until some condition is true or false.” Dec. on Inst. 21–23. Petitioner agrees with that interpretation. Reply 35. Patent Owner states that it “does not disagree with” the preliminary interpretation, but “proposes a small clarification to reflect the plain and ordinary meaning of a loop within the field of high-performance computing and in particular in the context of the ’324 Patent.” PO Resp. 70. Specifically, Patent Owner proposes the following interpretation: “a set of computations that is executed repeatedly *per datum*, either a fixed number of times or until some condition is true or false.”¹¹ *Id.* at 69–70 (emphasis added).

Patent Owner in its Response cites three technical dictionary definitions of “loop,” upon which we had based the preliminary interpretation as reflective of the plain meaning of the term. *See id.* at 70–71; Dec. on Inst. 21–23. None of those definitions, however, requires that a set of instructions or computations be executed repeatedly “per datum.” *See* Ex. 2024, 4 (“[a] sequence of

11. Patent Owner’s proposed interpretation in its Preliminary Response—“a sequence of computations that is repeated until a prescribed condition is satisfied”—did not include a “per datum” requirement. Paper 15, 21–22.

Appendix D

instructions that is repeated until a prescribed condition, such as agreement with a data element or completion of a count, is satisfied”); Ex. 2025, 5 (“a series of instructions being carried out repeatedly until a terminal condition prevails” or “[a] sequence of computer instructions that repeats itself until a predetermined count or other test is satisfied, or until the process is interrupted by operator intervention”); Ex. 2026, 8 (“[a] set of statements in a program executed repeatedly, either a fixed number of times or until some condition is true or false”); *see also* Ex. 2038, 3 (defining “computation” as “an act, process, or method of computing”). Indeed, Patent Owner cites the definitions in asserting that “[a] ‘computational loop’ is an iterative sequence of computations that repeats until a prescribed condition is satisfied” (without mentioning any “per datum” requirement). PO Resp. 80.

The Specification of the ’324 patent is consistent with the technical dictionary definitions, and does not require that the set of computations executed repeatedly as a “loop” be on a particular piece of data. *See* Ex. 1001, col. 7, ll. 3–5 (explaining that computational process 610 “loops over the depth slices” and “loops over the shots” of a seismic imaging application), Fig. 7A (depicting three loops each performing computations a particular number of times as “k = 1, nz”; “j = 1, ny”; and “i = 1, nx”), Fig. 8B (depicting two loops as “i = 1, l” and “k = 1, m”). Patent Owner does not cite—and we do not find—any specific support in the Specification (including a reference purportedly incorporated by reference in the ’324 patent (Ex. 2037, “Caliga”)) for imposing a “per datum” requirement. *See* PO Resp. 71–72 (citing Ex. 1001,

Appendix D

col. 4, ll. 59–63, col. 6, ll. 1–30, col. 6, l. 47–col. 7, l. 48, col. 8, ll. 27–45, Figs. 4A–4B, 6B–6G, 7A, 8A–8B; Ex. 2037, 4, 5, 7, 12, 13, 16–19). Indeed, the Specification and Caliga do not use the term “datum.” Caliga, in fact, describes the opposite of what Patent Owner proposes, namely a “loop” of a set of computations executed repeatedly using *different* data. *See* Ex. 1076 ¶¶ 5–8; Ex. 2037, 5 (describing a “[l]oop over filter coefficients” where index j is incremented and different values for “Data”, “ R_j ,” and “ R_{j+1} ” are used at each iteration); Reply 36.¹²

Patent Owner further cites as support U.S. Patent No. 8,589,666 B2 (Ex. 2027, “the ’666 patent”), which was originally assigned to the same original assignee as the ’324 patent (SRC Computers, Inc.) but is unrelated to the ’324 patent. PO Resp. 72–73. The cited portions of the ’666 patent generally describe a “loop body” with “new data” (plural) “fed in on every clock tick,” and do not describe a “loop” as requiring a set of computations to be executed repeatedly “per datum” (singular). *See* Ex. 2027, col. 2, l. 64–col. 3, l. 23, col. 6, ll. 6–28. Dr. Homayoun also testifies that “[a] computational loop evaluates each piece of data multiple times, ‘a fixed number of times or until some condition is true or false,’” but does not explain the basis

12. We do not agree with Patent Owner that Petitioner’s Reply and Dr. Stone’s reply testimony are improper, as they respond directly to Patent Owner’s arguments in the Response regarding claim interpretation, where Patent Owner argued a “per datum” requirement for the first time. *See* Sur-Reply 8–9; Reply 36; Ex. 1076 ¶¶ 5–8; 37 C.F.R. § 42.23(b) (“A reply may only respond to arguments raised in the corresponding . . . patent owner response.”).

Appendix D

for that interpretation, in the Specification of the '324 patent or otherwise. Ex. 2111 ¶ 207.

Finally, in its Sur-Reply and at the oral hearing, Patent Owner relied on Figure 4B of the '324 patent as support for its proposed interpretation, arguing that the figure “unambiguously describes . . . two loops A and B operating a number of times (looping) on *each* dimension of data” and that our preliminary interpretation “would exclude from the claims the '324 Patent’s embodiments and figures.” *See* Sur-Reply 15–17; Tr. 52:20–58:21, 82:8–84:2. Patent Owner also stated during the oral hearing that “per datum” in its proposed interpretation could be replaced with “per dimension,” as Patent Owner did not intend there to be a “difference” between the terms. Tr. 53:10–11, 55:3–10, 82:8–17.

We disagree for two reasons. First, claim 1 already recites two functional units operating on two data dimensions forming two computational loops. A first functional unit “operate[s] upon a subsequent data *dimension* of said calculation forming a first computational loop” and a second functional unit concurrently “operate[s] upon a previous data *dimension* of said calculation forming a second computational loop” (emphases added). Dependent claims 2–5 specify different options for what those dimensions may comprise, namely “multiple vectors,” “multiple planes,” “multiple time steps,” or “multiple grid points” in the calculation. To the extent Patent Owner’s proposed interpretation would require that a functional unit operate on a particular data “dimension” of the calculation forming a computational loop, that is already encompassed in the claim language.

Appendix D

Second, the Specification indicates that a data “dimension” is not the same thing as an individual “datum.”

Figure 4B of the '324 patent is reproduced below.




Fig. 4B

Figure 4B depicts “multi-dimensional process 410 . . . effectuated such that multiple dimensions of data are processed by both Loops A (first loop 412) and B (second loop 414) such that the computing system logic is operative on every clock cycle.” Ex. 1001, col. 6, ll. 13–20. An application

Appendix D

will “pass” a subsequent dimension of a given problem through the first loop 412 of logic concurrently with the previous dimension of data being processed through the second loop 414. In practice, a “*dimension*” of data can be: *multiple vectors of a problem, multiple planes of a problem, multiple time steps in a problem and so forth.*

Id. at col. 6, ll. 24–30 (emphasis added). The Specification uses the phrasing “dimension” “of data” or “dimension” “of a given problem,” not “datum” individually, and provides examples of possible dimensions that mirror the language of dependent claims 2–5. It does not describe loop calculations applied to the same individual “datum” repeatedly.¹³

Having reviewed all of the cited evidence, we do not find sufficient support for limiting the term “computational loop” to require that computations be executed repeatedly “per datum.” Applying the broadest

13. The parties do not refer to anything in the prosecution history as supporting their proposed interpretations of “computational loop.” During prosecution, the applicants initially proposed (during the examiner interview discussed above in connection with the term “systolic”) amending claim 1 to recite that “said first of said instantiated functional units and said second of said instantiated functional units are within a nested loop of said calculation.” Ex. 1002, 208–209. The applicants then instead amended the claim to recite, rather than a “nested loop,” first and second “computational loops” executing concurrently with computed data seamlessly passed between them. *Id.* at 214–215.

Appendix D

reasonable interpretation of the claims in light of the Specification, we interpret “computational loop” to mean “a set of computations that is executed repeatedly, either a fixed number of times or until some condition is true or false.”

3. “Pass Computed Data Seamlessly Between Said Computational Loops”

Claim 1 recites that “said systolic implementation of said calculation enables said first computational loop and said second computational loop execute¹⁴ concurrently and pass computed data seamlessly between said computational loops.”

Petitioner argues that “pass computed data seamlessly between said computational loops” in claim 1 means “communicate computed data directly between functional units that are calculating computational loops.” Pet. 19; Reply 20. Based on the record at the time, we preliminarily agreed with Petitioner’s proposed interpretation in the Decision on Institution. Dec. on Inst. 23–26. Patent Owner argues that the phrase instead should be interpreted to mean “communicating the computed data over the reconfigurable routing resources.” PO Resp. 34–35.

According to the plain language of the claim, which recites “pass[ing] computed data seamlessly between said computational loops,” “seamlessly” refers to how

14. We read the claim language to mean that the systolic implementation enables the loops “to” execute concurrently and pass computed data seamlessly between the computational loops.

Appendix D

computed data is passed *between* the first and second computational loops (performed by the first and second functional units of the reconfigurable processor)—not, for example, how computed data is passed from either of the computational loops to any other component unrelated to the functional units, or vice versa.

The only other time “seamlessly” appears in the ’324 patent is in independent claims 25 and 51, which mirror the language of claim 1, reciting that computed data is passed “seamlessly” between systolic walls or columns of a calculation. The written description does not use the terms “seam” or “seamlessly.” Figure 2, though, shows functional units 204 interconnected without any intervening structures between them, which is consistent with Petitioner’s proposed interpretation. *See* Ex. 1001, col. 5, ll. 41–47. Figures 7A–7B and 8A–8B similarly show direct communication of data from one computational loop to another. *See id.* at col. 7, l. 42–col. 8, l. 6 (loops 702, 704, and 706), col. 8, ll. 27–55 (loops 812 and 814). Although the description of the figures does not use the terms “seam” or “seamlessly,” a person of ordinary skill in the art would have understood the claim language in light of that description, given that it describes and shows the interaction between multiple “loops,” as recited in claim 1. *See id.*

The prosecution history of the ’324 patent also supports Petitioner’s proposed interpretation. The applicants added the “seamlessly” language to claim 1 in the Office Action response discussed above in connection with the term “systolic.” Ex. 1002, 214–215. In response

Appendix D

to the § 112, first paragraph, written description rejection, the applicants indicated that data from one functional unit performing a computational loop is passed directly to another functional unit:

Data flows across the array between functional units, usually with different data flowing in different directions. . . . [I]n the Applicant's invention Systolic implementation will connect computational loops such that data from one compute loop will be passed as input data to a concurrently executing compute loop. In the Applicant's invention data computed by computation units or groups of functional units flows seamlessly and concurrently with data being computed by other groups of functional units. Thus, the process claimed by the Applicant therefore significantly increases the computing processes taking place in a reconfigurable processor.

Ex. 1002, 226. Although the language above primarily relates to the term “[s]ystolic,” we also find it significant for purposes of interpreting the “seamlessly” phrase because it refers to the limitation expressly in describing “Applicant's invention.” *See id.*; Ex. 1003 ¶ 98; *Shire Dev., LLC v. Watson Pharms., Inc.*, 787 F.3d 1359, 1366 (Fed. Cir. 2015) (holding that even where “prosecution history statements do not rise to the level of unmistakable disavowal, they do inform the claim construction”).

Finally, Petitioner's proposed interpretation is consistent with the testimony of Mr. Huppenthal, one of

Appendix D

the named inventors of the '324 patent, who describes “interpret[ing] standard high level language program constructs, such as the fact that the output variable from one loop is the input to another, and then implement[ing] this as a *storage free, seamless connection between the two loops* implemented on the FPGA,” and states:

By seamlessly I mean that the results of one loop streamed from that loop’s output to the input of the next loop *without being placed in a circuit element that required explicit address based on read or write operations such as a data register or memory* or through a switch that requires additional non-data content for routing purposes.

Ex. 2100 ¶ 78 (emphases added).

Based on our review of the full record after trial, Patent Owner’s arguments regarding the “seamlessly” language in claim 1 are not persuasive. *See* PO Resp. 34–42; Sur-Reply 20–21. First, Patent Owner relies on portions of the Specification describing interaction of cells “at the boundary,” “the problem of passing data over numerous boundaries (or seams) between processing elements in typical multi-processor systems,” and the '324 patent’s solution to that problem of “staying on a single FPGA chip, effectively eliminating the associated boundaries or seams from chip-to-chip communication.” PO Resp. 35–36, 39 (citing Ex. 1001, col. 2, ll. 25–48) (emphasis omitted); *see* Ex. 2111 ¶¶ 160–162. The cited portions of the Specification do not use the terms “seam”

Appendix D

or “seamlessly” and, unlike the portions cited by Petitioner describing Figures 7A–7B and 8A–8B, do not describe interaction between “loops” in the disclosed method. We find them less relevant than the portions discussed above, particularly given that the claim language at issue recites passing computed data seamlessly “between” the computational loops (performed by the functional units of the reconfigurable processor).

Second, Patent Owner relies on certain statements made by the applicants during prosecution of the '324 patent. PO Resp. 36–39; *see* Ex. 2111 ¶¶ 163–166. Three of those statements mirror the Specification language cited by Patent Owner and, importantly, predate the addition of the “seamlessly” language to claim 1 (and likewise do not use the terms “seam” or “seamlessly”). *See* Ex. 1002, 117–118, 148–150, 174–175. In addition, Patent Owner cites the following statement by the applicants in the Office Action response discussed above:

The instantiation of the at least one reconfigurable processor with at least two functional units enables each functional unit to communicate with each other. Certainly communication between other reconfigurable processors within the system would require [a] communication protocol but communication between functional units within an individual reconfigurable processor is free of such a requirement. To alleviate any confusion, the reference to the term “protocol” has been replaced with an “interconnection” between functional units

Appendix D

that is established by reconfigurable routing resources inside each chip.

Id. at 224–25; *see* PO Resp. 37–39. The cited language pertains to the “instantiating” step recited earlier in the claim, which also was amended in the response, including replacing “communications . . . independent of external and internal communication protocols” with “interconnects . . . based on reconfigurable routing resources.” *See* Ex. 1002, 214, 224–25. By contrast, the applicants added the “seamlessly” limitation to specify the interaction between the computational loops (i.e., that computed data is passed “seamlessly” between them).

Third, Patent Owner disputes the “directly” aspect of Petitioner’s proposed interpretation because “it would exclude standard FPGAs . . . since standard FPGAs contain reconfigurable routing resources (comprising buffers and switches) between the configurable logic blocks,” and thus would “exclude the very embodiments of the ’324 Patent.” PO Resp. 40–41, 47–48 (citing Ex. 1035, 31; Ex. 2078, 19–29, 32–34, 37–41, 46–51, 59–65). We agree with Petitioner that “[j]ust because a standard FPGA may include memories does not mean that when functional units are instantiated within such an FPGA that the memories are necessarily placed *between* functional units.” *See* Reply 25. Claim 1 recites two functional units being instantiated and concurrently performing two computational loops, with computed data passed seamlessly between the computational loops. Petitioner’s proposed interpretation of the “seamlessly” phrase correctly pertains to how computed data is passed *between*

Appendix D

the computational loops when so instantiated, and is not inconsistent with the mere existence of memory on an FPGA.

Fourth, Patent Owner argues that the word “directly” makes Petitioner’s proposed interpretation unreasonable because it “introduce[s] ambiguity and confusion.” PO Resp. 36, 40–41; *see* Ex. 2111 ¶¶ 169–176. According to Patent Owner, Dr. Stone on cross-examination could not identify what “intervening” structures would prevent direct communication, and contradicted himself by stating that “an intervening memory would not be a direct connection” and later that “if a register . . . were between the two processing elements then the connection would still be direct, but if the intervening structure were a buffer then the connection would not be direct.” PO Resp. 41 (citing Ex. 2064, 86:13–91:24). Patent Owner argues that if Dr. Stone’s view is adopted, “[t]he same circuit would be both direct and indirect, depending on where the boundaries of the ‘processing element’ are arbitrarily drawn with respect to intervening structures.” Sur-Reply 20–21.

We disagree that the word “directly” in Petitioner’s proposed interpretation is unclear or that Dr. Stone’s testimony about it introduces ambiguity. Dr. Stone testified as follows:

Q. Okay. And you mention the word “directly,” it was passing data directly between processing elements. What does that phrase mean to you or what’s the context? What are you trying to describe there?

Appendix D

A. That the data goes from first to the second *without going to something intervening*. It directly go – is connected immediately. Indirectly we – you go through one or more intervening places to get there.

Q. Okay. So would memory, if the data was going from one processing element to memory and then back to a processing element, is that something you would consider as an intervening thing?

A. Well, that would not be a direct connection of the output of the cell to the next cell. It says, “Between processing elements you’re directly connected.” *If you’re saying you have a processing element outputting to memory and then coming back to another processing element, that would not be direct.*

Ex. 2064, 85:14–86:12 (emphases added). When asked about a register in particular, Dr. Stone testified as follows:

Q. Well, how about a – a register? Would that be an intervening structure?

A. I – I’m puzzled because that – that register would be within – *within the processing element in my mind.*

Appendix D

Q. Okay.

A. If it's *within the processing element as a register*, yeah, I would put it there, then *the output of that register, if it's connected directly to the input of the next processing element, would be direct.*

Id. at 86:19–87:5 (emphases added); *see also id.* at 87:23–25 (“If the register is part of the processing element, then the connection would be direct.”).

Data is not communicated “directly” between processing elements when it is communicated through an intervening structure *between* them. *See id.* at 85:14–91:24. Examples of such a structure are memory and another processing element, but logically others are possible as well depending on how they are situated with respect to the processing elements. *See id.* at 86:13–18. When a register is *within* a processing element, however, data can still be communicated “directly” (provided there is a direct connection between the processing elements), as Dr. Stone explained. *See id.* at 85:14–91:24. We do not see any ambiguity in making this distinction. The boundaries of a processing element are not arbitrary as Patent Owner contends, but rather, in the context of a particular written document, would depend on how the reference describes the processing element and the communication of data to and from the processing element. *See, e.g.,* Ex. 1001, Fig. 2 (depicting functional units 204 of adaptive processor chip 202). Logically, also, in an arrangement with multiple processing elements, there must be some division between

Appendix D

the processing elements, otherwise there would not be multiple elements. *See* Tr. 14:1–5.

Applying the broadest reasonable interpretation of the claims in light of the Specification, we interpret “pass computed data seamlessly between said computational loops” to mean “communicate computed data directly between functional units that are calculating computational loops.”

4. “Establishing a Stream Communication Connection Between Functional Units”

Claim 15 depends from claim 1 and recites that “instantiating includes establishing a stream communication connection between functional units” (the “stream communication” limitation). Petitioner argues that “stream communication” should be interpreted to mean “communication of a data sequence.” Pet. 19–21. Patent Owner argues that “stream communication” means “a data path that acts like a queue connecting via the reconfigurable routing resources a producer and a consumer of data that operate concurrently.” PO Resp. 50–69. We did not preliminarily interpret “stream communication” in the Decision on Institution.

We begin with the language of the claim. Claim 15 recites that “instantiating” includes establishing a stream communication connection “between functional units.” “[I]nstantiating” in claim 15 refers to the “instantiating” step of parent claim 1 (emphasis added):

Appendix D

instantiating at least two of said functional units at the at least one reconfigurable processor to perform said calculation wherein only functional units needed to solve the calculation are instantiated and wherein each instantiated functional unit at the at least one reconfigurable processor interconnects with each other instantiated functional unit at the at least one reconfigurable processor based on reconfigurable routing resources within the at least one reconfigurable processor as established at instantiation

Thus, according to the plain language of the claim, what is being “instantiat[ed]” is the at least two “functional units.” The “reconfigurable routing resources,” based on which the instantiated functional units are “interconnect[ed],” are established “at” the instantiation.

Turning to the language of claim 15, regardless of whether “stream communication” is used as an adjective (as Petitioner contends) or a noun (as Patent Owner contends), the term appears immediately before and modifies the phrase “connection between functional units.” See Reply 30–31; Sur-Reply 10 & n.2. The claim, therefore, requires establishing a connection of a particular type, namely a “stream communication” type of connection. Nothing further in the language of the claim limits or sheds light on what that type of connection entails.

Nor does the Specification. The only other time “stream communication” appears in the ’324 patent is in claim 40, which mirrors the language of claim 15. The

Appendix D

written description never uses the term.¹⁵ It also does not differentiate between different types of connections between functional units.

Both parties, however, refer to patents that the '324 patent purports to incorporate by reference. *See* Pet. 20–21; PO Resp. 62–66. The '324 patent states that it is “related to the subject matter of,” and “incorporate[s] in [its] entirety by . . . reference,” U.S. Patent No. 6,434,687 B1 (Ex. 1014, “the '687 patent”).¹⁶ Ex. 1001, col. 1, ll. 9–20. Petitioner points to the phrase “streams of operands” in the '687 patent as supporting its proposed interpretation. Pet. 20. The cited portion describes multi-adaptive (MAP) processor 112 with reconfigurable array 42 having one or more high performance FPGAs. Ex. 1014, col. 9, ll. 1–6. “After configuration, the user array 42 can perform whatever function it was programmed to do.” *Id.* at col. 9, ll. 11–12. “In order to maximize its performance for vector processing, the array 42 should be able to access two streams of operands simultaneously” by using a chain port “connector allow[ing] the MAP element 112 to use data provided to it by a previous MAP element 112.” *Id.* at col. 9, ll. 12–26.

15. Petitioner argues that its proposed interpretation is consistent with a portion of the Specification describing systolic wall of computation 712 in Figure 7C and stating that “communication of values between adjacent rows . . . in the virtual wall can occur without storing values to memory.” Ex. 1001, col. 8, ll. 7–12; *see* Pet. 20–21; Ex. 1003 ¶ 117. The cited portion, however, never uses the term “stream communication.”

16. Petitioner challenges the '687 patent in Case IPR2018-01594.

Appendix D

We find the cited portion to be of limited relevance. The '687 patent uses “streams of operands” rather than “stream communication” and does not describe the “streams of operands” other than the single sentence above. Also, given the length of the '687 patent disclosure, it is unclear exactly what aspects of the '687 patent are incorporated in the '324 patent or how they would apply in the context of the '324 patent invention. *See* Ex. 1001, col. 1, ll. 9–20 (merely stating that the '324 patent is “related to the subject matter of” the '687 patent). If anything, the cited portion of the '687 patent simply indicates that certain data (i.e., “operands”) can be streamed from one MAP element to another.

Patent Owner similarly quotes other excerpts from the '687 patent describing the use of a chain port connection, input buffer 40, and output First-In-First-Out (FIFO) buffer 74 in support of its contention that “stream communication” requires a “data path” acting like a “queue.” PO Resp. 63–66 (citing Ex. 1014, col. 7, ll. 43–57, col. 8, ll. 7–26, col. 9, ll. 2–67); *see* Ex. 1014, Figs. 6–7. We are not persuaded for the same reasons as Petitioner’s citation. The cited portions merely describe an exemplary embodiment of the '687 patent that uses memory buffers. Other than the single use of “streams of operands,” there is no connection between the '687 patent disclosure and anything pertaining to streaming. *See* Sur-Reply 12 (acknowledging that the '687 patent is “a different patent referring to a different invention pertaining to internet communications”). Certainly, the cited portions do not use or define what is meant by “establishing a stream communication connection between functional units”

Appendix D

in the context of the data processing method recited in claim 15.

Patent Owner also relies on U.S. Patent No. 6,339,819 B1 (Ex. 2085, “the ’819 patent”), of which the ’687 patent is a continuation-in-part, as incorporated by reference in the ’324 patent. PO Resp. 62–64. The ’324 patent describes “a representative systolic wavefront operation in the form of a reservoir simulation function” where

the computation of fluid flow properties are communicated to neighboring cells 710 and, importantly, this computation can be scheduled to eliminate the need for data storage. In accordance with the technique of the present invention, a set of cells can reside in an adaptive processor and the pipeline of computation can extend across multiple adaptive processors. Communication overhead between multiple adaptive processors may be advantageously minimized through the use of MAPTM adaptive processor chain ports as disclosed in U.S. Pat. No. 6,339,819 issued on Jan. 15, 2002 for: “Multiprocessor With Each Processor Element Accessing Operands in Loaded Input Buffer and Forwarding Results to FIFO Output Buffer,” assigned to SRC Computers, Inc., assignee of the present invention, the disclosure of which is herein specifically incorporated by this reference.

Ex. 1001, col. 7, ll. 42–46, col. 7, l. 59–col. 8, l. 6, Figs. 7A–7B. According to Patent Owner, the Specification

Appendix D

thereby discloses “the concept of using chain ports and a FIFO buffer for chip to chip communications,” such that a person of ordinary skill in the art “would recognize that this inter-chip communication concept can be adapted to intra-chip communications between functional units on the same chip—e.g., a FIFO within the chip.” PO Resp. 62–63. The cited disclosure is of a “representative” (i.e., exemplary) embodiment where communication “may” occur “between multiple adaptive processors” using chain ports and buffers. *See* Ex. 1001, col. 7, ll. 42–46, col. 7, l. 59–col. 8, l. 6. Again, there is no connection to anything pertaining to streaming data between functional units, as the disclosure does not use the word “stream” at all. The incorporated ’819 patent also includes much of the same disclosure as the ’687 patent and suffers from the same problems explained above. *Compare* Ex. 1014, col. 4, l. 15–col. 20, l. 35, Figs. 1–11B, *with* Ex. 2085, col. 4, l. 31–col. 21, l. 43, Figs. 1–11B. We are not persuaded that the incorporated patents support Patent Owner’s contention that “stream communication” in claim 15 requires a “data path” acting like a “queue.”

Turning next to the prosecution history, we find some guidance as to the scope of “stream communication.” Claim 15 was amended to its current form in the applicants’ first Office Action response. Ex. 1002, 110. The applicants, however, did not discuss the language of claim 15 in the response and instead attempted to distinguish the cited prior art as applied to parent claim 1. *Id.* at 116–122. Patent Owner points to various statements later in the prosecution history, only one of which addressed the “stream communication” language of claim 15 directly. *See*

Appendix D

PO Resp. 66–69 (citing Ex. 1002, 128–129, 147–150, 174–175, 208, 224–225). Specifically, in the subsequent Office Action, the examiner found the “stream communication” limitation taught by a particular reference, noting that the reference

taught minimiz[ing] interconnections of processing elements and the matrix and vector signal subsets are specifically formed so that they need to be inputted to only one row and one column[] and yet still [are] properly processing systolically along all dimensions within the array. . . . Consequently the stream of communication between functional units is established as the interconnections are made and data is transferred systolically in at least one stream between processors.

Ex. 1002, 128–129. This explanation indicates that establishing a stream communication connection allows data to be “transferred” as a “stream.” Because the applicants and examiner do not appear to have ever mentioned a “data path” acting like a “queue” when addressing the “stream communication” limitation, though, we are not persuaded that the prosecution history supports a narrower interpretation including those limitations, as Patent Owner contends.

Based on what information there is in the intrinsic record, in particular the claim language surrounding “stream communication” and the single reference to the “stream communication” limitation in the prosecution

Appendix D

history, we conclude that the plain meaning of “establishing a stream communication connection between functional units” is establishing a connection over which data is streamed between functional units. The interpretation gives meaning to the “stream communication” term and results in claim 15 appropriately having a different scope than claim 1, as claim 15 requires data to be communicated over a particular type of connection that permits streaming, whereas claim 1 would encompass any type of connection. We do not find any support in the intrinsic record for Patent Owner’s proposed interpretation requiring a “data path” acting like a “queue.”

Patent Owner points to a number of sources of extrinsic evidence in support of its narrower interpretation. As an initial matter, none of those references is a technical dictionary, textbook, or similar source that typically are referenced to help determine a term’s ordinary meaning. *See* Tr. 37:15–21, 38:10–17. Instead, Patent Owner refers to unrelated patents and other types of documents. *See* PO Resp. 53–62; Sur-Reply 12–14. For example, Patent Owner cites the ’666 patent, which discloses:

A stream is a data path between a producer and consumer of data, where the producer and consumer run concurrently. The path between the producer and consumer is made up of a data connection, a “valid” signal, and a reverse direction “stall” signal. FIG. 1 shows typical signals used in a stream connection as is well known and will be recognized by one skilled in the relevant art. The use of a First-In-First-

Appendix D

Out buffer 110, or “FIFO” buffer, removes the need for tight synchronization between the producer 120 and consumer 130. The producer 120 will generate data values 125 at its own rate, allowing them to accumulate in the FIFO buffer 110. As the FIFO buffer 110 approaches becoming full, it will issue a stall signal 140 to the producer 120 so that it will suspend the generation of data values 125 until the stall signal is released. The consumer 130 will take 150 values 145 from the FIFO buffer at its own rate and as the values 145 are available.

Ex. 2027, col. 2, ll. 39–54 (emphases added); *see* PO Resp. 53–56; Sur-Reply 12. According to Patent Owner and Dr. Homayoun, this disclosure of using a FIFO buffer to account for different data rates and allow the producer and consumer to run concurrently is “entirely consistent with . . . instantiating reconfigurable [routing] resources to seamlessly communicate computed data between processing elements” in claim 1 and “provid[ing] stream communication between those processing elements” in claim 15. PO Resp. 55–56; *see* Ex. 2111 ¶¶ 155–156.

Although sharing an assignee, the ’666 patent is unrelated to the ’324 patent, has an effective filing date (July 10, 2006) nearly four years after that of the ’324 patent (October 31, 2002), and has a different named inventor. *See* Reply 32; Ex. 1013, 2–3. We also do not see any relationship between the ’666 patent’s description of streaming in connection with a different invention and the ’324 patent’s systolic implementation.

Appendix D

We are not persuaded, therefore, that the '666 patent supports reading “stream communication” in claim 15 to require a “data path” acting like a “queue” as Patent Owner contends, particularly when the intrinsic record provides some indication of a broader interpretation. “[E]xtrinsic evidence may be used only to assist in the proper understanding of the disputed limitation; it may not be used to vary, contradict, expand, or limit the claim language from how it is defined, even by implication, in the specification or file history.” *Bell Atlantic Network Servs., Inc. v. Covad Comm’cns Group, Inc.*, 262 F.3d 1258, 1269 (Fed. Cir. 2001).

Likewise, Patent Owner cites its own product documentation, dated 2002–2007, as “describ[ing] a stream as a data structure that allows flexible communication between concurrent producer and consumer loops” using an internal buffer, as well as a 1993 Argonne National Laboratory paper. PO Resp. 56–59 (citing Ex. 2107, 94–98; Ex. 2028, 31; Ex. 2100 ¶ 79; Ex. 2111 ¶ 157). We do not see—and Patent Owner does not explain sufficiently—why such product documentation shows how a person of ordinary skill in the art would have understood the language of the “stream communication” limitation in claim 15. Moreover, whereas Patent Owner’s proposed interpretation is for “stream communication” to be a “data path,” the product documentation and paper indicate that a “stream” is a “data structure.” *See id.* at 50; Ex. 2107, 94; Ex. 2028, 31. They are different concepts. *See* Ex. 2065, 154–155 (defining “data path” as “[t]he route that a signal follows as it travels through a computer network” and “data structure” as “[a]n organizational scheme, such as

Appendix D

a record or array, that can be applied to data to facilitate interpreting the data or performing operations on it”). Finally, Patent Owner quotes a number of other patents and patent application publications as allegedly showing that “stream communication” requires a “data path” acting like a “queue.” PO Resp. 59–62. The references use the words “stream” or “streaming” (not “stream communication connection”) in the context of describing embodiments of their own disclosed inventions. We find that they have little probative value as well.

We are not persuaded that the extrinsic evidence cited by Patent Owner shows that there was a uniform understanding of persons of ordinary skill in the art of the meaning of the term “stream communication,” alone or in the context of a “stream communication connection” between functional units. Indeed, one technical dictionary, submitted but not cited by either party in connection with the term “stream communication,” defines “stream” more broadly as “[a]ny data transmission, such as the movement of a file between disk and memory, that occurs in a continuous flow” (as a noun) or “[t]o transfer data continuously, beginning to end, in a steady flow” (as a verb). Ex. 2065, 509. This is consistent with Petitioner’s view that the word “stream” in claim 15 “indicates some kind of movement, some kind of transfer of data,” as opposed to a particular data structure as Patent Owner contends. *See* Tr. 71:5–15, 72:9–10. Notably, the same technical dictionary that defines “stream” as a “data transmission” also defines “queue” as “[a] multi-element data structure from which . . . elements can be removed only in the same order in which they were inserted.”

Appendix D

Ex. 2065, 443, 509 (emphases added). This indicates that a “stream” and a “queue” are distinct concepts, and contradicts Dr. Homayoun’s testimony that “[a] person of ordinary skill in the art would understand that the term ‘stream’ refers to a specific type of structure called a queue.” *See* Ex. 2111 ¶ 152.

Applying the broadest reasonable interpretation of the claims in light of the Specification, we give “establishing a stream communication connection between functional units” its plain meaning, namely “establishing a connection over which data is streamed between functional units.” We are not persuaded that the extrinsic evidence relied upon by Patent Owner supports a different interpretation. No further interpretation is necessary to resolve the parties’ disputes over the asserted grounds of unpatentability in this proceeding.¹⁷ *See Nidec Motor Corp. v. Zhongshan*

17. Although our interpretation differs from those proposed by the parties for “stream communication,” we did not preliminarily interpret the term in the Decision on Institution, both parties had the opportunity to present arguments regarding the term in their papers and at the oral hearing, and the issue was discussed extensively during the oral hearing. *See* Tr. 20:21–22:18, 33:19–41:14, 68:19–73:17, 84:3–86:13; *TQ Delta, LLC v. DISH Network LLC*, 929 F.3d 1350, 1355 (Fed. Cir. 2019) (holding that the Board did not improperly “change course” under *SAS Institute, Inc. v. ComplementSoft, LLC*, 825 F.3d 1341 (Fed. Cir. 2016), *rev’d and remanded on other grounds*, *SAS Institute, Inc. v. Iancu*, 138 S. Ct. 1348 (2018), by “construing the limitation in the Final Written Decision because it did not construe the term in its Decision to Institute”); *WesternGeco LLC v. ION Geophysical Corp.*, 889 F.3d 1308, 1328–29 (Fed. Cir. 2018) (holding that “the Board is not bound to adopt either party’s preferred articulated

Appendix D

Broad Ocean Motor Co., 868 F.3d 1013, 1017 (Fed. Cir. 2017) (“Because we need only construe terms ‘that are in controversy, and only to the extent necessary to resolve the controversy,’ we need not construe [a particular claim limitation] where the construction is not ‘material to the . . . dispute.’” (citations omitted)).

D. Legal Standards

“Anticipation requires that every limitation of the claim in issue be disclosed, either expressly or under principles of inherency, in a single prior art reference,” *Corning Glass Works v. Sumitomo Elec. U.S.A., Inc.*, 868 F.2d 1251, 1255–56 (Fed. Cir. 1989), and that the claim limitations be “arranged or combined in the same way as recited in the claim,” *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1371 (Fed. Cir. 2008). However, “the reference need not satisfy an *ipsissimis verbis* test.” *In re Gleave*, 560 F.3d 1331, 1334 (Fed. Cir. 2009). “In an anticipation analysis, the dispositive question is whether a skilled artisan would ‘reasonably understand or infer’ from a prior art reference that every claim limitation is disclosed in that single reference.” *Acoustic Tech., Inc.*

construction of a disputed claim term” and “was permitted to issue a new construction in the final written decision given that claim construction was a disputed issue during the proceedings”); *Intellectual Ventures II LLC v. Ericsson Inc.*, 686 F. App’x 900, 906 (Fed. Cir. 2017) (unpublished) (“The Board is not constrained by the parties’ proposed constructions and is free to adopt its own construction,” but “after the Board adopts a construction, it may not change theories without giving the parties an opportunity to respond.”).

Appendix D

v. Itron Networked Solutions, Inc., 949 F.3d 1366, 1373 (Fed. Cir. 2020) (citation omitted). “Expert testimony may shed light on what a skilled artisan would reasonably understand or infer from a prior art reference.” *Id.*

A claim is unpatentable for obviousness if, to one of ordinary skill in the pertinent art, “the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007) (quoting 35 U.S.C. § 103(a)). The question of obviousness is resolved on the basis of underlying factual determinations, including “the scope and content of the prior art”; “differences between the prior art and the claims at issue”; and “the level of ordinary skill in the pertinent art.” *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). Additionally, secondary considerations, such as “commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. As indicia of obviousness or nonobviousness, these inquiries may have relevancy.” *Id.*

A patent claim “is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art.” *KSR*, 550 U.S. at 418. An obviousness determination requires finding “both ‘that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.’”

Appendix D

Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd., 821 F.3d 1359, 1367–68 (Fed. Cir. 2016) (citation omitted); see *KSR*, 550 U.S. at 418 (for an obviousness analysis, “it can be important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does”). “Although the *KSR* test is flexible, the Board ‘must still be careful not to allow hindsight reconstruction of references . . . without any explanation as to *how* or *why* the references would be combined to produce the claimed invention.’” *TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1066 (Fed. Cir. 2016) (citation omitted). Further, an assertion of obviousness “cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *KSR*, 550 U.S. at 418 (quoting *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)); accord *In re NuVasive, Inc.*, 842 F.3d 1376, 1383 (Fed. Cir. 2016) (stating that “conclusory statements” amount to an “insufficient articulation[] of motivation to combine”; “instead, the finding must be supported by a ‘reasoned explanation’” (citation omitted)); *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016) (“To satisfy its burden of proving obviousness, a petitioner cannot employ mere conclusory statements. The petitioner must instead articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness.”).

*Appendix D***E. Anticipation Ground Based on Splash2
(Claims 1, 15, 18, 21, and 22)****1. Splash2**

Splash2¹⁸ is a book describing the Splash 2 reconfigurable computer system created by the Supercomputing Research Center in the 1990s. Ex. 1007, xi. “Splash 2 is an attached processor system using Xilinx XC4010 FPGAs as its processing elements.” *Id.* “[T]he XC4010 contains a 20 x 20 array of Configurable Logic Blocks (CLBs)” *Id.* at 11.

Figure 2.3 of Splash2 is reproduced below.

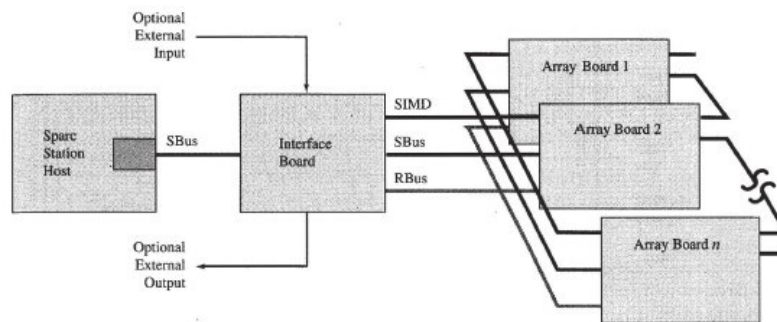


FIGURE 2.3 Splash 2 System Architecture

18. The first 11 pages of Splash2 and a four-page summary article about the Splash 2 system were made of record during prosecution of the '324 patent, but not the remainder of the book, including Chapter 8 on which Petitioner primarily relies. *See* Ex. 1001, code (56); Ex. 1002, 74; Ex. 1060; Pet. 1–2. RaPiD, Jeong, Chunky SLD, and Roccatano were not of record during prosecution of the '324 patent. *See* Ex. 1001, code (56).

Appendix D

Figure 2.3 depicts the system architecture of the Splash 2 system, including a set of array boards connected to a SPARCstation 2 host via an interface board. *Id.* at 12–13.

Figure 2.4 of Splash2 is reproduced below.

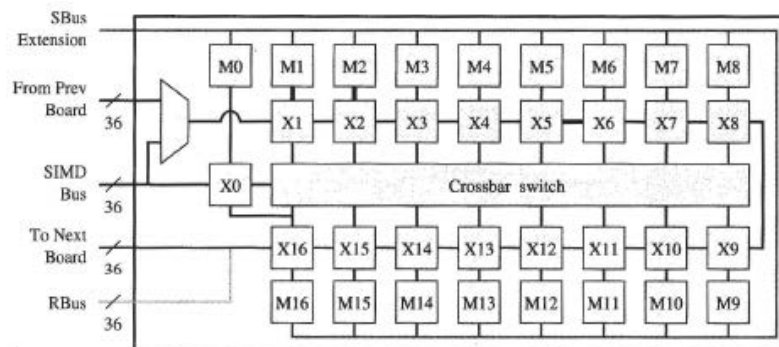


FIGURE 2.4 Array Board Architecture

Figure 2.4 depicts the architecture of one of the array boards. Each array board “contains 17 Xilinx XC4010 FPGA chips as its processing elements. Sixteen of these are connected in a linear array to create a linear data path and the seventeenth provides a broadcast capability to the other 16 chips.” *Id.* at 13 (citations omitted). “Viewed as a machine with a linear data path, the [Single Instruction, Multiple Data (SIMD)] Bus can be used to transmit data from the Interface Board to the first FPGA on the first Array Board. The data can then be moved through the linear data path on that board, then to the first FPGA on the second Array Board, and so on.” *Id.* at 14. The linear data path is bidirectional. *Id.*

Appendix D

Splash2 describes a number of applications programmed on the Splash 2 system. *Id.* at xi. Chapter 8 describes “two systolic array architectures for [genetic] sequence comparison and their implementations on the Splash 2 programmable logic array.” *Id.* at 97. Splash2 discloses that “[i]n comparing two sequences, it is useful to quantify their similarity in terms of a distance measure,” and one such measure is the “edit distance between two sequences,” which is “the minimum cost of transforming one sequence to the other with a sequence of the following operations: deletion of a character, insertion of a character, and substitution of one character for another.” *Id.* at 98. Each operation has a cost, and the transformation cost is “the sum of the costs of the individual operations.” *Id.* Splash2 discloses a “well-known dynamic programming algorithm” for computing edit distance, and describes how two example nucleotide sequences *TCTAGACC* and *GCATAAGC* would be compared using the equations for that algorithm. *Id.* at 98–99.

Splash2 discloses that the edit distance algorithm has “inherent parallelism” in that “each entry in the distance matrix” for comparing each character of the source sequence to each character of the target sequence “depends on adjacent entries,” and such parallelism can be “exploited to produce systolic algorithms in which communication is limited to adjacent processors.” *Id.* at 98–100, Figs. 8.2–8.3. Specifically, entries of the matrix on the same antidiagonal can be computed in parallel. *Id.* at 100, Fig. 8.4. Splash2 describes two architectures for doing so, one using a bidirectional systolic array and one using a unidirectional systolic array. *Id.* at 100–108.

Appendix D

Splash2 discloses that “[b]oth the bidirectional and unidirectional systolic arrays have been implemented on the Splash 2 programmable logic array, with versions for DNA and protein sequences.” *Id.* at 104.

a) Unidirectional Array Implementation

Figure 8.9 of Splash2 is reproduced below.




FIGURE 8.9 Data Flow through the Unidirectional Systolic Array. The source sequence is first loaded into the array. The target sequences are then streamed through the array. The tag acts as a simple instruction telling each PE how to process the incoming data. The SR tag instructs an empty PE to load the source character and distance from the input stream. The PR tag marks the end of the source stream. The TG tag signals a target character. Multiple source and target sequences can be carried on the input stream for uninterrupted pipelined processing.

As shown in Figure 8.9, “data flows through the unidirectional array in one direction. The source sequence is loaded once and stored in the array starting from the leftmost [processing element (PE)]. The target sequences are streamed through the array one at a time, separated by control characters.” Ex. 1007, 103. “In this configuration, each PE computes the distances in one row

Appendix D

of the distance matrix. At each time step, the PEs compute the distances along a single antidiagonal in the distance matrix . . .” *Id.* at 104. Splash2 discloses that the “DNA version of the unidirectional array,” for example, has 248 processing elements that would be capable of comparing millions of characters per second. *Id.* at 107.

Figure 8.12 of Splash2 is reproduced below.

```

loop
  if (TAGin = SR) then
    if (SRCch =  $\emptyset$ ) then
      SRCch  $\leftarrow$  CHRin
      CHRout  $\leftarrow$   $\emptyset$ 
      DSTout  $\leftarrow$  PDSTin
    else
      CHRout  $\leftarrow$  CHRin
    endif
    PDSTout  $\leftarrow$  PDSTin
  else-if (TAGin = PR) then
    if (SRCch =  $\emptyset$ ) then
      DSTout  $\leftarrow$  PDSTin
    endif
    PDSTout  $\leftarrow$  DSTin
    CHRout  $\leftarrow$  CHRin
  else-if (TAGin = TG) then
    if (SRCch  $\neq$   $\emptyset$ ) and (CHRin  $\neq$   $\emptyset$ ) then
      DSTout  $\leftarrow$  min  $\begin{cases} \text{PDSTout} + \psi(\text{SRCch}, \text{CHRin}), \\ \text{DSTin} + \psi(\text{SRCch}, \emptyset), \\ \text{DSTout} + \psi(\emptyset, \text{CHRin}) \end{cases}$ 
    else-if (SRCch =  $\emptyset$ ) then
      DSTout  $\leftarrow$  DSTin
    endif
    PDSTout  $\leftarrow$  DSTin
    CHRout  $\leftarrow$  CHRin
  endif
  TAGout  $\leftarrow$  TAGin
endloop

```

FIGURE 8.12 Code executed by each PE in the unidirectional array

Appendix D

Figure 8.12 depicts “[t]he algorithm executed by each PE in the unidirectional array” that “compares one source sequence to a single target sequence.” *Id.* at 104–105. As shown above, SRCch and CHRin represent the character of the source sequence and target sequence being processed, respectively; DSTin and PDSTin represent the distance stream being input; and DSTout and PDSTout represent the computed distance stream being output. *Id.* Splash2 discloses that “[a] unidirectional array of length n can compare a source sequence of length at most n to a target sequence of length m in $O(n + m)$ steps.” *Id.* at 104.

b) Bidirectional Array Implementation

Figure 8.5 of Splash2 is reproduced below.




FIGURE 8.5 Data Flow through the Bidirectional Systolic Array. The source and target sequences are streamed through the array in opposite directions. A comparison is performed when a source character and a target character meet in a PE.

As shown in Figure 8.5, “[t]he source and target sequences enter the array on opposite ends and flow in opposite

Appendix D

directions at the same speed.” Ex. 1007, 102. Splash2 discloses that

there is one distance stream associated with each character stream. At each step, the contents of the streams represent the characters to be compared and the distances along one of the antidiagonals of the distance matrix. At the end of the computation, the resulting edit distance is transported out of the array on the distance streams.

Id. (footnote omitted). In the bidirectional array implementation, each processing element “computes the distances along a particular diagonal of the distance matrix.” *Id.* at 100. Splash2 discloses that the “DNA version of the bidirectional array,” for example, has 384 processing elements that would be capable of comparing millions of characters per second. *Id.* at 107.

Figure 8.7 of Splash2 is reproduced below.

```

loop
  if (SCin  $\neq$   $\emptyset$ ) and (TCin  $\neq$   $\emptyset$ ) then
    PEDist  $\leftarrow$  min  $\begin{cases} \text{PEDist} + \psi(\text{SCin}, \text{TCin}), \\ \text{TDin} + \psi(\text{SCin}, \emptyset), \\ \text{SDin} + \psi(\emptyset, \text{TCin}) \end{cases}$ 
  else-if (SCin  $\neq$   $\emptyset$ ) then
    PEDist  $\leftarrow$  SDin
  else-if (TCin  $\neq$   $\emptyset$ ) then
    PEDist  $\leftarrow$  TDin
  endif
  SCout  $\leftarrow$  SCin
  TCout  $\leftarrow$  TCin
  SDout  $\leftarrow$  PEDist
  TDout  $\leftarrow$  PEDist
endloop

```

FIGURE 8.7 Code Executed by Each PE in the Bidirectional Array

Appendix D

Figure 8.7 depicts the algorithm executed by each processing element in the bidirectional array. *Id.* at 101. As shown above, SCin and TCin represent the characters of the source sequence and target sequence being processed, respectively; SCout and TCout represent respective characters being output; SDin and TDin represent the source and target distance stream being input, respectively; and SDout and TDout represent the computed source and target distance stream being output, respectively. *Id.*

Splash2 discloses that “[c]omparing sequences of lengths m and n requires at least $2\max(m + 1, n + 1)$ processors. The number of steps required to compute the edit distance and to transport it out of the array is proportional to the length of the array.” *Id.* at 103 (footnote omitted). Splash2 explains that the bidirectional array implementation had certain inefficiencies for database search operations, such as the fact that “[a]t each computational step, at most half of the PEs are active” and the source and target sequences were limited to half of the array’s length. *Id.* Those inefficiencies were remedied by the unidirectional array implementation according to Splash2. *Id.*

2. Claim 1

Petitioner argues that Splash2 is prior art under 35 U.S.C. §§ 102(a) and 102(b),¹⁹ and explains in detail how

19. Petitioner provides evidence supporting its contention that Splash2, as well as RaPiD, Jeong, Chunky SLD, and Roccatano, are

Appendix D

the reference discloses every limitation of claim 1, relying on the testimony of Dr. Stone as support. *See* Pet. 21–46;²⁰ Ex. 1003 ¶¶ 118–179, 184–193, 203–208, 210–216, 220–233, 235–279. Petitioner argues that both the description of the unidirectional array and the description of the bidirectional array in Splash2 disclose certain limitations of claim 1. Pet. 21–46. We address both implementations below.

a) Unidirectional Array Implementation**(1) Undisputed Limitations**

Petitioner argues that Splash2 discloses a “method for data processing,” as recited in claim 1, where “the characters of a source sequence of genetic information are compared against a target sequence of genetic information and an edit distance is calculated based on that comparison.” Pet. 29–30. Petitioner contends that Splash2 discloses a reconfigurable computing system comprising “at least

prior art printed publications under 35 U.S.C. §§ 102(a) and 102(b). *See* Pet. 21–22, 56, 68 (citing Ex. 1065 ¶¶ 8–12; Ex. 1066 ¶¶ 37–46, 75, 77, 80–92; Ex. 1068 ¶¶ 4–6, 9–12); -1602 Pet. 63 (citing Ex. 1066 ¶¶ 47–55; Ex. 1068 ¶¶ 7–12); -1603 Pet. 67 (citing Ex. 1066 ¶¶ 65–73; Ex. 1067); Ex. 1007, Cover 6 (Library of Congress stamp for Splash2). Patent Owner does not assert otherwise in its Response, and we agree that the references are prior art for the reasons stated by Petitioner.

20. Petitioner asserts many of the same arguments, such as those regarding claim 1, in all three of its Petitions. *See* Pet. 21–46; -1602 Pet. 22–52; -1603 Pet. 25–54. Where applicable, we cite the papers in Case IPR2018-01601 for convenience.

Appendix D

one reconfigurable processor” (i.e., the array boards with FPGAs) with “a plurality of functional units” (i.e., the matrix of CLBs configured as processing elements), and “transforming” an “algorithm” (i.e., the edit distance algorithm disclosed in Splash2) into a “calculation” (i.e., calculation of the edit distance for two sequences). *Id.* at 29–32. According to Petitioner, the calculation is “systolically implemented” in the unidirectional array implementation because “each Processing Element ‘computes the distances in one row of the distance matrix,’ with “a target sequence streamed through the array, and distance information continually output to the downstream neighboring Processing Element (via DSTout and PDSTout),” and a character comparison is triggered by the arrival of data rather than a program counter or clock. *Id.* at 32–34. We agree, given our interpretation of the term “systolic” as “the characteristic of rhythmically computing and passing data directly between processing elements in a manner that is transport triggered (i.e., by the arrival of a data object) rather than a program counter or clock driving movement of the data.” *See supra* Section II.C.1.

Petitioner further contends that Splash2 discloses “instantiating” at least two functional units to perform the calculation (i.e., loading information into the FPGAs to create the necessary processing elements) where “only functional units needed to solve the calculation are instantiated” because the FPGA components “clear all configuration memory before any configuration occurs” and the components instantiated for a particular calculation are tailored to that combination. Pet. 34–35.

Appendix D

Petitioner argues that “each instantiated functional unit . . . interconnects with each other instantiated functional unit . . . based on reconfigurable routing resources within the at least one reconfigurable processor as established at instantiation,” as shown by the target sequence streaming through the processing elements in Splash2. *Id.* at 35–36. Finally, “systolically linked lines of code” of the calculation are “instantiated as clusters of functional units” (i.e., the processing elements on each FPGA) according to Petitioner because “each Processing Element continually computes distance information for each character comparison and passes it to a neighboring Processing Element” “without storage of that information in memory between processing elements, without the data being driven by a clock or program counter, and triggered by the arrival of tags in the data stream.” *Id.* at 36–38.

Patent Owner in its Response does not dispute that Splash2 discloses the above limitations of claim 1. Petitioner’s analysis for each of the limitations, supported by the testimony of Dr. Stone, which we credit, is persuasive. *See id.* at 29–38; Ex. 1003 ¶¶ 153–179, 184–193, 203–208, 210–216, 220–233, 235–242.

(2) Disputed Limitation: “Computational Loops”

Claim 1 recites the following “computational loop” limitations:

utilizing a first of said instantiated functional units to operate upon a subsequent data dimension of said calculation forming a first computational loop; and

Appendix D

substantially concurrently utilizing a second of said instantiated functional units to operate upon a previous data dimension of said calculation forming a second computational loop.

Petitioner argues that a person of ordinary skill in the art would have understood that the loop/endloop code, shown in Figure 8.12 above, “constitutes a time step of a ‘data dimension’ comprising multiple time steps because each such instance of code executes in a single time step, and the systolic computation comprises multiple time steps.” Pet. 38 (citing Ex. 1003 ¶¶ 245–246). According to Petitioner, a processing element executes the code “in a single time step,” taking in certain inputs from the preceding processing element (e.g., TAG_{in}, CHR_{in}, distance information DST_{in} and PDST_{in}) and passing along outputs (e.g., TAG_{out}, CHR_{out}, distance information DST_{out} and PDST_{out}) to the downstream neighboring processing element for “each character comparison” performed. *Id.* at 26, 28, 33, 37–38, 46; *see* Ex. 1003 ¶¶ 135, 139, 167, 189, 192, 247–259.

Appendix D

Petitioner provides the following annotated version of Figure 8.13 of Splash2 (Pet. 39).




FIGURE 8.13 Trace of Unidirectional Array When Comparing the Sequences TCTAGACC and GCATAAGC

Annotated Figure 8.13 depicts the source sequence *TCTAGACC* pre-loaded in the processing elements (row SRCch), and the characters of the target sequence *GCATAAGC* streamed through the processing elements systolically (row CHRin). *Id.* at 39–41 (citing Ex. 1003 ¶¶ 247–259). For example, at Time Step 2, Processing Element 1 compares the first target character *G* to the first source character *T*. *Id.* at 40. At Time Step 3, (1) Processing Element 1 compares the second target

Appendix D

character *C* to the first source character *T*, and (2) Processing Element 2, which has received the first target character *G* streamed through the array, compares it to the second source character *C*. *Id.* Petitioner argues that, as shown in the example above, Processing Element 2 operates on a “subsequent data dimension” (i.e., comparing the second source character to the streamed target characters—the second row of the distance matrix) forming a “first computational loop” (i.e., “Processing Element 2 executing the loop instructions disclosed in Figure 8.12”); and Processing Element 1 operates on a “previous data dimension” (i.e., comparing the first source character to the streamed target characters—the first row of the distance matrix) forming a “second computational loop” (i.e., “Processing Element 1 executing the loop instructions disclosed in Figure 8.12”). *Id.* at 41 (emphases omitted). According to Petitioner, “[t]he same loop instructions are executed in each Processing Element, but they are still ‘first’ and ‘second’ computational loops because they are necessarily different instances of those loop instructions.” *Id.* Petitioner’s analysis for the “computational loop” limitations of claim 1, supported by the testimony of Dr. Stone, which we credit, is persuasive. *See* Ex. 1003 ¶¶ 243–259.

Patent Owner’s arguments are largely premised on its proposed interpretation of “computational loop” as requiring repeated execution “per datum,” which we do not adopt. *See* PO Resp. 80–87 (citing Ex. 2111 ¶¶ 127, 194–209; *supra* Section II.C.2. Patent Owner asserts that “[t]here is no disclosure [in Splash2] of looping or repeating of a computation multiple times *for each data*

Appendix D

until a condition is met or a number of repetitions has been satisfied,” as allegedly represented in Figure 4B of the ’324 patent. PO Resp. 80–81, 84–85 (emphasis added). We disagree that the claim requires such a process, and instead interpret “computational loop” to mean “a set of computations that is executed repeatedly, either a fixed number of times or until some condition is true or false.” *See supra* Section II.C.2.

Patent Owner further contends that the code shown in Figure 8.12 of Splash2 does not illustrate a “computational loop.” PO Resp. 81–87. Patent Owner first points to the if/else-if conditional statements in Figure 8.12, which select an execution path based on whether a condition is true or false. *Id.* at 82–83. Petitioner in its analysis, however, relies on the overall loop/endloop aspect of the code, not the if/else-if statements within that code. Pet. 26–28, 38–41. Petitioner first introduces the bidirectional array implementation and argues that Splash2 “discloses in Figure 8.7 the code executed by each processing element, . . . which *includes a loop for computing the edit distance* (i.e., the minimum cost of transforming one sequence to the other) between characters,” reproducing the code of Figure 8.7 (which includes the notations “loop” and “endloop”), then does the same for the unidirectional array code in Figure 8.12 (which also includes “loop” and “endloop”). *Id.* at 26, 28 (emphasis added).

What we must determine is how a person of ordinary skill in the art would have understood the code disclosed in Splash2, in particular “loop” and “endloop.” The parties and their respective declarants have different views.

Appendix D

See id. at 26–28, 38–41; PO Resp. 83–87; Reply 38–40. Patent Owner relies on Dr. Hodayoun, who testifies that “[t]o be defined as a loop, a loop exit condition must be specified. This is generally specified either in the first line of the code (begin loop and then condition) or the last line of [the] loop (loop end and then condition).” Ex. 2111 ¶ 200. According to Dr. Hodayoun, the code in Figure 8.12 lacks an “exit condition” at the beginning or end, so if the code was a loop as Petitioner contends, it “would run forever.” *Id.* ¶¶ 201–202. “This means nothing would be calculated by the code because the first piece of data would enter each PE and then run forever. Nothing would be passed to the next PE.” *Id.* ¶ 203. Therefore, “[t]he only reasonable interpretation of this pseudocode is to assume that something replaces the ‘loop-endloop’ syntax so that the pseudocode executes once then passes the data to the next PE.” *Id.* ¶ 204. In that case, the processing element executing the code “would never evaluate the same data more than once” and would not perform a “computational loop” under Patent Owner and Dr. Hodayoun’s claim interpretation. *Id.* ¶¶ 205–208.

Dr. Stone testifies that the loop/endloop code in Figure 8.12 represents “a loop for computing the edit distance (i.e., the minimum cost of transforming one sequence to the other) between characters.” Ex. 1003 ¶¶ 135, 139, 166–169. At a particular time step, the processing element “execut[es] the loop instructions disclosed in Figure 8.12” to compare the target character that was streamed in to the source character loaded in that processing element. *Id.* ¶¶ 139, 247–257. The processing element then repeatedly does the same for each input character of the target

Appendix D

sequence (i.e., “multiple comparisons”). *Id.* Responding to Dr. Homyoun’s testimony regarding the loop/endloop code,²¹ Dr. Stone points out that the repeated comparisons result in the code being executed a fixed number of times because the target and source sequences are of limited length (m and n , respectively) and each processing element executes the loop instructions until it completes computing distances along the diagonal matrix. Ex. 1076 ¶¶ 12–13.

Based on our review of the full trial record, we find that Dr. Stone’s reading, explained in detail in his declarations, is most consistent with the disclosure of Splash2 and that a person of ordinary skill in the art would have read the reference in the manner he describes. Figure 8.12 is labeled as a “loop,” indicating that the code within the “loop” and “endloop” designations is executed repeatedly by the processing element. Splash2 expressly describes the reason *why* it is executed repeatedly: the target sequence is “streamed through the array” so that a single incoming target character can be compared to the single source character loaded in that processing element. Ex. 1007, 103. The target sequence and source sequence both contain a large set of characters of “length m ” and “length n ,” respectively. *Id.* at 104. The target sequence can be any “length,” whereas the “maximum length” of the

21. Again, we do not agree with Patent Owner that Petitioner’s Reply and Dr. Stone’s reply testimony are improper, as they respond directly to Patent Owner’s arguments in the Response and Dr. Homyoun’s testimony regarding how a person of ordinary skill in the art would have understood the loop/endloop code in Splash2. *See* Sur-Reply 8–9; Reply 15–18, 38–40; Ex. 1076 ¶¶ 9–19; 37 C.F.R. § 42.23(b).

Appendix D

source sequence is limited by the “length” of the array. *Id.* Naturally, when the end of the m -length target sequence is reached, no more characters can be streamed in and the loop would complete for that processing element (i.e., after a fixed number of times executing the code). *See* Ex. 1076 ¶ 13. Thus, a person of ordinary skill in the art would have understood the code in Splash2 to loop and have an exit condition—it executes repeatedly (i.e., once for each target character streamed through the array) a fixed number of times (i.e., until the last target character in the target sequence is streamed through and compared).

Indeed, Splash2 expressly discloses limits on the number of times a processing element executes the code. In the unidirectional array implementation, “each PE computes the distances in one row of the distance matrix,” and “[a]t each time step, the PEs compute the distances along a single antidiagonal in the distance matrix,” indicating that processing would complete when there are no more distances to compute. *See* Ex. 1007, 104; *see also id.* at 100 (disclosing that each processing element “computes the distances along a particular diagonal of the distance matrix” in the bidirectional array implementation), 102 (“At the end of the computation, the resulting edit distance is transported out of the array on the distance streams.”). Dr. Stone’s understanding of the loop/endloop code is consistent with the text of Splash2 and the “loop” and “endloop” designations shown in the code itself. *See* Ex. 1076 ¶¶ 9–13.

Importantly, we also note that Splash2’s description of the unidirectional array implementation would disclose

Appendix D

the “computational loop” limitations even under Patent Owner’s proposed interpretation. Because the loop/endloop code is executed repeatedly for the same source character (e.g., source character *T* for Processing Element 1 and source character *C* for Processing Element 2 in annotated Figure 8.13 above), comparing it to each target character as it is streamed in, computations are executed repeatedly “per datum” (i.e., for the same source character “datum”) a fixed number of times. *See id.* ¶¶ 9–11, 14.

Patent Owner’s next argument is that Splash2 does not disclose “utilizing” the two “instantiated functional units” on the reconfigurable processor to operate on data dimensions forming “computational loops,” as recited in claim 1, because “even if Splash2 could be read to disclose two computational loops, it does not disclose instantiating those loops in the FPGAs of Splash2.” PO Resp. 87–90; Sur-Reply 17–19. According to Patent Owner, the external Sun workstation, not the FPGAs, would “handle any looping.” PO Resp. 88 (emphasis omitted). As support, Patent Owner relies on testimony from its declarants (Ex. 2111 ¶ 209; Ex. 2164 ¶¶ 42–43) and one paragraph from the thesis of Richard Peyton Halverson, Jr., Ph.D., “The Functional Memory Approach to the Design of Custom Computing Machines,” Ph.D. diss., University of Hawaii, 1994 (Ex. 2167, “the Halverson thesis”). PO Resp. 88.

We disagree. Splash2 expressly discloses that “[b]oth the bidirectional and unidirectional systolic arrays have been implemented *on the Splash 2 programmable logic array*, with versions for DNA and protein sequences.” Ex. 1007, 104 (emphasis added). “In the DNA version of

Appendix D

the unidirectional array, each of the 16 array FPGAs (X1 to X16) holds 14 PEs.” *Id.* at 107. Splash2’s discussion of the two implementations for the edit distance algorithm repeatedly refers to the functionality of an individual processing element (PE). *See, e.g., id.* at 100 (“There are several ways to map the edit distance computation onto a linear systolic array. We describe two such mappings.”), 104 (“At each time step, the PEs compute the distances along a single antidiagonal in the distance matrix The algorithm executed by each PE in the unidirectional array is listed in Figure 8.12.”). Thus, the processing elements that perform the looping are instantiated within the FPGAs in Splash2. *See* Ex. 1003 ¶¶ 238, 247–259; Ex. 1076 ¶¶ 18–19.

The Halverson thesis does not support a different reading. It discloses:

Splash 2 contains one or more boards each with an array of 16 well connected XILINX 4010 chips [Gokhale and Minnich, 1993]. The architecture does an excellent job supporting pipelined and SIMD processor configurations. Splash 2, for example, can be programmed in dbC, which is a superset of C used on other SIMD computers. The dbC preprocessor produces C that runs on the Sun and VHDL which define SIMD processors with an instruction set tailored to the application, one or more of which fit into each XILINX chip. When the actual program executes, looping is still handled in the Sun, which transmits SIMD instructions to the Splash 2 board(s).

Appendix D

Ex. 2167, 37–38. This isolated reference is to the Splash 2 system and how it “can be programmed in [Data-parallel Bit-serial C (dbC)].” *See id.* It does not mention or relate to the particular systolic array architectures and edit distance calculations described in Chapter 8 of Splash2. Further, whereas the paragraph above references the Splash 2 system being programmed in dbC, Splash2 describes programming the edit distance calculations in VHDL. *See id.* at 37–38, 208; Ex. 1007, 70, 106; Ex. 1074, 94–95, 97 (“Gokhale and Minnich” paper referenced in the Halverson thesis describing dbC and VHDL); Ex. 1003 ¶ 172; Ex. 1076 ¶¶ 15–17.

Dr. Stone conducted a detailed analysis of Splash2, citing specific figures and discussion in the reference, and explained why a person of ordinary skill in the art would have understood Splash2 to disclose utilizing two functional units to operate on different data dimensions of a calculation forming two “computational loops,” as recited in claim 1. After reviewing the full trial record, we credit that testimony. For the foregoing reasons, we find persuasive the arguments and evidence presented by Petitioner and Dr. Stone and, therefore, find that Splash2, in its description of the unidirectional array implementation, discloses the “computational loop” limitations of claim 1.

(3) Disputed Limitation: “Seamlessly”

Claim 1 recites that “said systolic implementation of said calculation enables said first computational loop and said second computational loop execute concurrently


Appendix D

and pass computed data seamlessly between said computational loops” (the “seamlessly” limitation).

Petitioner contends that the computational loops in Splash2 “execute concurrently” and “pass computed data seamlessly between” the computational loops because each processing element “execute[s] an instance of” the loop instructions, shown in Figure 8.12 above, and “each Processing Element directly communicates computed distance information to a neighbor.” Pet. 46 (citing Ex. 1003 ¶¶ 274–279). Dr. Stone explains that “[t]he output of a functional unit *passes directly* to the input of the next functional unit with *no intervening gaps or additional interfaces*.” Ex. 1003 ¶ 276 (emphases added); *see also id.* ¶¶ 186, 239 (stating that “distance information for each character comparison” is passed to “a neighboring Processing Element . . . without storage of that information in memory *between* processing elements” (emphasis added)). Each functional unit takes in “DSTin from the previous functional unit in the series” and outputs “DSTout for the next functional unit in the series.” *Id.* ¶ 276 (citing Ex. 1007, 104, Fig. 8.10).

Appendix D

Dr. Stone provides the following annotated version of a portion of Figure 8.13 of Splash2 (*id.* ¶ 277).



The annotated portion above of a trace of the unidirectional array depicts “the seamless communication of the DST data in the systolic calculation as a sequence of red arrows.” *Id.* Dr. Stone explains that

[t]he red arrows indicate . . . that the computed output DSTout from a functional unit is *directly connected* to the next functional unit input DSTin. Similarly, all output data from one

Appendix D

functional unit are seamlessly connected to the corresponding inputs of the functional unit in the sequence. There are *no intermediate interfaces between modules* to translate the output of one module into a form where it can be used as input data to the next module. Such interfaces, if they were to exist, would be “seams” in the communication links between adjacent functional units.

Id. (emphases added). As explained above, we interpret “pass computed data seamlessly between said computational loops” to mean “communicate computed data directly between functional units that are calculating computational loops.” *See supra* Section II.C.3. Petitioner’s analysis for the “seamlessly” limitation of claim 1, consistent with that interpretation and supported by the testimony of Dr. Stone, which we credit, is persuasive. *See* Ex. 1003 ¶¶ 272–277, 279.

Patent Owner makes four arguments in its Response. First, Patent Owner contends that in passing computed data between the computational loops operating on data dimensions of the calculation, claim 1 “specifically require[s] a sequence of computations *per datum*,” but the loop/endloop code in Splash2 “merely selects an execution path for the processor to execute once per datum.” PO Resp. 96–97. We disagree for the reasons stated above. *See supra* Sections II.C.2, II.E.2.a.2.

Second, Patent Owner argues that “the FPGAs [in Splash2] must communicate with the Sun workstation

Appendix D

(which is handling any looping) through the Sbus” and the “boundary between the FPGAs and the workstation . . . clearly constitutes a ‘seam,’” citing the same declarant testimony and Halverson thesis discussed above in connection with the “computational loop” limitations. PO Resp. 96–97, 100. Again, we disagree that the Sun workstation handles the looping discussed in Chapter 8 for the edit distance calculations. *See supra* Section II.E.2.a.2.

Third, Patent Owner asserts that “Splash2 is, at best, ambiguous on whether memory is used to store the results from each processing element after each time step to preserve it for output and later use.” PO Resp. 97–100 (citing Ex. 2111 ¶¶ 210–219). Patent Owner contends that “storage is likely necessary to preserve the values calculated at each timestep,” where the computed data is overwritten at each time step and the overall edit distance is based on the repeated calculations and overwriting (rather than just the final time step). *Id.* at 98; *see* Ex. 2111 ¶ 210 (identifying, for example, SDout, TDout, and PEDist from the bidirectional array implementation). Patent Owner points out that Splash2 has “local memory at each FPGA for storage purposes,” which “can be used for storage of results.” PO Resp. 98–99 (citing Ex. 1007, 88, 95, 102 n.3; Ex. 1035, 5; Ex. 2156, 205–206). According to Patent Owner, “it is equally (if not more) plausible for a [person of ordinary skill in the art] to interpret Splash2 to use the local memory due to the known timing problems in systolic systems.” *Id.* at 99–100.

We are not persuaded that Splash2 is ambiguous in the manner Patent Owner contends. The disclosure and

Appendix D

figures discussed above indicate that computed distance information is communicated directly between processing elements. In the context of the unidirectional array implementation, for example, a processing element takes in DSTin, executes the loop/endloop code to compute DSTout, and outputs DSTout to the next processing element in the series. *See* Ex. 1007, 104 (Fig. 8.10 depicting the inputs and outputs for the unidirectional array PE), 105 (Fig. 8.12 showing how DSTout is calculated), 106 (Fig. 8.13 depicting direct connections between processing elements); Ex. 1003 ¶¶ 275–279; *see also* Ex. 1007, 101–102 (Figs. 8.6–8.8 showing inputs of SDin and TDin and output of the computed PEDist as SDout and TDout in either direction for the bidirectional array implementation). Claim 1 recites that the systolic implementation enables the computational loops to “pass computed data seamlessly between said computational loops.” There is no indication in Splash2 itself that there is any memory or other intervening structure *between* the disclosed processing elements. That data may be stored temporarily *within* a processing element (e.g., in a storage register) does not automatically mean that memory between the processing element and another processing element is used for communicating that data. *See* Ex. 1007, 104 (stating that the unidirectional array PE “stores two distances, DST and PDST,” but never stating that memory between the PE and another PE is used to do so).

Fourth, Patent Owner argues that Splash2 does not disclose the “seamlessly” limitation even under Petitioner’s proposed interpretation because “the Xilinx FPGAs . . . clearly contain structure (such as the buffered switch

Appendix D

matrix) within the internal routing resources to connect processing elements.” PO Resp. 100–105 (citing Ex. 1035, 28–31; Ex. 2078, 19–29, 32–34, 37–41, 46–51, 59–65, 70). Again, just because an FPGA has memory does not mean that when functional units are instantiated within the FPGA, memory is necessarily placed *between* functional units or used to transfer data from one functional unit to another. *See* Reply 25, 47; *supra* Section II.C.3.

For the foregoing reasons, we find persuasive the arguments and evidence presented by Petitioner and Dr. Stone and, therefore, find that Splash2, in its description of the unidirectional array implementation, discloses the “seamlessly” limitation of claim 1.

b) Bidirectional Array Implementation

Petitioner’s explanation for why a person of ordinary skill in the art would have understood Splash2’s description of the bidirectional array implementation to disclose the limitations of claim 1 is similar to its explanation for the unidirectional array implementation. Petitioner argues that the calculation is “systolically implemented” in the bidirectional array implementation because “each Processing Element ‘computes the distances along a particular diagonal of the distance matrix,’” where the calculation is based on the source and target sequences streamed in from the processing element’s “upstream and downstream neighbors” and “[t]he resulting distance information for each character comparison is continually passed to the upstream and downstream neighboring Processing Elements, without storage of that information in memory between processing elements.” Pet. 32.

Appendix D

With respect to the “computational loop” limitations, Petitioner argues that a person of ordinary skill in the art would have understood that the loop/endloop code shown in Figure 8.7 above “constitutes a time step of a ‘data dimension’ comprising multiple time steps because each such instance of code executes in a single time step, and the systolic computation comprises multiple time steps.” *Id.* at 38 (citing Ex. 1003 ¶¶ 245–246). Petitioner provides the following annotated version of Figure 8.8 of Splash2 (*id.* at 43).

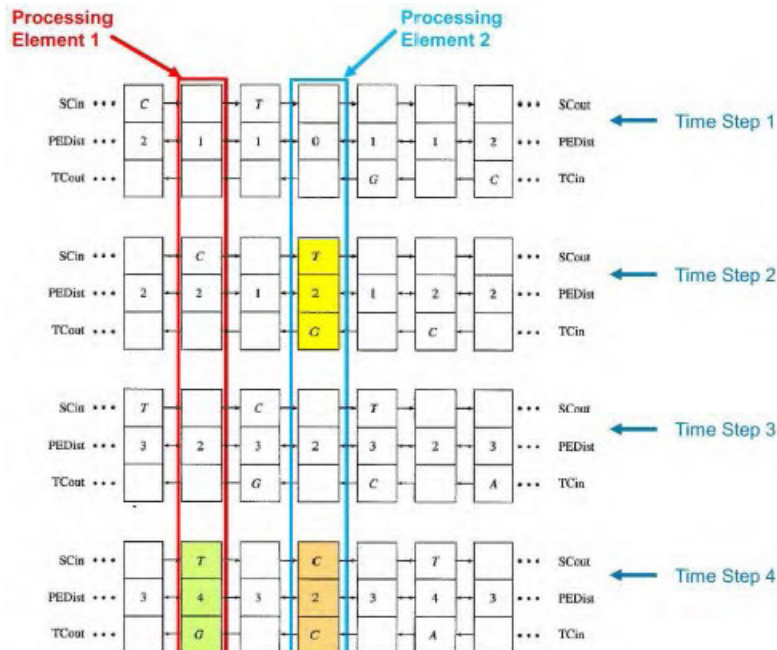


FIGURE B.B Trace of Bidirectional Array When Comparing the Sequences TCTAGACC and GCATAAGC

Appendix D

Annotated Figure 8.8 depicts the characters of the source sequence *TCTAGACC* streamed left to right through the processing elements (row SCin), and the characters of the target sequence *GCATAAGC* streamed right to left through the processing elements (row TCin). *Id.* at 42–46 (citing Ex. 1003 ¶¶ 260–271). For example, at Time Step 2, Processing Element 2 compares the first target character *G* to the first source character *T* (shown in yellow). *Id.* at 43. At Time Step 4, (1) Processing Element 2 compares the second target character *C* to the second source character *C* (shown in orange), and (2) Processing Element 1, which has received the first target character *G* streamed through the array, compares it to the third source character *T* (shown in green). *Id.*

Petitioner argues that, as shown in the example above, Processing Element 2 operates on a “subsequent data dimension” (i.e., comparing streamed target and source characters on a particular diagonal of the distance matrix) forming a “first computational loop” (i.e., “Processing Element 2 executing the loop instructions disclosed in Figure 8.7”); and Processing Element 1 operates on a “previous data dimension” (i.e., comparing streamed target and source characters on a particular diagonal of the distance matrix) forming a “second computational loop” (i.e., “Processing Element 1 executing the loop instructions disclosed in Figure 8.7”). *Id.* at 44–45 (emphases omitted). As with the unidirectional array implementation, “the loop instructions executed in each Processing Element are first and second computational loops because they are necessarily different instances of those loop instructions.” *Id.* at 45.

Appendix D

With respect to the “seamlessly” limitation, Petitioner argues that each processing element “execute[s] an instance of” the loop instructions shown in Figure 8.7 above and “each Processing Element directly communicates computed distance information to a neighbor.” *Id.* at 46 (citing Ex. 1003 ¶¶ 274–279). Dr. Stone explains that the same passing of data occurs in the unidirectional and bidirectional array implementations, where “[t]he output of a functional unit *passes directly* to the input of the next functional unit with *no intervening gaps or additional interfaces.*” Ex. 1003 ¶¶ 276, 278 (emphases added). Dr. Stone provides the following annotated version of a portion of Figure 8.8 of Splash2 (*id.* ¶ 278).




FIGURE 8.8 Trace of Bidirectional Array When Comparing the Sequences *TCTAGACC* and *GCATAAGC*

The annotated portion above of a trace of the bidirectional array depicts the seamless communication of distance information. *Id.* Dr. Stone explains that

[t]he blue arrows indicate that the computed output PEDist from a functional unit is directly connected to the next functional unit input PEDist in both directions. The code for the

Appendix D

Bidirectional implementation discloses that the computed PEDist is conveyed to the left on output TDout, and to the right on output SDout.

Id. (citing Ex. 1007, 101, Fig. 8.6).

Patent Owner argues in its Response that Splash2 does not disclose the “computational loop” and “seamlessly” limitations of claim 1. PO Resp. 80–90, 96–105. Patent Owner’s arguments apply to both the unidirectional and bidirectional array implementations described in Splash2, and we disagree with those arguments for the reasons explained above.²² *See id.*; *supra* Section II.E.2.a. Petitioner’s analysis of how Splash2’s description of the bidirectional array implementation discloses the limitations of claim 1, supported by the testimony of Dr. Stone, which we credit, is persuasive. *See* Ex. 1003 ¶¶ 153–179, 184–193, 203–208, 210–216, 220–233, 235–246, 260–276, 278–279. For the reasons stated above, we find persuasive the arguments and evidence presented by Petitioner and Dr. Stone and, therefore, find that Splash2, in its description of the bidirectional array implementation, discloses the limitations of claim 1.

22. The loop/endloop code executed by a processing element operates similarly in both implementations. We note that although the unidirectional array implementation would meet the “computational loop” limitation even under Patent Owner’s proposed interpretation (because the code is executed repeatedly for the same source character), the bidirectional array implementation would not (because the target and source sequences are both streamed through).

*Appendix D***c) Conclusion**

We are persuaded by Petitioner's explanations and supporting evidence regarding both the unidirectional and bidirectional implementations of Splash2, and find that Splash2 discloses every limitation of claim 1 under both of Petitioner's theories. Petitioner has proven, by a preponderance of the evidence, that claim 1 is anticipated by Splash2 under 35 U.S.C. §§ 102(a) and 102(b).

3. Claims 18, 21, and 22

Dependent claims 18, 21, and 22 limit the "calculation" of parent claim 1 to a particular type. Claim 18 recites that the calculation comprises a "search algorithm for data mining," claim 21 recites that it comprises a "genetic pattern matching function," and claim 22 recites that it comprises a "protein folding function." Petitioner argues that a person of ordinary skill in the art would have understood Splash2's edit distance calculation to be each type of calculation. -1603 Pet. 54-57. For example, Petitioner argues that the edit distance calculation is a "genetic pattern matching function" because it "compares characters of two genetic sequences." *Id.* at 55; *see* Ex. 1007, 100-104; Ex. 1003 ¶¶ 126-140, 462-465. Patent Owner does not argue separately dependent claims 18, 21, and 22 in its Response, only disputing Petitioner's contentions with respect to parent claim 1. PO Resp. 77-90, 96-105. We have reviewed Petitioner's contentions regarding claims 18, 21, and 22, which are consistent with the disclosure of Splash2 and supported by the testimony of Dr. Stone, and are persuaded that Petitioner has proven,

Appendix D

by a preponderance of the evidence, that dependent claims 18, 21, and 22 are anticipated by Splash2 under 35 U.S.C. §§ 102(a) and 102(b). *See* -1603 Pet. 54–57.

4. Claim 15

Claim 15 recites that the “instantiating” step of parent claim 1 includes “establishing a stream communication connection between functional units.” Petitioner argues that Splash2 discloses establishing such a connection “between adjacent Processing Elements in both the unidirectional and bidirectional edit distance calculation because it discloses streaming target and source characters between processing elements both unidirectionally and bidirectionally.” -1602 Pet. 52–53. Patent Owner responds that Splash2 does not disclose the limitation because it does not disclose a “queue” between processing elements or “signaling” for the processing elements to interact with such a queue. PO Resp. 77–80. Patent Owner’s arguments are premised on its proposed interpretation of “stream communication” as “a data path that acts like a queue connecting via the reconfigurable routing resources a producer and a consumer of data that operate concurrently.” *Id.* at 50, 77–80. As explained above, we disagree and instead interpret “establishing a stream communication connection between functional units” to mean “establishing a connection over which data is streamed between functional units.” *See supra* Section II.C.4.

Given our interpretation, Petitioner has made a sufficient showing that Splash2 discloses the limitation.

Appendix D

In the unidirectional array implementation, “[t]he source sequence is loaded once” and “target sequences are streamed through the array one at a time, separated by control characters.” Ex. 1007, 103. “[D]ata flows through the unidirectional array in one direction.” *Id.* In the bidirectional array implementation, “[t]he source and target sequences are streamed through the array in opposite directions.” *Id.* at 101–102. “[T]here is one distance stream associated with each character stream. At each step, the contents of the streams represent the characters to be compared and the distances along one of the antidiagonals of the distance matrix.” *Id.* at 102. Thus, Splash2 discloses in the various implementations establishing a connection over which data (e.g., characters of the source and target sequences, computed distance information) is streamed between functional units. Petitioner has proven, by a preponderance of the evidence, that dependent claim 15 is anticipated by Splash2 under 35 U.S.C. §§ 102(a) and 102(b).

**F. Obviousness Ground Based on Splash2
(Claims 1, 15, 18, 21, and 22)**

1. Obviousness Contentions

Petitioner argues, in the alternative, that claims 1, 15, 18, 21, and 22 would have been obvious over Splash2. Pet. 47–52; -1602 Pet. 53–59; -1603 Pet. 57–63. Petitioner argues that “[t]o the extent one might argue” that the chapters of Splash2 cannot be considered together for purposes of anticipation, a person of ordinary skill in the art would have found claim 1 to be obvious based on the

Appendix D

full disclosure of Splash2. Pet. 47–48. Petitioner further contends that “[t]o the extent one might argue” that Splash2 does not disclose four limitations of claim 1 (i.e., the “transforming” step, portions of the “instantiating” step, the limitation that “only functional units needed to solve the calculation are instantiated,” and the limitation that “each instantiated functional unit at the at least one reconfigurable processor interconnects with each other instantiated functional unit”), a person of ordinary skill in the art would have been motivated to modify Splash2 to include them. *Id.* at 48–52. For example, Petitioner argues that if Splash2 does not disclose the limitation that “only functional units needed to solve the calculation are instantiated,” doing so would have been obvious because the FPGAs in Splash2 are “cleared of any configuration data before a new configuration can be entered,” the arrays are customized for a particular application, and a person of ordinary skill in the art would have been motivated to instantiate only the processing elements necessary to solve a calculation to “avoid the additional work” of writing code for processing elements that “performed no function [and] served no purpose.” *Id.* at 50–51. Splash2 is a single reference, and Patent Owner does not dispute that the chapters of Splash2 can be considered together. *See* Ex. 1007, 104 (Chapter 8 indicating that “the bidirectional and unidirectional systolic arrays have been implemented on the Splash 2 programmable logic array,” which was described in earlier chapters). Nor does Patent Owner dispute that Splash2 discloses the four identified limitations. We find that the four limitations are disclosed by Splash2. *See supra* Section II.E.2.a.1. Thus, we need not address the alternative arguments made by Petitioner.

Appendix D

Patent Owner argues that Splash2 fails to disclose certain other limitations of claims 1 and 15. *See supra* Sections II.E.2, II.E.4. We disagree, find that Splash2 expressly teaches all limitations of claims 1, 15, 18, 21, and 22, and conclude that Petitioner has proven anticipation by a preponderance of the evidence. *See id.* Accordingly, absent a persuasive showing of secondary considerations, which we discuss below, the claims also would have been obvious based on Splash2. *See Realtime Data, LLC v. Iancu*, 912 F.3d 1368, 1373 (Fed. Cir. 2019) (“[I]t is well settled that a disclosure that anticipates under § 102 also renders the claim invalid under § 103, for anticipation is the epitome of obviousness.” (citations and internal quotation marks omitted)).

2. Secondary Considerations of Nonobviousness

Patent Owner also argues that secondary considerations of nonobviousness demonstrate that claims 1–5, 7–9, 15, 17, 18, and 20–24 would not have been obvious to a person of ordinary skill in the art. PO Resp. 121–125. Although some of the challenged claims are anticipated by—and, therefore, would have been obvious based on—Splash2, because the parties refer to the challenged claims collectively in their arguments regarding secondary considerations, we do so as well and now address those arguments as applied to all of the claims challenged as obvious. *See id.*; Reply 52–55.

“In order to accord substantial weight to secondary considerations in an obviousness analysis, the evidence of secondary considerations must have a nexus to the claims, *i.e.*, there must be a legally and factually sufficient

Appendix D

connection between the evidence and the patented invention.” *Fox Factory, Inc. v. SRAM, LLC*, 944 F.3d 1366, 1373 (Fed. Cir. 2019) (citations and internal quotation marks omitted). “The patentee bears the burden of showing that a nexus exists.” *Id.* “To determine whether the patentee has met that burden, we consider the correspondence between the objective evidence and the claim scope.” *Id.* A patentee is entitled to a rebuttable presumption of nexus “when the patentee shows that the asserted objective evidence is tied to a specific product and that product ‘embodies the claimed features, and is coextensive with them.’” *Id.* (citation omitted). However, “[a] finding that a presumption of nexus is inappropriate does not end the inquiry into secondary considerations.” *Id.* “To the contrary, the patent owner is still afforded an opportunity to prove nexus by showing that the evidence of secondary considerations is the ‘direct result of the unique characteristics of the claimed invention.’” *Id.* at 1373–1374 (citation omitted).

“Where the offered secondary consideration actually results from something other than what is both claimed and *novel* in the claim, there is no nexus to the merits of the claimed invention,” meaning that “there must be a nexus to some aspect of the claim not already in the prior art.” *In re Kao*, 639 F.3d 1057, 1068–69 (Fed. Cir. 2011). On the other hand, there is no requirement that “objective evidence must be tied exclusively to claim elements that are not disclosed in a particular prior art reference in order for that evidence to carry substantial weight.” *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1331 (Fed. Cir. 2016). A patent owner may show, for example, “that it is the

Appendix D

claimed combination as a whole that serves as a nexus for the objective evidence; proof of nexus is not limited to only when objective evidence is tied to the supposedly ‘new’ feature(s).” *Id.* at 1330. Ultimately, the fact finder must weigh the secondary considerations evidence presented in the context of whether the claimed invention as a whole would have been obvious to an ordinarily skilled artisan. *Id.* at 1331–32.

As objective evidence of nonobviousness, Patent Owner cites the declarations of Dr. Homayoun (Ex. 2111), Mr. Huppenthal (Ex. 2100), and Dr. El-Ghazawi (Ex. 2164); the deposition transcript of Stephen M. Trimberger, Ph.D. in Cases IPR2018-01599 and IPR2018-01600 (Ex. 2076); the deposition transcript of Dr. Stone in Case IPR2018-01594 (Ex. 2066); and two publications authored by, among others, Dr. El-Ghazawi and one of the authors of Splash2, Duncan Buell, Ph.D.: Duncan Buell *et al.*, “High-Performance Reconfigurable Computing,” *IEEE Computer Society*, pp. 23–27 (Mar. 2007) (Ex. 2166), and Tarek El-Ghazawi *et al.*, “The Promise of High-Performance Reconfigurable Computing,” *IEEE Computer Society*, pp. 69–76 (Feb. 2008) (Ex. 2165). PO Resp. 121–125.

Patent Owner asserts that there was commercial success and praise by others for “DirectStream’s patented products (SRC-6, SRC-6e),” which involved “instantiating the reconfigurable resources necessary for an algorithm so as to maximize the speed data can be passed between different looping portions of the program” running concurrently “to achieve a particular

Appendix D

systolic system that processes data without the associated delays of communications protocols or the Von Neumann bottleneck.” *Id.* at 14–15, 125. Patent Owner, however, does not provide any explanation or analysis demonstrating that its “SRC-6” or “SRC-6e” products were used to perform the methods recited in any of the challenged claims, or are coextensive with the claimed methods. Mr. Huppenthal discusses “SRC-6 products and the SRC-7 . . . under development” and testifies that “[p]roduction systems would incorporate . . . AI and Machine Learning as described in [the ’324 patent]” and “Army and Air Force applications would make extensive use of program loops and streams as defined in [the ’324 patent].” Ex. 2100 ¶¶ 80, 83. Importantly, though, Mr. Huppenthal never refers to the claims in his testimony and acknowledged that he did not perform “any kind of comparison of the claims [of the ’324 patent] to any particular system.” *See id.* ¶¶ 80–87; Ex. 1073, 106:14–107:14, 108:17–109:5; PO Opp. 2 (acknowledging that Mr. Huppenthal only provides “high-level, general testimony regarding . . . specific hardware features of interest” to “customers who purchased SRC-6 computers” and “does not specifically discuss any applications any customers ran on the SRC-6”). Therefore, we find that a presumption of nexus is inappropriate.

Nor does the other evidence cited by Patent Owner establish a nexus between DirectStream’s products and any of the challenged claims. Patent Owner cites one of the publications and Dr. Homayoun’s declaration mentioning the “SRC-6,” “SRC-6E,” and “SRC-7” products, but does not point to anything in either document explaining how

Appendix D

the products embody the claims. *See* PO Resp. 125 (citing Ex. 2165, 2–7, Fig. 4; Ex. 2111 ¶¶ 108–109, 118–119). Absent some explanation of how the cited products worked, there is no basis to say that they were used to practice the methods recited in the challenged claims. Accordingly, we find that Patent Owner has not established a sufficient nexus between the claims and the alleged commercial success and industry praise.

Patent Owner further asserts that there was a long-felt but unmet need for “very fast processing of large volumes of data, and improvements in conventional systems focused on reducing the delays associated with chip-to-chip communications protocols and also the Von Neumann processing bottleneck.” *Id.* at 123 (citing Ex. 2164 ¶¶ 17–25; Ex. 2166, 3–5; Ex. 2100 ¶¶ 36–41; Ex. 2111 ¶¶ 40–106). According to Patent Owner, there was skepticism in the industry that “reconfigurable processors could adequately satisfy these needs compared to other well-established solutions in the same technology space.” *Id.* (citing Ex. 2076, 129:24–130:20; Ex. 2066, 168:9–169:4, 179:6–13, 197:8–11).

To support a conclusion of nonobviousness, an alleged long-felt need must have been a persistent one that was recognized by those of ordinary skill in the art, must not have been satisfied by another before the challenged patent, and must have been satisfied by the claimed invention. *Perfect Web Techs., Inc. v. InfoUSA, Inc.*, 587 F.3d 1324, 1332–33 (Fed. Cir. 2009). Similarly, evidence showing that “skilled artisans were initially skeptical about the [claimed] invention” is relevant to

Appendix D

nonobviousness. *Metabolite Labs., Inc. v. Laboratory Corp. of Am. Holdings*, 370 F.3d 1354, 1368 (Fed. Cir. 2004). We have reviewed all of the cited evidence and do not find it persuasive, as Patent Owner does not provide any explanation establishing a nexus to the challenged claims. Patent Owner does not explain in its Response, for example, how exactly the claimed methods solved the identified problems of speed and reducing communication delays. Nor does Patent Owner explain how the cited testimony from Dr. Trimberger and Dr. Stone allegedly shows skepticism in the industry that the reconfigurable processor approach *of the challenged claims* would reduce such issues, as opposed to mere views on reconfigurable processors in general. *See* PO Resp. 123.

Patent Owner also contends that others in the industry failed to solve the problems allegedly solved by the '324 patent. PO Resp. 124–125 (citing Ex. 2167, 12–20; Ex. 2174, 4; Ex. 2164 ¶¶ 32–43; Ex. 2100 ¶¶ 30–31, 36–41; Ex. 2111 ¶¶ 179–184; Ex. 2066, 168:9–169:4, 179:6–13, 197:8–11; Ex. 2076, 129:24–130:20). We are not persuaded for similar reasons. Patent Owner does not explain in sufficient detail how the cited evidence demonstrates a “failure of others to do *that which the patent claims*.” *See id.* at 124 (quoting *Para-Ordnance Mfg., Inc. v. SGS Importers Int'l, Inc.*, 73 F.3d 1085, 1094 (Fed. Cir. 1995)) (emphasis added). Without further explanation, we are not persuaded that Patent Owner’s evidence of the failure of others supports the nonobviousness of the challenged claims.

Finally, Patent Owner points to the publications authored by Dr. Buell and Dr. El-Ghazawi in 2007 and

Appendix D

2008 (well after the effective filing date of the '324 patent, October 31, 2002) that “survey[ed] . . . the state of the art” and discussed “DirectStream’s SRC-6 and SRC-6e systems” but “did not include [the Splash 2 system] as an example of a viable [High-Performance Computing Technology (HPRC)] system.” *Id.* at 124–125 (citing Exs. 2165, 2166). Patent Owner does not point to any authority—and we are not aware of any—for the proposition that a reference’s *silence* as to a particular implementation indicates a failure of that implementation to solve a particular problem. *See* Tr. 66:1–67:11. Because the cited documents do not mention the Splash 2 system, we find that they do not support Patent Owner’s contention that the system represented a failure of others.

For the reasons explained above, we conclude that Patent Owner’s evidence purportedly showing commercial success, industry praise, long-felt need, skepticism in the industry, and failure of others does not weigh in favor of nonobviousness of the challenged claims.

3. Conclusion

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, we determine that claims 1, 15, 18, 21, and 22 would have been obvious based on Splash2 under 35 U.S.C. § 103(a).

*Appendix D***G. Obviousness Ground Based on Splash2 and RaPiD (Claims 8 and 9)**

Petitioner asserts that dependent claims 8 and 9 are unpatentable over Splash2 and RaPiD. Pet. 55–67; *see* Ex. 1003 ¶¶ 330–369. Claim 8 recites that the calculation comprises “a JPEG image compression calculation” and claim 9 recites that the calculation comprises “an MPEG image compression calculation.” With respect to claim 8, for example, Petitioner argues that RaPiD teaches a discrete cosine transform (DCT) “implemented on a systolic array of a reconfigurable computing system called the Reconfigurable Pipelined Datapath, or ‘RaPiD.’” Pet. 56. RaPiD discloses that “[t]he datapaths constructed in RaPiD are linear arrays of functional units communicating in mostly nearest-neighbor fashion” and “[s]ystolic algorithms, for example, map very well into RaPiD datapaths.” *Id.* (quoting Ex. 1009, 106). RaPiD explains that one application implemented on the disclosed system was a “2-D DCT . . . used in JPEG/MPEG data compression to convert an image from the spatial domain to the frequency domain.” *Id.* at 56–57 (quoting Ex. 1009, 110).


Petitioner asserts that a person of ordinary skill in the art would have combined the teachings of Splash2 and RaPiD such that “the functionality of the cell of RaPiD’s Figure 10 would be instantiated into Processing Elements of the Splash2 Unidirectional Systolic Array,” where “[t]he Processing Elements would perform the matrix multiply calculations described in RaPiD, and then pass the resulting data to the next Processing Element in a systolic fashion.” *Id.* at 58–59, 64–65. Petitioner provides

Appendix D

numerous reasons why a person of ordinary skill in the art would have been motivated to do so. *Id.* at 64–67.

Patent Owner does not dispute that RaPiD teaches JPEG and MPEG image compression calculations, as recited in claims 8 and 9, but makes other arguments. First, Patent Owner argues that the cited calculations in RaPiD do not constitute a “computational loop,” as recited in parent claim 1. PO Resp. 93–95.

Petitioner asserts in the Petition that each processing element (“cell”) in the RaPiD system carries out a “computational loop” in performing the 2-D DCT calculation, citing the following annotated version of Figure 10 of RaPiD provided by Dr. Stone (Pet. 63; Ex. 1003 ¶ 357).



*Figure 10: Netlist for one cell of 2-D DCT. The top pipelined bus streams in the **A** matrix while the bottom bus streams out resulting 1-D DCT, transposed. The top bus also streams the **W** columns into the local memories prior to the computation.*

Annotated Figure 10 depicts a netlist for one cell of the 2-D DCT. Dr. Stone provides a detailed explanation of how RaPiD performs its calculations. Ex. 1003 ¶¶ 333–360.

Appendix D

In relevant part for purposes of addressing Figure 10, “[a] 2-D DCT can be decomposed into two sequential 1-D DCTs.” *Id.* ¶ 333 (quoting Ex. 1009, 110). The 1-D DCT can be carried out in “cells” of an 8-cell pipeline by performing matrix multiply calculations on two matrices: input vector A and weight W. *Id.* ¶¶ 334–336. RaPiD discloses that the 2-D DCT can be reduced to the following dot product equations (Ex. 1009, 111).

$$z_{mj} = \sum_{n=0}^{N-1} a_{mn} w_{nj}, \quad (5)$$

and thus

$$y_{ji} = \sum_{m=0}^{N-1} z_{mj} w_{mi} \quad (6)$$

As illustrated in the equations above, “both z_{mj} and y_{ji} are equivalent to . . . matrix multiplies,” but because “the z_{mj} values are produced in row-major order but required in column-major order, the results from the z_{mj} DCT must be transposed prior to computing y_{ji} ,” *Id.* Figure 8 of RaPiD is reproduced below.




Figure 8: 2-D $N \times N$ DCT

Figure 8 depicts the 2-D DCT as two 1-D DCT calculations with a transpose in between. Dr. Stone explains in detail

Appendix D

how the calculation of the dot product equations above requires multiple iterations of (1) multiplications between elements of a first matrix (a or z) with elements of a second matrix (w), and (2) addition of the result of that multiplication to a running sum from previous iterations. Ex. 1003 ¶¶ 351–359; Ex. 1076 ¶¶ 21–22.²³ According to Dr. Stone, the portion of annotated Figure 10 above highlighted in yellow shows “the output [of] the [arithmetic logic unit (ALU) being] looped back to the ALU input.” Ex. 1003 ¶ 357. That path is what “loops the intermediate, or running, sum from the output of the ALU to its input so that it may be used in the next iteration of the DCT calculation.” Ex. 1076 ¶ 25.

With respect to Figure 10, the 2-D DCT partitions an $M \times N$ image into 8×8 sub-images and performs two

23. Again, we do not agree with Patent Owner that Petitioner’s Reply and Dr. Stone’s reply testimony are improper, as they respond directly to Patent Owner’s argument in the Response that RaPiD teaches a “bypass or forwarding path” rather than a “computational loop.” See Sur-Reply 8–9; Reply 40–44; Ex. 1076 ¶¶ 20–27; 37 C.F.R. § 42.23(b). Petitioner did not change its theory as to how RaPiD teaches a “computational loop,” relying in both its Petition and in its Reply on the portion of Figure 10 highlighted in yellow above. See *Apple Inc. v. Andrea Elecs. Corp.*, 949 F.3d 697, 706–707 (Fed. Cir. 2020) (concluding that a reply was proper where it did not change the “legal ground” asserted in the petition and “relie[d] on the same [teaching] from the same prior art reference to support the same legal argument”); *Chamberlain Grp., Inc. v. One World Techs., Inc.*, 944 F.3d 919, 925 (Fed. Cir. 2019) (“Parties are not barred from elaborating on their arguments on issues previously raised.”). Patent Owner also had the opportunity to cross-examine Dr. Stone about his reply declaration. See Ex. 2176.

Appendix D

1-D DCTs for each, using an 8 x 8 weight matrix, with the transpose in between. Ex. 1003 ¶¶ 337–341 (citing Ex. 1009, 111, Figs. 8–10). RaPiD discloses that

[s]ince a 2-D DCT performs two multiplies by the same weight matrix, W is loaded only once: one column per cell in both the first 8 cells and last 8 cells. The transpose in between matrix multiplies is performed with two local memories per cell: one to store products of the current sub-image and the other to store the products of the *previous* sub-image. During the computation of the current sub-image, the transpose of the previous sub-image computation is passed to the next 8 cells. The datapath for one RaPiD cell of a 2-D DCT is shown in Figure 10.

Ex. 1009, 111. Thus, as shown in Figure 10, “each cell in the array stores one column of weight data and receives both a stream of row image data and a stream of 1-D DCT data. The cell calculates the DCT for the current sub-image and passes on the DCT data for the previous sub-image along with row image data directly to the next cell.” Ex. 1003 ¶ 343; *see* Ex. 1009, 111, Fig. 10 (depicting the “[c]olumn of matrix W ” stored in RAM and the “[r]ow of matrix A ” and “[c]olumn of 1-D DCT results” received from the previous cell).

Patent Owner responds that the portion of annotated Figure 10 above highlighted in yellow is not a “computational loop” but rather a “bypass or forwarding path,” which “support[s] back to back execution of operations without

Appendix D

stall, by forwarding (or bypassing) the output of an ALU to an input of the same or other ALU.” PO Resp. 94–95 (citing Ex. 2029 ¶¶ 41–48; Ex. 2111 ¶¶ 234–241; Ex. 2043, 301; Ex. 2044, 2, 6–7). According to Patent Owner, Dr. Stone also conceded during cross-examination that Figure 10 shows a bypass or forwarding path. *Id.* at 95 (citing Ex. 2064, 201:21–202:1).

Based on our review of the full record after trial, we find Patent Owner’s arguments unavailing. Dr. Stone’s explanation for how the calculation of the dot product equations requires adding products to a running sum is consistent with the figures and text of RaPiD, and we see no structure in Figure 10 other than the portion highlighted in yellow that would be capable of performing such functionality. Nor does Patent Owner point to any. RaPiD teaches, via the looping of the ALU output in Figure 10, a set of computations that is executed repeatedly a fixed number of times (i.e., until the end of the matrices is reached). *See supra* Section II.C.2.

We also are not persuaded that Dr. Stone conceded otherwise during cross-examination. He was asked simply whether “Figure 10 show[s] a bypass path or a forwarding path” and responded “Yes.” Ex. 2064, 201:21–202:2 (emphasis added). He was not asked to—and did not—identify any particular structure shown in Figure 10. Nor was he asked specifically whether the yellow highlighted portion is a bypass or forwarding path. Moreover, there is a bypass or forwarding path shown in Figure 10 that is not the yellow highlighted portion, namely the bottom wire labeled “Column of 1-D DCT results flows out” where

Appendix D

DCT data from the previous cell is passed along to the next cell, bypassing the ALU functionality and memory shown in the top portion of the figure. Ex. 1003 ¶ 343; Ex. 1076 ¶ 27. Dr. Stone states that that wire was what he was referring to during cross-examination. Ex. 1076 ¶ 27. We see no contradiction between Dr. Stone’s declaration testimony and his statement that a bypass or forwarding path exists in Figure 10.

Second, Patent Owner argues that RaPiD does not teach passing computed data “seamlessly” between computational loops because Figure 10 “shows storage of results in memory (RAM) before being passed onto the next cell.” PO Resp. 105–106. As an initial matter, we note that Petitioner’s position, supported by the testimony of Dr. Stone, is that a person of ordinary skill in the art would have had reason to combine the references’ teachings such that “each Processing Element of the Splash2 Unidirectional Systolic Array would carry out the calculations of RaPiD Figure 10” and computed data would be “communicated directly from one Processing Element to the next.” Pet. 64 (emphases omitted); *see* Ex. 1003 ¶¶ 344–346, 349, 351, 360, 363. We find, for the reasons explained above, that in the unidirectional array implementation of Splash2, computed data is passed “seamlessly” (i.e., communicated directly, without intervening memory or other structures) between processing elements performing the computational loops. *See supra* Section II.E.2.a.3. Thus, Patent Owner’s argument attacking RaPiD individually with respect to the “seamlessly” limitation is not persuasive. *See In re Mouttet*, 686 F.3d 1322, 1332–33 (Fed. Cir. 2012) (holding

Appendix D

that the test for obviousness is “what the combined teachings of the references would have suggested to those having ordinary skill in the art”); *Medichem, S.A. v. Rolabo, S.L.*, 437 F.3d 1157, 1166 (Fed. Cir. 2006) (explaining that in an obviousness analysis, “the prior art must be considered *as a whole* for what it teaches”); *In re Merck & Co., Inc.*, 800 F.2d 1091, 1097 (Fed. Cir. 1986) (“Non-obviousness cannot be established by attacking references individually where the rejection is based upon the teachings of a combination of references.”). Regardless, though, the RAM that Patent Owner points to in Figure 10 is part of the cell (i.e., functional unit) itself, not between cells.²⁴ *See* Pet. 60–64; Reply 48; Ex. 1009, 106 (“RaPiD is a linear array of functional units which is configured to form a mostly linear computational pipeline. This array of functional units is divided into identical cells which are replicated to form a complete array.”); Ex. 2064, 201:11–18.

Third, Patent Owner raises various challenges to all of Petitioner’s obviousness grounds collectively, in particular to Petitioner’s reliance on the testimony of Dr. Stone. Patent Owner argues, for example, that Dr.

24. As explained above, for purposes of determining whether memory is within a processing element or between processing elements, the boundaries of a processing element are not arbitrary, but rather depend on the disclosure of the particular reference and how the reference describes the processing element and the communication of data to and from the processing element. *See supra* Section II.C.3. In this case, Figure 10 of RaPiD shows the “[n]etlist for one *cell* of [the] 2-D DCT.” Ex. 1009, 111 (emphasis added).

Appendix D

Stone's testimony is conclusory, grounded in hindsight bias, fails to disclose the underlying facts or data on which his opinions are based under 37 C.F.R. § 42.65(a), fails to "articulate reasons why or how a [person of ordinary skill in the art] would combine the references" or consider whether the asserted combinations were "feasible," and "assumes only the benefits of his combinations and ignores any problems or drawbacks." PO Resp. 17–29, 110–115, 119–121 (emphasis omitted). Many of Patent Owner's arguments are premised on Petitioner failing to "acknowledge" purported "shortcomings" of Splash2, such as the "need for looping to be performed by [the] Sun workstation" and Splash2's lack of seamless passing of data and stream communication. *See, e.g., id.* at 27, 112–115, 120. Because we are not persuaded that Splash2 has any of those alleged deficiencies, those arguments are not persuasive. *See supra* Section II.E.2. Other arguments as to how Dr. Stone allegedly failed to understand the '324 patent, prior art, and problems and issues in the art are premised on his testimony in a different proceeding challenging the '687 patent and belied by the detailed explanations provided in his declarations, as explained herein. *See, e.g.,* PO Resp. 17–19, 21, 23–26, 113–114, 119–120 (citing Ex. 2066). We have reviewed Patent Owner's arguments and cited evidence in the record and disagree with the remainder of Patent Owner's assertions as well. We address the combination of Splash2 and RaPiD as representative.

Petitioner explains in detail what teachings of the two references it is relying on as teaching the various limitations of claims 8 and 9 and, importantly, explains

Appendix D

exactly how a person of ordinary skill in the art would have combined those teachings, i.e., instantiating the cell functionality of Figure 10 in the processing elements of the Splash2 unidirectional array implementation. Pet. 58–59, 64–65. Petitioner asserts that the references are combinable in an obviousness combination because they are both analogous art to the '324 patent and provides multiple reasons a person of ordinary skill in the art would have had for making the combination. *Id.* at 65–66. For example, RaPiD expressly cites the Splash 2 system as a “very successful example[.]” of a field-programmable custom computing machine. Ex. 1009, 106 (citing a reference describing the Splash 2 system), 115; *see* Pet. 64–67. Splash2 also expressly states that the Splash 2 system would be advantageous for image processing applications (of which JPEG and MPEG image compression are examples), suggesting that implementing the RaPiD calculations in the Splash 2 system “would have been successful and efficient.” *See* Ex. 1007, 141–142, 162 (“Even though Splash 2 was not designed specifically for image processing, this platform possesses architectural properties that make it well suited for the computation and data transfer rates that are characteristic of this class of problems. Furthermore, the price/performance of this system makes it a competitive alternative to conventional real-time image processing systems.”); Pet. 66–67. Other documentation indicates that image compression techniques using DCT were being used at the time, such that an ordinarily skilled artisan would have been motivated to find “more efficient and powerful solutions” to do so (e.g., by implementing the RaPiD DCT calculation in the Splash 2 system). *See* Ex. 1049, 73; Pet. 66.

Appendix D

Petitioner's assertions are supported by the testimony of Dr. Stone. *See* Ex. 1003 ¶¶ 330–369. That testimony is not based on hindsight bias, but rather the content of both asserted references, which Dr. Stone cites extensively in his analysis. *See id.* He explains the exact combination of teachings and the reasons for making the combination. *See id.* The reasons are not conclusory; they are explained in sufficient detail and premised on express statements in Splash2 and RaPiD connecting the two disclosures and suggesting that their combination would be advantageous for multiple reasons. Patent Owner never addresses those identified reasons to combine in its papers or explains in any way why they are factually incorrect. Finally, we note that Dr. Stone stated during cross-examination that in addition to advantages of the various obviousness combinations, he considered potential disadvantages where relevant to the combination. *See, e.g.,* Ex. 2066, 145:9–146:5. We do not find any evidence in the record that combining RaPiD's teaching of the DCT calculation with Splash2's teaching of the unidirectional array implementation would have had any disadvantages that would have outweighed Petitioner's stated advantages.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claims 8 and 9 would have been obvious based on Splash2 and RaPiD under 35 U.S.C. § 103(a).

*Appendix D***H. Obviousness Ground Based on Splash2 and Jeong
(Claim 20)**

Petitioner asserts that dependent claim 20 is unpatentable over Splash2 and Jeong. Pet. 67–77; *see* Ex. 1003 ¶¶ 425–461. Claim 20 recites that the calculation comprises “an encryption algorithm.” Petitioner relies on Jeong for this limitation, arguing that Jeong teaches “a systolic implementation of a modular multiplication algorithm for encryption suitable for implementation in FPGAs.” Pet. 68–69 (citing Ex. 1061, 214, Fig. 2(a)). Petitioner provides numerous reasons why a person of ordinary skill in the art would have been motivated to combine the teachings of Jeong and Splash2, including, for example, that Jeong “expressly cites the systolic structures as the platform to which its computing algorithms are mapped, and discloses the intent to make use of FPGAs as a possible means for implementation.” *Id.* at 75–77 (citing Ex. 1061, 211).

Patent Owner does not dispute that Jeong teaches “an encryption algorithm.” Instead, Patent Owner argues that Jeong does not teach two “computational loops,” as recited in parent claim 1, because the reference does not teach repeating a computation multiple times “for each piece of data.” PO Resp. 92–93. Patent Owner’s argument is premised on its proposed interpretation requiring repeated execution “per datum,” which we do not adopt. *See supra* Section II.C.2. Also, as explained above, we find that Splash2 teaches two “computational loops.” *See supra* Section II.E.2.

Appendix D

Likewise, with respect to the “seamlessly” limitation of claim 1, Patent Owner argues that similar to Splash2, “Jeong is ambiguous” regarding where results and inputs are stored. PO Resp. 109. Patent Owner further disputes the combination of references because Splash2 discloses a “linear” system and Jeong discloses a “non-linear” system. *Id.* at 118–119. Petitioner’s position, though, is that a person of ordinary skill in the art would have had reason to implement the “modular multiplication algorithm” of Jeong using the unidirectional systolic array described in Splash2—not the entirety of Jeong’s disclosed system. Pet. 70–71, 75–76; *see* Ex. 1003 ¶¶ 443, 444, 455, 458. That assertion is supported sufficiently by the testimony of Dr. Stone, which is not conclusory or based on hindsight bias as Patent Owner contends. *See* Ex. 1003 ¶¶ 456–461; *supra* Section II.G.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claim 20 would have been obvious based on Splash2 and Jeong under 35 U.S.C. § 103(a).

**I. Obviousness Ground Based on Splash2
and Chunky SLD (Claims 7, 17, and 24)**

Petitioner asserts that dependent claims 7, 17, and 24 are unpatentable over Splash2 and Chunky SLD. -1602 Pet. 63–79; *see* Ex. 1003 ¶¶ 373–417, 523–531. Claims 7, 17, and 24 limit the “calculation” in claim 1 to a particular type of calculation. For example, claim 17 recites that the calculation comprises “a search algorithm for an image search.” Petitioner relies on Chunky SLD

Appendix D

for the limitation, arguing that the reference teaches an “automatic target resolution algorithm . . . used in the search for partially obscured images in synthetic aperture radar data” and “implemented on a linear systolic array instantiated in a Splash 2 system.” -1602 Pet. 64 (citing Ex. 1011, 192). Petitioner provides numerous reasons why a person of ordinary skill in the art would have been motivated to combine the teachings of Chunky SLD with Splash2’s teachings regarding the unidirectional array implementation. *Id.* at 74–76. For example, Petitioner notes that Chunky SLD “expressly cites the Splash 2 system as the platform to which its computing algorithms are mapped.” *Id.* at 75 (citing Ex. 1011, 192).

Patent Owner does not dispute that Chunky SLD teaches the limitations of dependent claims 7, 17, and 24 or dispute Petitioner’s asserted reasons to combine Splash2 and Chunky SLD (other than the general arguments regarding motivation addressed above). *See supra* Section II.G. Instead, Patent Owner argues that Chunky SLD does not teach two “computational loops,” as recited in parent claim 1, because the reference does not teach repeating a computation multiple times “for each piece of data.” PO Resp. 91. Patent Owner’s argument is premised on its proposed interpretation requiring repeated execution “per datum,” which we do not adopt. *See supra* Section II.C.2. Also, as explained above, we find that Splash2 teaches two “computational loops.” *See supra* Section II.E.2.

Likewise, with respect to the “seamlessly” limitation of claim 1, Patent Owner argues that “Chunky SLD is

Appendix D

simply an algorithm deployed on Splash2, and therefore the same ambiguity present in Splash2 [regarding where data is stored] is also present.” PO Resp. 108–109. We disagree that Splash2 is ambiguous. *See supra* Section II.E.2.a.3. Further, Petitioner’s arguments regarding the combined teachings of Splash2 and Chunky SLD are supported sufficiently by the testimony of Dr. Stone, which is not conclusory or based on hindsight bias as Patent Owner contends. *See* -1602 Pet. 69–70, 73–76; Ex. 1003 ¶¶ 409–417; *supra* Section II.G.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claims 7, 17, and 24 would have been obvious based on Splash2 and Chunky SLD under 35 U.S.C. § 103(a).

**J. Obviousness Ground Based on Splash2
and Roccatano (Claims 2–5, 22, and 23)**

Petitioner asserts that dependent claims 2–5, 22, and 23 are unpatentable over Splash2 and Roccatano. -1603 Pet. 67–86; *see* Ex. 1003 ¶¶ 280–329, 466–522. Claims 22 and 23 limit the “calculation” in claim 1 to a particular type of calculation, and claims 2–5 further limit the “subsequent and previous data dimensions” of the calculation in claim 1. For example, claim 2 recites that “said subsequent and previous data dimensions of said calculation comprise multiple vectors in said calculation.” Petitioner relies on Roccatano for the limitation, arguing that Roccatano teaches “a program for simulating

Appendix D

molecular dynamics interactions of molecules and compounds” using “a systolic loop of processors as laid out on a 3-D grid” shown in Figure 2 of Roccatano. -1603 Pet. 67–69. Petitioner provides numerous reasons why a person of ordinary skill in the art would have been motivated to combine the teachings of Roccatano with Splash2’s teaching of the unidirectional array implementation, including, for example, that “Splash2 discloses that it is one of two reconfigurable systems that have achieved ‘supercomputer performance’ on applications that include molecular biology, which is the underlying application for Roccatano.” *Id.* at 77–79 (citing Ex. 1007, 6).

Patent Owner does not dispute that Roccatano teaches the limitations of dependent claims 2–5, 22,²⁵ and 23 or dispute Petitioner’s asserted reasons to combine Splash2 and Roccatano (other than the general arguments regarding motivation addressed above). *See supra* Section II.G. Instead, Patent Owner argues that Roccatano does not teach two “computational loops,” as recited in parent claim 1, because the reference does not teach repeating a computation multiple times “for each piece of data.” PO Resp. 91–92. Patent Owner’s argument is premised on its proposed interpretation requiring repeated execution “per datum,” which we do not adopt. *See supra* Section II.C.2. Also, as explained above, we find that Splash2 teaches two “computational loops.” *See supra* Section II.E.2.

25. We also conclude that Petitioner has proven that claim 22 is anticipated by Splash2. *See supra* Section II.E.3.

Appendix D

Likewise, with respect to the “seamlessly” limitation of claim 1, Patent Owner argues that Roccatano does not teach passing computed data “seamlessly” because it uses multiple processors and, even if not, it is “ambiguous” where the processors store intermediate results, similar to the alleged teachings of Splash2. PO Resp. 106–107. Petitioner’s position, though, is that a person of ordinary skill in the art would have had reason to implement the calculations of Roccatano using the unidirectional systolic array described in Splash2—not using the multiple processors of Roccatano. -1603 Pet. 70–72, 76–81, 85–86; *see* Ex. 1003 ¶¶ 291, 294, 315, 493, 495–497, 501, 504, 517. Those assertions are supported sufficiently by the testimony of Dr. Stone, which is not conclusory or based on hindsight bias as Patent Owner contends. *See* Ex. 1003 ¶¶ 299–318, 514–520; *supra* Section II.G. Patent Owner’s arguments attacking Roccatano individually are not persuasive.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claims 2–5, 22, and 23 would have been obvious based on Splash2 and Roccatano under 35 U.S.C. § 103(a).

**K. Obviousness Grounds Based on Combinations
with Gaudiot (Claims 1–5, 7–9, 15, 17, 18, and 20–24)**

Petitioner contends that claim 1 is unpatentable over Splash2 and Gaudiot, relying on Gaudiot in the alternative for one limitation of claim 1. *See* Pet. 52–55; -1602 Pet.

Appendix D

60–63; -1603 Pet. 63–67. Specifically, Petitioner states that “[t]o the extent one might argue that Splash2 does not sufficiently disclose that the edit distance calculation ‘is systolically implemented by said reconfigurable computing system at the at least one reconfigurable processor,’” as recited in claim 1, it would have been obvious to modify Splash2 to do so based on the teachings of Gaudiot. Pet. 52–55 (emphasis omitted). Petitioner does not rely on Gaudiot for any of the limitations of claim 1 that are disputed by Patent Owner, namely “computational loops” and passing computed data “seamlessly” between the computational loops. *See id.*; Tr. 51:4–18. As explained above, we conclude that Splash2 discloses the systolic implementation limitation of claim 1 (which is not disputed by Patent Owner in its Response) and that Splash2 anticipates claim 1. *See supra* Section II.E.2. As such, we need not address Petitioner’s alternative ground of unpatentability of claim 1 based on the combination of Splash2 and Gaudiot. Likewise, for all of Petitioner’s other asserted obviousness grounds, Petitioner relies on certain references “with or without” Gaudiot. *See* Pet. 5; -1602 Pet. 6; -1603 Pet. 5. We address the grounds “without” Gaudiot above and need not address Petitioner’s alternative grounds based on combinations “with” Gaudiot. *See supra* Sections II.F–II.J.

L. Constitutionality Argument

Patent Owner argues that this proceeding should be dismissed because the instant panel “was unconstitutionally appointed in violation of the Appointments Clause of the U.S. Constitution, Article II, Section 2, Clause 2,”

Appendix D

and “lacks the constitutional authority to enter a final decision.” Sur-Reply 21–23. We decline to consider Patent Owner’s constitutional challenge, as the issue has been addressed in *Arthrex, Inc. v. Smith & Nephew, Inc.*, 941 F.3d 1320 (Fed. Cir. 2019), *reh’g denied*, 933 F.3d 760 (Fed. Cir. 2020).

III. CONCLUSION²⁶

Petitioner has demonstrated, by a preponderance of the evidence, that claims 1–5, 7–9, 15, 17, 18, and 20–24 are unpatentable. In summary:

26. Should Patent Owner wish to pursue amendment of the challenged claims in a reissue or reexamination proceeding subsequent to the issuance of this Decision, we draw Patent Owner’s attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. See 84 Fed. Reg. 16,654 (Apr. 22, 2019). If Patent Owner chooses to file a reissue application or a request for reexamination of the challenged patent, we remind Patent Owner of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. See 37 C.F.R. §§ 42.8(a)(3), 42.8(b)(2).

Appendix D

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not shown Unpatentable
1, 15, 18, 21, 22	102(a), 102(b)	Splash2	1, 15, 18, 21, 22	
1, 15, 18, 21, 22	103(a)	Splash2	1, 15, 18, 21, 22	
1, 15, 18, 21, 22	103(a)	Splash2, Gaudiot ²⁷		
8, 9	103(a)	Splash2, RaPiD	8, 9	
8, 9	103(a)	Splash2, RaPiD, Gaudiot		
20	103(a)	Splash2, Jeong	20	
20	103(a)	Splash2, Jeong, Gaudiot		

27. As explained above, we do not reach Petitioner's alternative grounds based on combinations with Gaudiot. *See supra* Section II.K.

Appendix D

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not shown Unpatentable
7, 17, 24	103(a)	Splash2, Chunky SLD	7, 17, 24	
7, 17, 24	103(a)	Splash2, Chunky SLD, Gaudiot		
2-5, 22, 23	103(a)	Splash2, Roccatano	2-5, 22, 23	
2-5, 22, 23	103(a)	Splash2, Roccatano, Gaudiot		
Overall Outcome			1-5, 7-9, 15, 17, 18, 20-24	

192a

Appendix D

IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that claims 1–5, 7–9, 15, 17, 18, and 20–24 of the '324 patent have been shown to be unpatentable;

FURTHER ORDERED that Petitioner's Motion to Exclude (Paper 60) is *granted-in-part, denied-in-part, and dismissed-in-part*;

FURTHER ORDERED that Patent Owner's Motion to Exclude (Paper 61) is *denied-in-part and dismissed-in-part*; and

FURTHER ORDERED that Exhibit 2168 is expunged from the record of this proceeding.

This is a final decision. Parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2018-01601
Patent 7,225,324 B2

FOR PETITIONER:

Joseph A. Micallef
Jason Greenhut
Scott Border

193a

Appendix D

SIDLEY AUSTIN LLP
jmicallef@sidley.com
jgreenhut@sidley.com
sborder@sidley.com

FOR PATENT OWNER

Alfonso Chan
Joseph DePumpo
SHORE CHAN DePUMPO LLP
achan@shorechan.com
jdepumpo@shorechan.com

**APPENDIX E — JUDGMENT OF THE UNITED
STATES PATENT AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL AND APPEAL
BOARD, DATED APRIL 9, 2020**

UNITED STATES PATENT
AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL
AND APPEAL BOARD

MICROSOFT CORPORATION,

Petitioner,

v.

FG SRC LLC,¹

Patent Owner.

IPR2018-01605²
Patent 7,620,800 B2

Before KALYAN K. DESHPANDE, JUSTIN T. ARBES,
and CHRISTA P. ZADO, *Administrative Patent Judges.*

ARBES, *Administrative Patent Judge.*

1. Patent Owner filed updated mandatory notice information indicating that DirectStream, LLC (“DirectStream”) assigned the challenged patent to FG SRC LLC. Paper 69, 1. Accordingly, the caption for this proceeding has been changed.

2. Cases IPR2018-01606 and IPR2018-01607 have been consolidated with this proceeding.

*Appendix E***JUDGMENT**

Final Written Decision
 Determining All Challenged Claims Unpatentable
35 U.S.C. § 318(a)

I. INTRODUCTION*A. Background and Summary*

Petitioner Microsoft Corporation filed three Petitions, collectively requesting *inter partes* review of claims 1–5, 7–9, 15, 17, 18, and 20–24 of U.S. Patent No. 7,620,800 B2 (Ex. 1005, “the ’800 patent”) pursuant to 35 U.S.C. § 311(a), as listed in the following chart.³

Case Number	Challenged Claims	Petition
IPR2018-01605	1, 8, 9, and 20	Paper 1 (“Pet.”)
IPR2018-01606	1, 7, 15, 17, and 24	Paper 1 (“-1606 Pet.”)
IPR2018-01607	1–5, 18, and 21–23	Paper 1 (“-1607 Pet.”)

On April 12, 2019, we instituted an *inter partes* review as to all challenged claims on all grounds of unpatentability asserted in the Petitions, and exercised our authority under 35 U.S.C. § 315(d) to consolidate the three proceedings and conduct the proceedings as one trial.

3. Unless otherwise noted, references herein are to the exhibits filed in Case IPR2018-01605.

Appendix E

Paper 21 (“Decision on Institution” or “Dec. on Inst.”). Patent Owner FG SRC LLC subsequently filed a Patent Owner Response (Paper 36, “PO Resp.”), Petitioner filed a Reply (Paper 49, “Reply”), and Patent Owner filed a Sur-Reply (Paper 59, “Sur-Reply”). Petitioner filed a Motion to Exclude (Paper 60, “Pet. Mot.”) certain evidence submitted by Patent Owner, to which Patent Owner filed an Opposition (Paper 63, “PO Opp.”) and Petitioner filed a Reply (Paper 66, “Pet. Mot. Reply”). Patent Owner filed a Motion to Exclude (Paper 61, “PO Mot.”) certain evidence submitted by Petitioner, to which Petitioner filed an Opposition (Paper 62, “Pet. Opp.”) and Patent Owner filed a Reply (Paper 65, “PO Mot. Reply”). An oral hearing was held on February 4, 2020, and a transcript of the hearing is included in the record (Paper 71, “Tr.”).

We have jurisdiction under 35 U.S.C. § 6. This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a). For the reasons that follow, we determine that Petitioner has shown by a preponderance of the evidence that claims 1–5, 7–9, 15, 17, 18, and 20–24 are unpatentable.

B. Related Matters

The parties indicate that the ’800 patent is the subject of the following district court cases: *SRC Labs, LLC v. Microsoft Corp.*, No. 2:18-cv-00321 (W.D. Wash.), and *SRC Labs, LLC v. Amazon Web Servs., Inc.*, No. 2:18-cv-00317 (W.D. Wash.). See Pet. 4–5; Paper 69, 1.

*Appendix E**C. The '800 Patent*

The '800 patent⁴ discloses “multi-adaptive processing systems and techniques for enhancing parallelism and performance of computational functions.” Ex. 1005, col. 1, ll. 40–43. Parallel processing “allows multiple processors to work simultaneously on the same problem to achieve a solution” in less time than it would take a single processor. *Id.* at col. 1, ll. 44–49. “[A]s more and more performance is required, so is more parallelism, resulting in ever larger systems” and associated difficulties, including “facility requirements, power, heat generation and reliability.” *Id.* at col. 1, ll. 53–61. The '800 patent discloses that

if a processor technology could be employed that offers orders of magnitude more parallelism per processor, these systems could be reduced in size by a comparable factor. Such a processor or processing element is possible through the use of a reconfigurable processor. Reconfigurable processors instantiate only the functional units needed to solve a particular application, and as a result, have available space to instantiate as many functional units as may be required to solve the problem up to the total capacity of the integrated circuit chips they employ.


Id. at col. 1, l. 65–col. 2, l. 7. The '800 patent describes a known issue where each processor in a multi-processor

4. The '800 patent is a continuation of U.S. Patent No. 7,225,324 B2 (Ex. 1001, “the '324 patent”), is challenged by Petitioner in Case IPR2018-01601. We enter a Final Written Decision in Case IPR2018-01601 concurrently with this Decision.

Appendix E

system is allocated a portion of a problem called a “cell” and “to solve the total problem, results of one processor are often required by many adjacent cells because their cells interact at the boundary.” *Id.* at col. 2, ll. 26–32. Passing intermediate results around the system to complete the problem requires using “numerous other chips and busses that run at much slower speeds than the microprocessor,” diminishing performance. *Id.* at col. 2, ll. 32–38, col. 5, ll. 16–28, Fig. 1 (depicting a conventional multi-processor arrangement). In an adaptive processor-based system, however, “any boundary data that is shared between . . . functional units need never leave a single integrated circuit chip,” reducing “data moving around the system” and improving performance. *Id.* at col. 2, ll. 39–49.

Figure 2 of the '800 patent is reproduced below.



Appendix E

Figure 2 is “a functional block diagram of an adaptive processor 200 communications path for implementing the technique of the present invention.” *Id.* at col. 5, ll. 29–32. Adaptive processor 200 includes adaptive processor chip 202, which is coupled to memory element 206, interconnect 208, and additional adaptive processor chips 210. *Id.* at col. 5, ll. 32–37. Adaptive processor chip 202 includes thousands of functional units (“FU”) 204 interconnected by “reconfigurable routing resources” inside adaptive processor chip 202, allowing functional units 204 to “exchange data at much higher data rates and lower latencies than a standard microprocessor.” *Id.* at col. 5, ll. 39–45.

Figures 4A and 4B of the '800 patent are reproduced below.




Figure 4A depicts conventional sequential processing operation 400 where “nested Loops A (first loop 402) and B (second loop 404) are alternately active on different

Appendix E

phases of the process.” *Id.* at col. 5, l. 65–col. 6, l. 2. Because first loop 402 must be completed before beginning second loop 404, “all of the logic that has been instantiated is not being completely utilized.” *Id.* at col. 6, ll. 3–9. Figure 4B depicts “multi-dimensional process 410 in accordance with the technique of the present invention.” *Id.* at col. 6, ll. 11–14. “[M]ulti-dimensional process 410 is effectuated such that multiple dimensions of data are processed by both Loops A (first loop 412) and B (second loop 414) such that the computing system logic is operative on every clock cycle.” *Id.* at col. 6, ll. 14–18. A “dimension” of data can be “multiple vectors of a problem, multiple plans of a problem, multiple time steps in a problem and so forth.” *Id.* at col. 6, ll. 25–28. The ’800 patent discloses that available resources are utilized more effectively in the multi-dimensional process by “hav[ing] an application evaluate a problem in a data flow sense. That is, it will ‘pass’ a subsequent dimension of a given problem through the first loop 412 of logic concurrently with the previous dimension of data being processed through the second loop.” *Id.* at col. 6, ll. 19–25.

The ’800 patent states that the disclosed process can be utilized for a variety of applications. *Id.* at col. 9, ll. 10–20. For example, seismic imaging applications, which “process echo data to produce detailed analysis of subsurface features” for oil and gas exploration, would “particularly benefit from the tight parallelism that can be found in the use of adaptive or reconfigurable processors” because they “use data collected at numerous points and consisting of many repeated parameters” and “the results of the computation on one data point are used

Appendix E

in the computation of the next.” *Id.* at col. 9, ll. 25–34; *see id.* at col. 6, l. 30–col. 7, l. 37, Figs. 5A–5B, 6A–6B (describing a seismic imaging function that can be adapted to utilize the disclosed parallelism, where computational process 610 “loops over the depth slices as indicated by reference number 622 and loops over the shots as indicated by reference number 624”). Also, reservoir simulation applications, which “process fluid flow data in . . . oil and gas subsurface reservoirs to produce extraction models,” would benefit from the disclosed process because they define a three dimensional set of cells for the reservoir, utilize repeated operations on each cell, and “information computed for each cell is then passed to neighboring cells.” *Id.* at col. 9, l. 59–col. 10, l. 2; *see id.* at col. 7, l. 38–col. 8, l. 20, Figs. 7A–7D (describing “process 700 for performing a representative systolic wavefront operation in the form of a reservoir simulation function” in which “the computation of fluid flow properties are communicated to neighboring cells 710” without storing data in memory, “a set of cells can reside in an adaptive processor,” and “the pipeline of computation can extend across multiple adaptive processors,” where the process involves nested loops and systolic walls 712 and 714 of computation at different time sets), col. 8, l. 21–col. 9, l. 9, Figs. 8A–8C, 9A–9C (describing two other processes for performing “a representative systolic wavefront operation”). Finally, the disclosed process may be used for genetic pattern matching applications, which “look[] for matches of a particular genetic sequence (or model) to a database of genetic records,” performing repeated operations to “compare[] each character in the model to the characters in [a particular] genetic record.” *Id.* at col. 11, ll. 42–52.

*Appendix E**D. Illustrative Claims*

Claim 1 of the '800 patent is independent. Claims 2–5, 7–9, 15, 17, 18, and 20–24 each depend from claim 1. Claims 1 and 15 recite:

1. A method for data processing in a reconfigurable computing system, the reconfigurable computing system comprising at least one reconfigurable processor, the reconfigurable processor comprising a plurality of functional units, said method comprising:

transforming an algorithm into a data driven calculation that is implemented by said reconfigurable computing system at the at least one reconfigurable processor;

forming at least two of said functional units at the at least one reconfigurable processor to perform said calculation wherein only functional units needed to solve the calculation are formed and wherein each formed functional unit at the at least one reconfigurable processor interconnects with each other formed functional unit at the at least one reconfigurable processor based on reconfigurable routing resources within the at least one reconfigurable processor as established at formation, and wherein lines of code of said calculation are formed as clusters of functional units within the at least one reconfigurable processor;

Appendix E

utilizing a first of said formed functional units to operate upon a subsequent data dimension of said calculation forming a first computational loop; and

substantially concurrently utilizing a second of said formed functional units to operate upon a previous data dimension of said calculation generating a second computational loop wherein said implementation of said calculation enables said first computational loop and said second computational loop execute concurrently and pass computed data seamlessly between said computational loops.

15. The method of claim 1 wherein instantiating includes establishing a stream communication connection between functional units.

E. Evidence

The pending grounds of unpatentability in the instant *inter partes* review are based on the following prior art:

Jean-Luc Gaudiot, “Data-Driven Multicomputers in Digital Signal Processing,” *Proceedings of the IEEE, Special Issue on Hardware and Software for Digital Signal Processing*, vol. 75, no. 9, Sept. 1987, pp. 1220–1234 (Ex. 1010, “Gaudiot”);

Duncan A. Buell, Jeffrey M. Arnold, & Walter J. Kleinfelder, SPLASH2: FPGAS IN A

Appendix E

CUSTOM COMPUTING MACHINE (1996) (Ex. 1007, “Splash2”);

Carl Ebeling *et al.*, “Mapping Applications to the RaPiD Configurable Architecture,” *Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines*, Apr. 16–18, 1997, pp. 106–115 (Ex. 1009, “RaPiD”);

Michael Rencher & Brad L. Hutchings, “Automated Target Recognition on SPLASH 2,” *Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines*, Apr. 16–18, 1997, pp. 192–200 (Ex. 1011, “Chunky SLD”);

Yong-Jin Jeong & Wayne P. Burleson, “VLSI Array Algorithms and Architectures for RSA Modular Multiplication,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 2, June 1997, pp. 211–217 (Ex. 1061, “Jeong”); and

D. Roccatano *et al.*, “Development of a Parallel Molecular Dynamics Code on SIMD Computers: Algorithm for Use of Pair List Criterion,” *Journal of Computational Chemistry*, vol. 19, no. 7, May 1998, pp. 685–694 (Ex. 1012, “Roccatano”).⁵

5. When citing the prior art references and other exhibits, we refer to the page numbers in the bottom-right corner added by the filing party. *See* 37 C.F.R. § 42.63(d)(2).

Appendix E

Petitioner filed a declaration from Harold Stone, Ph.D. (Ex. 1003) with its Petition and a reply declaration from Dr. Stone (Ex. 1076) with its Reply. Patent Owner filed declarations from Jon Huppenthal (Ex. 2101), Houman Homayoun, Ph.D. (Exs. 2029, 2112), and Tarek El-Ghazawi, Ph.D. (Ex. 2166).

F. Asserted Grounds

The instant *inter partes* review involves the following grounds of unpatentability:

Petition(s)	Claim(s) Challenged	35 U.S.C. §	Reference(s)/ Basis
IPR2018-01605, IPR2018-01606, IPR2018-01607	1, 15, 18, 21, 22	102(a), 102(b) ^{<?>}	Splash2
IPR2018-01605, IPR2018-01606, IPR2018-01607	1, 15, 18, 21, 22	103(a)	Splash2

Appendix E

Petition(s)	Claim(s) Challenged	35 U.S.C. §	Reference(s)/ Basis
IPR2018-01605, IPR2018-01606, IPR2018-01607	1, 15, 18, 21, 22	103(a)	Splash2, Gaudiot
IPR2018-01605	8, 9	103(a)	Splash2, RaPiD
IPR2018-01605	8, 9	103(a)	Splash2, RaPiD, Gaudiot
IPR2018-01605	20	103(a)	Splash2, Jeong
IPR2018-01605	20	103(a)	Splash2, Jeong, Gaudiot
IPR2018-01606	7, 17, 24	103(a)	Splash2, Chunky SLD
IPR2018-01606	7, 17, 24	103(a)	Splash2, Chunky SLD, Gaudiot
IPR2018-01607	2-5, 22, 23	103(a)	Splash2, Roccatano
IPR2018-01607	2-5, 22, 23	103(a)	Splash2, Roccatano, Gaudiot

*Appendix E***II. ANALYSIS***A. Motions to Exclude*

The party moving to exclude evidence bears the burden of proof to establish that it is entitled to the relief requested—namely, that the material sought to be excluded is inadmissible under the Federal Rules of Evidence. *See* 37 C.F.R. §§ 42.20(c), 42.62(a). For the reasons discussed below, Petitioner’s Motion is granted-in-part, denied-in-part, and dismissed-in-part, and Patent Owner’s Motion is denied-in-part and dismissed-in-part.

1. Petitioner’s Motion to Exclude

Exhibits 2101, 2065, 2075, and 2091: Petitioner moves to exclude the entirety of the declaration (Ex. 2101) of Mr. Huppenthal, one of the named inventors of the ’800 patent, “as not being relevant to any issue on which trial has been instituted, and for lacking foundation, containing hearsay, and/or causing undue prejudice.” Pet. Mot. 3–6. Petitioner argues that the declaration includes “irrelevant narrative discussion of [Mr. Huppenthal’s] participation in reconfigurable computing” and statements “either based on hearsay or lack of personal knowledge.” *Id.* at 3–4. Petitioner also moves to exclude paragraphs 80 and 82–86 of the declaration based on Mr. Huppenthal’s alleged “refusal to answer questions concerning those portions of the declaration” during cross-examination. *Id.* at 1–3.

Petitioner also moves to exclude three transcripts (Exs. 2065, 2075, 2091) of depositions of Petitioner’s

Appendix E

declarants from other *inter partes* reviews as “not being relevant to any issue on which trial has been instituted, for containing hearsay, and/or causing undue prejudice.” *Id.* at 6–7. Petitioner argues that allowing the transcripts in the record would be “highly prejudicial as they present themselves with the indicia of expert testimony while being totally devoid from the necessary context of the matters from which they originate.” *Id.* Patent Owner cites Exhibits 2065 and 2075 in its Response, but does not cite Exhibit 2091 in its Response or Sur-Reply.

Petitioner’s Motion is dismissed as moot, as we do not rely on the cited portions of the testimony in a manner adverse to Petitioner in this Decision. As explained below, even if the testimony is considered, we are not persuaded by Patent Owner’s arguments regarding the state of the art or alleged nonobviousness of the challenged claims, and Patent Owner has not shown proof of secondary considerations that would support a conclusion of nonobviousness. *See infra* Sections II.E–II.J.

Exhibit 2112: Petitioner moves to exclude paragraphs 41, 145, 180, 181, and 247 of the declaration of Dr. Homayoun, which refer to Exhibits 2066 and 2101. Pet. Mot. 8–9. Because we do not exclude those exhibits, we also dismiss as moot Petitioner’s Motion with respect to Exhibit 2112.

Exhibits 2066–2074, 2076, 2078–2099, 2102–2104, 2106, 2107, 2110, 2113–2134, 2140–2152, 2156, 2163, 2165, and 2170: Petitioner moves to exclude a number of exhibits as “not being relevant to any issues on which trial has

Appendix E

been instituted, lacking foundation, and/or causing undue prejudice” because the exhibits were not discussed or cited, or “only cited superficially,” in Patent Owner’s Response and Sur-Reply. Pet. Mot. 7–8. Petitioner’s Motion is dismissed as moot, as we do not rely on the exhibits in a manner adverse to Petitioner in this Decision. We note, however, that in evaluating Petitioner’s asserted grounds of unpatentability, we only consider substantive arguments made by the parties in their papers during trial (i.e., the Petitions, Response, Reply, and Sur-Reply). To the extent a document is filed in the record but never discussed in a paper, there is no substantive argument pertaining to that document that can be considered.

Exhibit 2170: Petitioner moves to exclude Exhibit 2170 under Federal Rules of Evidence 401–403 as cumulative of Exhibit 1007. Pet. Mot. 8. Both exhibits are copies of Splash2. Dr. El-Ghazawi refers to the document in his declaration with the numeral “1007.” Ex. 2166 ¶ 40. To ensure a clear record, we grant Petitioner’s Motion, expunge Exhibit 2170, and refer herein to Exhibit 1007. *See* 37 C.F.R. §§ 42.7(a), 42.6(d) (“A document already in the record of the proceeding must not be filed again, not even as an exhibit or an appendix, without express Board authorization.”).

Portions of Patent Owner’s Response: Petitioner moves to exclude portions of Patent Owner’s Response referring to the exhibits that Petitioner seeks to exclude. Pet. Mot. 9. Patent Owner’s Response is a paper with

Appendix E

attorney arguments, not evidence that may be excluded.⁶ Further, we do not exclude any of the exhibits referred to in the identified portions of the Response. Petitioner’s Motion is denied as to Patent Owner’s Response.

2. *Patent Owner’s Motion to Exclude*

Exhibits 1074, 1077, and 1079: Patent Owner moves to exclude three technical documents filed by Petitioner with its Reply.

First, Patent Owner moves to exclude Exhibits 1074 and 1079 as unauthenticated under Federal Rule of Evidence 901. PO Mot. 6–7. “To satisfy the requirement of authenticating or identifying an item of evidence, the proponent must produce evidence sufficient to support a finding that the item is what the proponent claims it is.” Fed. R. Evid. 901(a). Certain evidence, though, is “self-authenticating” and “require[s] no extrinsic evidence of authenticity in order to be admitted.” Fed. R. Evid. 902. Exhibit 1074 is an article by Maya Gokhale and Ron Minnich titled “FPGA Computing in a Data Parallel C,” and includes an IEEE trade inscription, copyright symbol, and International Standard Book Number (ISBN) on the first page (“0-8186-3890-7/93 \$03.00 © 1993 IEEE”). It is self-authenticating under at least Federal Rule of Evidence 902(6) (“Printed material purporting to be a . . .

6. Petitioner did not seek authorization to file a motion to strike Patent Owner’s Response. *See* Patent Trial and Appeal Board Consolidated Trial Practice Guide (Nov. 2019), 80–81, *available at* <https://www.uspto.gov/TrialPracticeGuideConsolidated> (“Trial Practice Guide”).

Appendix E

periodical.”) and 902(7) (“An inscription, sign, tag, or label purporting to have been affixed in the course of business and indicating origin, ownership, or control.”). Exhibit 1079 is an excerpt from a book by Kevin Skahill titled “VHDL for Programmable Logic.” It includes a copyright notice (“Copyright © 1996 by Addison-Wesley Publishing, Inc.”), Library of Congress catalogue information, and hand-written library catalogue information. Ex. 1079, Cover 3, v. Petitioner points out where the book is available from “several well-known book sellers.” Pet. Opp. 4 & n.2. It is authenticated at least as an ancient document under Federal Rule of Evidence 901(b)(8) because it “is in a condition that creates no suspicion about its authenticity,” “was in a place where, if authentic, it would likely be,” and “is at least 20 years old when offered.” Nothing about either exhibit suggests that it is not what it points to be, and Patent Owner does not point to anything in particular in the exhibits that would indicate otherwise. *See* PO Mot. 6–7; PO Mot. Reply 1–4.

Second, Patent Owner moves to exclude Exhibits 1074, 1077, and 1079 as containing inadmissible hearsay under Federal Rule of Evidence 802. PO Mot. 7. Patent Owner states that Petitioner in its Reply “cites each of these documents to prove the truth of technical matters allegedly asserted in such documents, *i.e.* to support Petitioner’s specific factual assertions regarding a technical issue.” *Id.* We are not persuaded. Patent Owner does not identify any particular “statement” in any of the exhibits that is being offered “to prove the truth of the matter asserted in the statement,” and thus fails to meet its burden to prove inadmissibility as hearsay. *See* Fed.

Appendix E

R. Evid. 801(c); 37 C.F.R. § 42.20(c). Even if Patent Owner had done so, Petitioner cites the exhibits to show what a person of ordinary skill in the art would have known at the time of the '800 patent and, with respect to Exhibit 1074 in particular, how a person of ordinary skill in the art would have understood another reference relied on by Patent Owner that cites the article. *See* Reply 7, 17, 33; Pet. Opp. 5–6; Ex. 2169, 37–38, 208. The exhibits are not being offered for the truth of any particular matter discussed in the references. Finally, Exhibits 1074 and 1079 also are admissible under Federal Rule of Evidence 803(16), which provides as an exception to the hearsay rule “[a] statement in a document that was prepared before January 1, 1998, and whose authenticity is established.” *See* 1074, 94 (“1993” date); Ex. 1079, Cover 3, v (“1996” date).

Third, Patent Owner moves to exclude Exhibit 1077 as “irrelevant to the patent and claim construction issues in dispute.” PO Mot. 7–8. Exhibit 1077 is an excerpt of certain pages of the Microsoft Computer Dictionary (5th ed. 2002). Patent Owner argues that Exhibit 1077 “is extrinsic evidence pertaining to Petitioner’s proffered definitions from the Microsoft Computer Dictionary of ‘data structure’ and ‘data path,’ neither of which are claim terms in the patent.” PO Mot. 7. Federal Rule of Evidence 401 provides that “[e]vidence is relevant if: (a) it has any tendency to make a fact more or less probable than it would be without the evidence; and (b) the fact is of consequence in determining the action.”

Petitioner submits Exhibit 1077 in support of its arguments regarding the interpretation of disputed claim

Appendix E

language, specifically the term “stream communication.” Reply 29. The meaning of this phrase is “of consequence in determining” whether challenged claim 15 is unpatentable over the asserted prior art, and Exhibit 1077 provides insight as to the meaning of words used in both parties’ proposed interpretations. *See infra* Section II.C.4. Exhibit 1077 has some “tendency to make a fact more or less probable than it would be without the evidence” and is relevant under Federal Rule of Evidence 401. Thus, there is no basis to exclude Exhibit 1077. We also note that Patent Owner appears to have filed a full copy of the dictionary as Exhibit 2065, and the same pages filed as Exhibit 1077 (pages 144–145) that Patent Owner seeks to exclude are in the exhibit that Patent Owner filed. Further, the record contains numerous other dictionary references filed by both parties, including Exhibits 1025, 1059, 2024–2026, and 2038. Patent Owner’s Motion is denied as to Exhibits 1074, 1077, and 1079.

Exhibit 1076: Patent Owner also moves to exclude paragraphs 15–17 of the reply declaration of Dr. Stone, which refer to Exhibit 1074 “for the first time,” under Federal Rules of Evidence 402 and 403. PO Mot. 8. Because we find no basis to exclude Exhibit 1074, we also deny Patent Owner’s Motion with respect to Exhibit 1076. To the extent Patent Owner’s position is that Petitioner’s Reply and Dr. Stone’s reply declaration exceed the proper scope of a reply, we address those arguments below. *See infra* Section II.E.2.a.2; Trial Practice Guide, 79 (“A motion to exclude is not a vehicle for addressing the weight to be given evidence—arguments regarding weight should appear only in the merits documents. Nor should a motion to exclude address arguments or evidence that a party

Appendix E

believes exceeds the proper scope of reply or sur-reply.”).

Exhibits 1075 and 1078: Patent Owner moves to exclude certain portions of the transcript of the deposition of Dr. Homayoun because the questions asked were “vague, ambiguous, call[] for a legal conclusion, and misleading.” PO Mot. 9–11 (citing Ex. 1075, 27:4–12, 65:5–17). Patent Owner also moves to exclude a portion of the transcript of the deposition of Dr. El-Ghazawi because the question asked was “vague, ambiguous, and calls for a speculative answer.” *Id.* at 11 (citing Ex. 1078, 65:12–17). Patent Owner’s Motion is dismissed as moot, as we do not rely on the disputed portions of the testimony in rendering our Decision.

B. Level of Ordinary Skill in the Art

In determining the level of ordinary skill in the art for a challenged patent, we look to “1) the types of problems encountered in the art; 2) the prior art solutions to those problems; 3) the rapidity with which innovations are made; 4) the sophistication of the technology; and 5) the educational level of active workers in the field.” *Ruiz v. A.B. Chance Co.*, 234 F.3d 654, 666–667 (Fed. Cir. 2000). “Not all such factors may be present in every case, and one or more of them may predominate.” *Id.*

Petitioner’s declarant, Dr. Stone, testifies that a person of ordinary skill in the art at the time of the ’800 patent would have had “an advanced degree in electrical or computer engineering, or computer science with substantial study in computer architecture, hardware design, and computer algorithms,” and “at least three

Appendix E

years' experience working in the field," or alternatively "a bachelor's degree covering those disciplines and at least four years working [in] the field." Ex. 1003 ¶ 45. According to Dr. Stone,

[s]uch a person would also have been knowledgeable about the programming, design and operation of computer systems based on reconfigurable components such as FPGAs (field programmable gate arrays) and CPLDs (complex programmable logic devices), including computer systems for performing systolic and data driven calculations. That person would also have been familiar with hardware description languages such as [Very High Speed Integrated Circuit Hardware Description Language (VHDL)] that could be used to configure FPGAs and CPLDS that serve as components of reconfigurable computer systems. Finally, as demonstrated by many of the references discussed [in Dr. Stone's declaration], such a person would also have been familiar with various other areas of technology that by 2002 had relied on high performance and parallel computing systems, such as genetic sequence comparisons, image processing, data mining, and processing related to proteins and organic structures.

Id.

Patent Owner states that it "does not dispute the level of education and skill promoted by [Dr. Stone]," and

Appendix E

Patent Owner’s declarant, Dr. Homayoun, “agree[d] with Dr. Stone’s assessment of the level of ordinary skill in the art,” noting that such an individual also would have been “experienced in developing with high-level languages (C and Fortran), hardware description languages, and the unique problems involved with programming FPGAs and FPGA based systems.” *See* PO Resp. 29; Ex. 2029 ¶ 17; *see also* Ex. 2112 ¶ 133 (“In general, I would agree to the level of education and skill promoted by [Petitioner’s] expert [for the ’800 patent].”). Dr. Homayoun further expands on his understanding of what a person of ordinary skill in the art at the time of the ’800 patent would have known and considered. Ex. 2112 ¶¶ 133–146. For example, Dr. Homayoun testifies that in addition to the technical background set forth in Dr. Stone’s definition, a person of ordinary skill in the art would have “considered all of the state of the art [described in Dr. Homayoun’s declaration] in the design of computer architecture, . . . [i]ncluding the issues of reconfigurable programming, processor speed, FPGA speed, and cost/benefit analysis of overhead introduction as applied to [high performance computing (HPC)] applications.” *Id.* ¶ 133. Also, according to Dr. Homayoun, an ordinarily skilled artisan would have “consider[ed] the technical problems [the ’800 patent was] attempting to solve” without using the patent “as a technical road-map to then conflate the technical problem with the solution,” and “would have considered the drawbacks in HPC computing and the deficiencies in FPGA systems and computer architecture design considerations at the time of the disclosed inventions.” *Id.* ¶ 140. Patent Owner similarly argues in its Response that a person of ordinary skill in the art would have “considered” all of these issues. PO Resp. 18–30.

Appendix E

We have evaluated all of Patent Owner’s arguments and supporting evidence regarding what a person of ordinary skill in the art allegedly would have considered when reading the asserted references. Patent Owner’s arguments pertain more to its criticism of Dr. Stone’s analysis as allegedly failing to understand the problems solved by the ’800 patent and being based on “hindsight bias” than a dispute over the “level” of ordinary skill in the art. *See id.* (also arguing that “Petitioner essentially uses the patent itself as a roadmap for stitching together various prior art references”); Sur-Reply 1–8; Tr. 64:4–65:15. It suffices at this point to conclude that a person of ordinary skill in the art would have had the *technical education and work experience* set forth in Dr. Stone’s declaration (and agreed to by Dr. Homayoun). *See* Tr. 64:13–14 (Patent Owner stating that there is no dispute as to “the level of education”). Among other things, such an individual would have had *knowledge* of “multi-adaptive processing systems and techniques,” “parallel processing,” and field-programmable gate arrays (FPGAs). *See* Ex. 1005, col. 1, ll. 37–61; Ex. 1003 ¶ 45; Ex. 2029 ¶ 17; Ex. 2112 ¶¶ 107–119, 133; Ex. 1007, 1–9; *Custom Accessories, Inc. v. Jeffrey-Allan Indus., Inc.*, 807 F.2d 955, 962 (Fed. Cir. 1986) (“The person of ordinary skill is a hypothetical person who is presumed to be aware of all the pertinent prior art.”). What that individual would have considered in evaluating particular prior art references and making potential combinations, though, is an issue we address below in evaluating Petitioner’s grounds of unpatentability.⁷

7. We note that for Petitioner’s anticipation ground based on Splash2, the level of ordinary skill in the art is relevant to

Appendix E

Based on the full record developed during trial, including our review of the '800 patent and the types of problems and prior art solutions described in the '800 patent, as well as the sophistication of the technology described in the '800 patent, we conclude that a person of ordinary skill in the art would have had (1) an advanced degree in electrical or computer engineering, or computer science with substantial study in computer architecture, hardware design, and computer algorithms, and at least three years of experience working in the field, or (2) a bachelor's degree covering those disciplines and at least four years working in the field. We apply that level of skill for purposes of this Decision.

C. Claim Interpretation

In this proceeding, we interpret the claims of the unexpired '800 patent using the “broadest reasonable construction in light of the specification of the patent.”

understanding how a person of ordinary skill in the art would have interpreted the claims and understood the reference, but not with respect to, for example, any motivations or problems a person of ordinary skill in the art would have had in combining teachings. See *Wasica Finance GmbH v. Continental Auto. Sys., Inc.*, 853 F.3d 1272, 1284 (Fed. Cir. 2017) (“Anticipation is an inquiry viewed from the perspective of one skilled in the art.”); *Vivint, Inc. v. Alarm. com Inc.*, 741 F. App'x 786, 791–792 (Fed. Cir. 2018) (unpublished) (“While anticipation is proven based on the express and inherent teachings of a single prior art reference, an obviousness analysis reaches beyond the prior art reference and takes into account other considerations such as the level of ordinary skill in the art and any objective indicia of nonobviousness.”).

Appendix E

37 C.F.R. § 42.100(b) (2017).⁸ Under this standard, we interpret claim terms using “the broadest reasonable meaning of the words in their ordinary usage as they would be understood by one of ordinary skill in the art, taking into account whatever enlightenment by way of definitions or otherwise that may be afforded by the written description contained in the applicant’s specification.” *In re Morris*, 127 F.3d 1048, 1054 (Fed. Cir. 1997); see *In re Smith Int’l, Inc.*, 871 F.3d 1375, 1382–83 (Fed. Cir. 2017) (“[The] broadest reasonable interpretation . . . is an interpretation that corresponds with what and how the inventor describes his invention in the specification.”). “Under a broadest reasonable interpretation, words of the claim must be given their plain meaning, unless such meaning is inconsistent with the specification and prosecution history.” *TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1062 (Fed. Cir. 2016). Our interpretation “‘cannot be divorced from the specification and the record evidence,’ and ‘must be consistent with the one that those skilled in the art would reach.’ A construction that is ‘unreasonably broad’ and which does

8. The Petitions in this proceeding were filed on September 5, 2018, prior to the effective date of the rule change that replaces the broadest reasonable interpretation standard with the federal court claim interpretation standard. See Changes to the Claim Construction Standard for Interpreting Claims in Trial Proceedings Before the Patent Trial and Appeal Board, 83 Fed. Reg. 51,340, 51,340, 51,358 (Oct. 11, 2018) (amending 37 C.F.R. § 42.100(b) effective November 13, 2018) (now codified at 37 C.F.R. § 42.100(b) (2019)); *Game & Tech. Co. v. Wargaming Grp. Ltd.*, 942 F.3d 1343, 1351 (Fed. Cir. 2019) (“If, as here, the [*inter partes* review] stems from a petition filed before November 13, 2018, the claims are given the ‘broadest reasonable interpretation’ consistent with the specification.” (citation omitted)).

Appendix E

not ‘reasonably reflect the plain language and disclosure’ will not pass muster.” *Microsoft Corp. v. Proxyconn, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015) (citations omitted), *overruled on other grounds by Aqua Prods., Inc. v. Matal*, 872 F.3d 1290 (Fed. Cir. 2017).

Upon review of the parties’ arguments during trial and the evidence as a whole, we conclude that three terms in claim 1 (“data driven,” “computational loop,” and “pass computed data seamlessly between said computational loops”) and one term in claim 15 (“establishing a stream communication connection between functional units”) require interpretation.

1. “*Data Driven*”

Claim 1 recites “transforming an algorithm into a *data driven* calculation that is implemented by said reconfigurable computing system at the at least one reconfigurable processor” (emphases added). Petitioner argued in the Petition that “data driven” should be interpreted to mean “the scheduling of operations upon the availability of their operands.” Pet. 11. We preliminarily adopted that interpretation in the Decision on Institution. Dec. on Inst. 17–18. Patent Owner argues that Petitioner’s proposed interpretation is “highly problematic” because it “introduces the extraneous limitation of ‘passing data directly.’” PO Resp. 34–45. Although Petitioner proposes such a limitation for the phrase “pass computed data seamlessly between said computational loops,” it does not do so for the term “data driven.” *See infra* Section II.C.3. Applying the broadest reasonable interpretation of the

Appendix E

claims in light of the Specification, we interpret “data driven” to mean “the scheduling of operations upon the availability of their operands.”

2. “*Computational Loop*”

Claim 1 recites “utilizing a first of said formed functional units to operate upon a subsequent data dimension of said calculation forming a *first computational loop*” and “substantially concurrently utilizing a second of said formed functional units to operate upon a previous data dimension of said calculation generating a *second computational loop*” (emphases added).

In the Decision on Institution, we preliminarily interpreted “computational loop” in claim 1 to mean “a set of computations that is executed repeatedly, either a fixed number of times or until some condition is true or false.” Dec. on Inst. 18–19. Petitioner agrees with that interpretation. Reply 31. Patent Owner states that it “does not disagree with” the preliminary interpretation, but “proposes a small clarification to reflect the plain and ordinary meaning of a loop within the field of high-performance computing and in particular in the context of the ’800 Patent.” PO Resp. 65. Specifically, Patent Owner proposes the following interpretation: “a set of computations that is executed repeatedly *per datum*, either a fixed number of times or until some condition is true or false.”⁹ *Id.* (emphasis added).

9. Patent Owner’s proposed interpretation in its Preliminary Response—“a sequence of computations that is repeated until a prescribed condition is satisfied”—did not include a “per datum”

Appendix E

Patent Owner in its Response cites three technical dictionary definitions of “loop,” upon which we had based the preliminary interpretation as reflective of the plain meaning of the term. *See id.* at 66; Dec. on Inst. 18–19. None of those definitions, however, requires that a set of instructions or computations be executed repeatedly “per datum.” *See* Ex. 2024, 4 (“[a] sequence of instructions that is repeated until a prescribed condition, such as agreement with a data element or completion of a count, is satisfied”); Ex. 2025, 5 (“a series of instructions being carried out repeatedly until a terminal condition prevails” or “[a] sequence of computer instructions that repeats itself until a predetermined count or other test is satisfied, or until the process is interrupted by operator intervention”); Ex. 2026, 8 (“[a] set of statements in a program executed repeatedly, either a fixed number of times or until some condition is true or false”); *see also* Ex. 2038, 3 (defining “computation” as “an act, process, or method of computing”). Indeed, Patent Owner cites the definitions in asserting that “[a] ‘computational loop’ is an iterative sequence of computations that repeats until a prescribed condition is satisfied” (without mentioning any “per datum” requirement). PO Resp. 76.

The Specification of the ’800 patent is consistent with the technical dictionary definitions, and does not require that the set of computations executed repeatedly as a “loop” be on a particular piece of data. *See* Ex. 1005, col. 6, l. 67–col. 7, l. 3 (explaining that computational process 610 “loops over the depth slices” and “loops over the shots” of a seismic imaging application), Fig. 7A (depicting three

requirement. Paper 15, 21–22.

Appendix E

loops each performing computations a particular number of times as “ $k = 1, nz$ ”; “ $j = 1, ny$ ”; and “ $i = 1, nx$ ”), Fig. 8B (depicting two loops as “ $i = 1, l$ ” and “ $k = 1, m$ ”). Patent Owner does not cite—and we do not find—any specific support in the Specification (including a reference purportedly incorporated by reference in the ’800 patent (Ex. 2037, “Caliga”)) for imposing a “per datum” requirement. *See* PO Resp. 66–67 (citing Ex. 1005, col. 5, l. 65–col. 6, l. 28, col. 6, l. 46–col. 7, l. 37, col. 8, ll. 21–39, Figs. 4A–4B, 6B–6G, 7A, 8A–8B; Ex. 2037, 4, 5, 7, 12, 13, 16–19). Indeed, the Specification and Caliga do not use the term “datum.” Caliga, in fact, describes the opposite of what Patent Owner proposes, namely a “loop” of a set of computations executed repeatedly using *different* data. *See* Ex. 1076 ¶¶ 5–8; Ex. 2037, 5 (describing a “[l]oop over filter coefficients” where index j is incremented and different values for “Data”, “ R_j ,” and “ R_{j+1} ” are used at each iteration); Reply 32.¹⁰

Patent Owner further cites as support U.S. Patent No. 8,589,666 B2 (Ex. 2027, “the ’666 patent”), which was originally assigned to the same original assignee as the ’800 patent (SRC Computers, Inc.) but is unrelated to the ’800 patent. PO Resp. 67–68. The cited portions of the ’666

10. We do not agree with Patent Owner that Petitioner’s Reply and Dr. Stone’s reply testimony are improper, as they respond directly to Patent Owner’s arguments in the Response regarding claim interpretation, where Patent Owner argued a “per datum” requirement for the first time. *See* Sur-Reply 8–9; Reply 32; Ex. 1076 ¶¶ 5–8; 37 C.F.R. § 42.23(b) (“A reply may only respond to arguments raised in the corresponding . . . patent owner response.”).

Appendix E

patent generally describe a “loop body” with “new data” (plural) “fed in on every clock tick,” and do not describe a “loop” as requiring a set of computations to be executed repeatedly “per datum” (singular). *See* Ex. 2027, col. 2, l. 64–col. 3, l. 23, col. 6, ll. 6–28. Dr. Homayoun also testifies that “[a] computational loop evaluates each piece of data multiple times, ‘a fixed number of times or until some condition is true or false,’” but does not explain the basis for that interpretation, in the Specification of the ’800 patent or otherwise. Ex. 2112 ¶ 207.

Finally, in its Sur-Reply and at the oral hearing, Patent Owner relied on Figure 4B of the ’800 patent as support for its proposed interpretation, arguing that the figure “unambiguously describes . . . two loops A and B operating a number of times (looping) on *each* dimension of data” and that our preliminary interpretation “would exclude from the claims the ’800 Patent’s embodiments and figures.” *See* Sur-Reply 15–17; Tr. 52:20–58:21, 82:8–84:2. Patent Owner also stated during the oral hearing that “per datum” in its proposed interpretation could be replaced with “per dimension,” as Patent Owner did not intend there to be a “difference” between the terms. Tr. 53:10–11, 55:3–10, 82:8–17.

We disagree for two reasons. First, claim 1 already recites two functional units operating on two data dimensions forming two computational loops. A first functional unit “operate[s] upon a subsequent data *dimension* of said calculation forming a first computational loop” and a second functional unit concurrently “operate[s] upon a previous data *dimension* of said calculation

Appendix E

generating a second computational loop” (emphases added). Dependent claims 2–5 specify different options for what those dimensions may comprise, namely “multiple vectors,” “multiple planes,” “multiple time steps,” or “multiple grid points” in the calculation. To the extent Patent Owner’s proposed interpretation would require that a functional unit operate on a particular data “dimension” of the calculation forming a computational loop, that is already encompassed in the claim language.

Second, the Specification indicates that a data “dimension” is not the same thing as an individual “datum.”

Figure 4B of the '800 patent is reproduced below.




Fig. 4B

Appendix E

Figure 4B depicts “multi-dimensional process 410 . . . effectuated such that multiple dimensions of data are processed by both Loops A (first loop 412) and B (second loop 414) such that the computing system logic is operative on every clock cycle.” Ex. 1005, col. 6, ll. 11–18. An application

will “pass” a subsequent dimension of a given problem through the first loop 412 of logic concurrently with the previous dimension of data being processed through the second loop 414. In practice, a “dimension” of data can be: *multiple vectors of a problem, multiple planes of a problem, multiple time steps in a problem and so forth.*

Id. at col. 6, ll. 22–28 (emphasis added). The Specification uses the phrasing “dimension” “of data” or “dimension” “of a given problem,” not “datum” individually, and provides examples of possible dimensions that mirror the language of dependent claims 2–5. It does not describe loop calculations applied to the same individual “datum” repeatedly.¹¹

11. The parties do not refer to anything in the prosecution history as supporting their proposed interpretations of “computational loop.” During prosecution of the parent ’324 patent, the applicants initially proposed during an examiner interview amending its similarly worded claim 1 to recite that “said first of said instantiated functional units and said second of said instantiated functional units are within a nested loop of said calculation.” Ex. 1002, 208–209. The applicants then instead amended the claim to recite, rather than a “nested loop,” first

Appendix E

Having reviewed all of the cited evidence, we do not find sufficient support for limiting the term “computational loop” to require that computations be executed repeatedly “per datum.” Applying the broadest reasonable interpretation of the claims in light of the Specification, we interpret “computational loop” to mean “a set of computations that is executed repeatedly, either a fixed number of times or until some condition is true or false.”

3. “Pass Computed Data Seamlessly Between Said Computational Loops”

Claim 1 recites that “said implementation of said calculation enables said first computational loop and said second computational loop execute¹² concurrently and pass computed data seamlessly between said computational loops.”

Petitioner argues that “pass computed data seamlessly between said computational loops” in claim 1 means “communicate computed data directly between functional units that are calculating computational loops.” Pet. 15–16; Reply 20. Based on the record at the time, we preliminarily agreed with Petitioner’s proposed interpretation in the Decision on Institution. Dec. on Inst. 20–22. Patent Owner

and second “computational loops” executing concurrently with computed data seamlessly passed between them. *Id.* at 214–215.

12. We read the claim language to mean that the implementation enables the loops “to” execute concurrently and pass computed data seamlessly between the computational loops.

Appendix E

argues that the phrase instead should be interpreted to mean “communicating the computed data over the reconfigurable routing resources.” PO Resp. 35.

According to the plain language of the claim, which recites “pass[ing] computed data seamlessly between said computational loops,” “seamlessly” refers to how computed data is passed *between* the first and second computational loops (performed by the first and second functional units of the reconfigurable processor)—not, for example, how computed data is passed from either of the computational loops to any other component unrelated to the functional units, or vice versa.

The only other time “seamlessly” appears in the ’800 patent is in independent claims 25 and 51, which mirror the language of claim 1, reciting that computed data is passed “seamlessly” between systolic walls or columns of a calculation. The written description does not use the terms “seam” or “seamlessly.” Figure 2, though, shows functional units 204 interconnected without any intervening structures between them, which is consistent with Petitioner’s proposed interpretation. *See* Ex. 1005, col. 5, ll. 39–45. Figures 7A–7B and 8A–8B similarly show direct communication of data from one computational loop to another. *See id.* at col. 7, ll. 38–67 (loops 702, 704, and 706), col. 8, ll. 27–55 (loops 812 and 814). Although the description of the figures does not use the terms “seam” or “seamlessly,” a person of ordinary skill in the art would have understood the claim language in light of that description, given that it describes and shows the interaction between multiple “loops,” as recited in claim 1. *See id.*

Appendix E

The prosecution history of the parent '324 patent also supports Petitioner's proposed interpretation. *See Teva Pharms. USA, Inc. v. Sandoz, Inc.*, 789 F.3d 1335, 1343 (Fed. Cir. 2015) ("A statement made during prosecution of related patents may be properly considered in construing a term common to those patents, regardless of whether the statement pre- or post-dates the issuance of the particular patent at issue."); PO Resp. 37 n.10 ("The '800 Patent is a continuation of the '324 Patent, so the file histories for both are relevant."). Following an Office Action rejecting similarly worded claim 1 under 35 U.S.C. § 112, first paragraph, and 35 U.S.C. § 103, the applicants conducted an interview with the examiner and filed an Office Action response amending the claim in a number of ways, including adding the "pass computed data seamlessly between said computational loops" limitation. Ex. 1002, 214–215. In response to the § 112, first paragraph, written description rejection, the applicants indicated that data from one functional unit performing a computational loop is passed directly to another functional unit:

Data flows across the array between functional units, usually with different data flowing in different directions. . . . [I]n the Applicant's invention Systolic implementation will connect computational loops such that data from one compute loop will be passed as input data to a concurrently executing compute loop. In the Applicant's invention data computed by computation units or groups of functional units flows seamlessly and concurrently with data being computed by other groups of

Appendix E

functional units. Thus, the process claimed by the Applicant therefore significantly increases the computing processes taking place in a reconfigurable processor.

Ex. 1002, 226. Although the language above primarily relates to the term “[s]ystolic” (which is present in claim 1 of the ’324 patent but not claim 1 of the ’800 patent), we also find it significant for purposes of interpreting the “seamlessly” phrase because it refers to the limitation expressly in describing “Applicant’s invention.” *See id.*; Ex. 1003 ¶ 76; *Shire Dev., LLC v. Watson Pharms., Inc.*, 787 F.3d 1359, 1366 (Fed. Cir. 2015) (holding that even where “prosecution history statements do not rise to the level of unmistakable disavowal, they do inform the claim construction”).

Finally, Petitioner’s proposed interpretation is consistent with the testimony of Mr. Huppenthal, one of the named inventors of the ’800 patent, who describes “interpret[ing] standard high level language program constructs, such as the fact that the output variable from one loop is the input to another, and then implement[ing] this as a *storage free, seamless connection between the two loops* implemented on the FPGA,” and states:

By seamlessly I mean that the results of one loop streamed from that loop’s output to the input of the next loop *without being placed in a circuit element that required explicit address based on read or write operations such as a data register or memory* or through a switch

Appendix E

that requires additional non-data content for routing purposes.

Ex. 2101 ¶ 78 (emphases added).

Based on our review of the full record after trial, Patent Owner's arguments regarding the "seamlessly" language in claim 1 are not persuasive. *See* PO Resp. 35–43; Sur-Reply 20–21. First, Patent Owner relies on portions of the Specification describing interaction of cells "at the boundary," "the problem of passing data over numerous boundaries (or seams) between processing elements in typical multi-processor systems," and the '800 patent's solution to that problem of "staying on a single FPGA chip, effectively eliminating the associated boundaries or seams from chip-to-chip communication." PO Resp. 36–37, 40 (citing Ex. 1005, col. 2, ll. 26–49) (emphasis omitted); *see* Ex. 2112 ¶¶ 160–162. The cited portions of the Specification do not use the terms "seam" or "seamlessly" and, unlike the portions cited by Petitioner describing Figures 7A–7B and 8A–8B, do not describe interaction between "loops" in the disclosed method. We find them less relevant than the portions discussed above, particularly given that the claim language at issue recites passing computed data seamlessly "between" the computational loops (performed by the functional units of the reconfigurable processor).

Second, Patent Owner relies on certain statements made by the applicants during prosecution of the parent '324 patent. PO Resp. 37–40; *see* Ex. 2112 ¶¶ 163–166. Three of those statements mirror the Specification

Appendix E

language cited by Patent Owner and, importantly, predate the addition of the “seamlessly” language to claim 1 (and likewise do not use the terms “seam” or “seamlessly”). *See* Ex. 1002, 117–118, 148–150, 174–175. In addition, Patent Owner cites the following statement by the applicants in the Office Action response discussed above:

The instantiation of the at least one reconfigurable processor with at least two functional units enables each functional unit to communicate with each other. Certainly communication between other reconfigurable processors within the system would require [a] communication protocol but communication between functional units within an individual reconfigurable processor is free of such a requirement. To alleviate any confusion, the reference to the term “protocol” has been replaced with an “interconnection” between functional units that is established by reconfigurable routing resources inside each chip.

Id. at 224–25; *see* PO Resp. 38–40. The cited language pertains to the “instantiating” step recited earlier in the claim, which also was amended in the response, including replacing “communications . . . independent of external and internal communication protocols” with “interconnects . . . based on reconfigurable routing resources.” *See* Ex. 1002, 214, 224–25. By contrast, the applicants added the “seamlessly” limitation to specify the interaction between the computational loops (i.e., that computed data is passed “seamlessly” between them).

Appendix E

Third, Patent Owner disputes the “directly” aspect of Petitioner’s proposed interpretation because “it would exclude standard FPGAs . . . since standard FPGAs contain reconfigurable routing resources (comprising buffers and switches) between the configurable logic blocks,” and thus would “exclude the very embodiments of the ’800 Patent.” PO Resp. 40–41 (citing Ex. 1035, 31; Ex. 2077, 19–29, 32–34, 37–41, 46–51, 59–65). We agree with Petitioner that “[j]ust because a standard FPGA may include memories does not mean that when functional units are instantiated within such an FPGA that the memories are necessarily placed *between* functional units.” See Reply 25. Claim 1 recites two functional units being formed and concurrently performing two computational loops, with computed data passed seamlessly between the computational loops. Petitioner’s proposed interpretation of the “seamlessly” phrase correctly pertains to how computed data is passed *between* the computational loops when so formed, and is not inconsistent with the mere existence of memory on an FPGA.

Fourth, Patent Owner argues that the word “directly” makes Petitioner’s proposed interpretation unreasonable because it “introduce[s] ambiguity and confusion.” PO Resp. 36–37, 41–42; see Ex. 2112 ¶¶ 169–176. According to Patent Owner, Dr. Stone on cross-examination could not identify what “intervening” structures would prevent direct communication, and contradicted himself by stating that “an intervening memory would not be a direct connection” and later that “if a register . . . were between the two processing elements then the connection would still be direct, but if the intervening structure were a

Appendix E

buffer then the connection would not be direct.” PO Resp. 42 (citing Ex. 2063, 86:13–91:24). Patent Owner argues that if Dr. Stone’s view is adopted, “[t]he same circuit would be both direct and indirect, depending on where the boundaries of the ‘processing element’ are arbitrarily drawn with respect to intervening structures.” Sur-Reply 20–21.

We disagree that the word “directly” in Petitioner’s proposed interpretation is unclear or that Dr. Stone’s testimony about it introduces ambiguity. Dr. Stone testified as follows:

Q. Okay. And you mention the word “directly,” it was passing data directly between processing elements. What does that phrase mean to you or what’s the context? What are you trying to describe there?

A. That the data goes from first to the second *without going to something intervening*. It directly go – is connected immediately. Indirectly we – you go through one or more intervening places to get there.

Q. Okay. So would memory, if the data was going from one processing element to memory and then back to a processing element, is that something you would consider as an intervening thing?

Appendix E

A. Well, that would not be a direct connection of the output of the cell to the next cell. It says, “Between processing elements you’re directly connected.” *If you’re saying you have a processing element outputting to memory and then coming back to another processing element, that would not be direct.*

Ex. 2063, 85:14–86:12 (emphases added). When asked about a register in particular, Dr. Stone testified as follows:

Q. Well, how about a – a register? Would that be an intervening structure?

A. I – I’m puzzled because that – that register would be within – *within the processing element in my mind.*

Q. Okay.

A. If it’s *within the processing element as a register*, yeah, I would put it there, then *the output of that register, if it’s connected directly to the input of the next processing element, would be direct.*

Id. at 86:19–87:5 (emphases added); *see also id.* at 87:23–25 (“If the register is part of the processing element, then the connection would be direct.”).

Appendix E

Data is not communicated “directly” between processing elements when it is communicated through an intervening structure *between* them. *See id.* at 85:14–91:24. Examples of such a structure are memory and another processing element, but logically others are possible as well depending on how they are situated with respect to the processing elements. *See id.* at 86:13–18. When a register is *within* a processing element, however, data can still be communicated “directly” (provided there is a direct connection between the processing elements), as Dr. Stone explained. *See id.* at 85:14–91:24. We do not see any ambiguity in making this distinction. The boundaries of a processing element are not arbitrary as Patent Owner contends, but rather, in the context of a particular written document, would depend on how the reference describes the processing element and the communication of data to and from the processing element. *See, e.g.,* Ex. 1005, Fig. 2 (depicting functional units 204 of adaptive processor chip 202). Logically, also, in an arrangement with multiple processing elements, there must be some division between the processing elements, otherwise there would not be multiple elements. *See* Tr. 14:1–5.

Applying the broadest reasonable interpretation of the claims in light of the Specification, we interpret “pass computed data seamlessly between said computational loops” to mean “communicate computed data directly between functional units that are calculating computational loops.”

*Appendix E**4. “Establishing a Stream Communication Connection Between Functional Units”*

Claim 15 depends from claim 1 and recites that “instantiating includes establishing a stream communication connection between functional units” (the “stream communication” limitation). Petitioner argues that “stream communication” should be interpreted to mean “communication of a data sequence.” Pet. 17–19. Patent Owner argues that “stream communication” means “a data path that acts like a queue connecting via the reconfigurable routing resources a producer and a consumer of data that operate concurrently.” PO Resp. 45–65. We did not preliminarily interpret “stream communication” in the Decision on Institution.

We begin with the language of the claim. Claim 15 recites that “instantiating” includes establishing a stream communication connection “between functional units.” As the parties do, we read “instantiating” in claim 15 to refer to the “forming” step of parent claim 1¹³ (emphasis added):

forming at least two of said functional units at the at least one reconfigurable processor to perform said calculation wherein only functional units needed to solve the calculation are formed and wherein each formed functional unit at the at least one reconfigurable processor

13. The parent ’324 patent includes an “instantiating” step in claim 1 very similar to the language of the “forming” step of claim 1 of the ’800 patent. Claim 15 is identical in both patents.

Appendix E

interconnects with each other formed functional unit at the at least one reconfigurable processor based on *reconfigurable routing resources within the at least one reconfigurable processor as established at formation*

Thus, according to the plain language of the claim, what is being “form[ed]” is the at least two “functional units.” The “reconfigurable routing resources,” based on which the formed functional units are “interconnect[ed],” are established “at” the formation.

Turning to the language of claim 15, regardless of whether “stream communication” is used as an adjective (as Petitioner contends) or a noun (as Patent Owner contends), the term appears immediately before and modifies the phrase “connection between functional units.” *See* Reply 26–27; Sur-Reply 10 & n.2. The claim, therefore, requires establishing a connection of a particular type, namely a “stream communication” type of connection. Nothing further in the language of the claim limits or sheds light on what that type of connection entails.

Nor does the Specification. The only other time “stream communication” appears in the ’800 patent is in claim 40, which mirrors the language of claim 15. The written description never uses the term.¹⁴ It also does

14. Petitioner argues that its proposed interpretation is consistent with a portion of the Specification describing systolic wall of computation 712 in Figure 7C and stating that “communication of values between adjacent rows . . . in the virtual wall can occur without storing values to memory.” Ex. 1005, col. 8,

Appendix E

not differentiate between different types of connections between functional units.

Both parties, however, refer to patents that the '800 patent purports to incorporate by reference. *See* Pet. 17; PO Resp. 57–61. The '800 patent states that it is “related to the subject matter of,” and “incorporate[s] in [its] entirety by . . . reference,” U.S. Patent No. 6,434,687 B1 (Ex. 1014, “the '687 patent”).¹⁵ Ex. 1005, col. 1, ll. 9–22. Petitioner points to the phrase “streams of operands” in the '687 patent as supporting its proposed interpretation. Pet. 17. The cited portion describes multi-adaptive (MAP) processor 112 with reconfigurable array 42 having one or more high performance FPGAs. Ex. 1014, col. 9, ll. 1–6. “After configuration, the user array 42 can perform whatever function it was programmed to do.” *Id.* at col. 9, ll. 11–12. “In order to maximize its performance for vector processing, the array 42 should be able to access two streams of operands simultaneously” by using a chain port “connector allow[ing] the MAP element 112 to use data provided to it by a previous MAP element 112.” *Id.* at col. 9, ll. 12–26.

We find the cited portion to be of limited relevance. The '687 patent uses “streams of operands” rather than “stream communication” and does not describe the “streams of operands” other than the single sentence

ll. 1–6; *see* Pet. 17–19; Ex. 1003 ¶ 95. The cited portion, however, never uses the term “stream communication.”

15. Petitioner challenges the '687 patent in Case IPR2018-01594.

Appendix E

above. Also, given the length of the '687 patent disclosure, it is unclear exactly what aspects of the '687 patent are incorporated in the '800 patent or how they would apply in the context of the '800 patent invention. *See* Ex. 1005, col. 1, ll. 9–22 (merely stating that the '800 patent is “related to the subject matter of” the '687 patent). If anything, the cited portion of the '687 patent simply indicates that certain data (i.e., “operands”) can be streamed from one MAP element to another.

Patent Owner similarly quotes other excerpts from the '687 patent describing the use of a chain port connection, input buffer 40, and output First-In-First-Out (FIFO) buffer 74 in support of its contention that “stream communication” requires a “data path” acting like a “queue.” PO Resp. 58–61 (citing Ex. 1014, col. 7, ll. 43–57, col. 8, ll. 7–26, col. 9, ll. 2–67); *see* Ex. 1014, Figs. 6–7. We are not persuaded for the same reasons as Petitioner’s citation. The cited portions merely describe an exemplary embodiment of the '687 patent that uses memory buffers. Other than the single use of “streams of operands,” there is no connection between the '687 patent disclosure and anything pertaining to streaming. *See* Sur-Reply 12 (acknowledging that the '687 patent is “a different patent referring to a different invention pertaining to internet communications”). Certainly, the cited portions do not use or define what is meant by “establishing a stream communication connection between functional units” in the context of the data processing method recited in claim 15.

Appendix E

Patent Owner also relies on U.S. Patent No. 6,339,819 B1 (Ex. 2084, “the ’819 patent”), of which the ’687 patent is a continuation-in-part, as incorporated by reference in the ’800 patent. PO Resp. 58–59. The ’800 patent describes “a representative systolic wavefront operation in the form of a reservoir simulation function” where

the computation of fluid flow properties are communicated to neighboring cells 710 and, importantly, this computation can be scheduled to eliminate the need for data storage. In accordance with the technique of the present invention, a set of cells can reside in an adaptive processor and the pipeline of computation can extend across multiple adaptive processors. Communication overhead between multiple adaptive processors may be advantageously minimized through the use of MAP™ adaptive processor chain ports as disclosed in U.S. Pat. No. 6,339,819 issued on Jan. 15, 2002 for: “Multiprocessor With Each Processor Element Accessing Operands in Loaded Input Buffer and Forwarding Results to FIFO Output Buffer,” assigned to SRC Computers, Inc., assignee of the present invention, the disclosure of which is herein specifically incorporated by this reference.

Ex. 1005, col. 7, ll. 42–46, col. 7, l. 59–col. 8, l. 6, Figs. 7A–7B. According to Patent Owner, the Specification thereby discloses “the concept of using chain ports and a FIFO buffer for chip to chip communications,” such that

Appendix E

a person of ordinary skill in the art “would recognize that this inter-chip communication concept can be adapted to intra-chip communications between functional units on the same chip—e.g., a FIFO within the chip.” PO Resp. 58–61. The cited disclosure is of a “representative” (i.e., exemplary) embodiment where communication “may” occur “between multiple adaptive processors” using chain ports and buffers. *See* Ex. 1005, col. 7, ll. 38–42, col. 7, ll. 55–65. Again, there is no connection to anything pertaining to streaming data between functional units, as the disclosure does not use the word “stream” at all. The incorporated ’819 patent also includes much of the same disclosure as the ’687 patent and suffers from the same problems explained above. *Compare* Ex. 1014, col. 4, l. 15–col. 20, l. 35, Figs. 1–11B, *with* Ex. 2084, col. 4, l. 31–col. 21, l. 43, Figs. 1–11B. We are not persuaded that the incorporated patents support Patent Owner’s contention that “stream communication” in claim 15 requires a “data path” acting like a “queue.”

Turning next to the prosecution history, we find some guidance as to the scope of “stream communication.” During prosecution of the parent ’324 patent, claim 15 was amended to recite “stream communication” in the applicants’ first Office Action response. Ex. 1002, 110. The applicants, however, did not discuss the language of claim 15 in the response and instead attempted to distinguish the cited prior art as applied to parent claim 1. *Id.* at 116–122. Patent Owner points to various statements later in the prosecution history, only one of which addressed the “stream communication” language of claim 15 directly. *See* PO Resp. 61–65 (citing Ex. 1002, 128–129, 147–150, 174–

Appendix E

175, 208, 224–225). Specifically, in the subsequent Office Action, the examiner found the “stream communication” limitation taught by a particular reference, noting that the reference

taught minimiz[ing] interconnections of processing elements and the matrix and vector signal subsets are specifically formed so that they need to be inputted to only one row and one column[] and yet still [are] properly processing systolically along all dimensions within the array. . . . Consequently the stream of communication between functional units is established as the interconnections are made and data is transferred systolically in at least one stream between processors.

Ex. 1002, 128–129. This explanation indicates that establishing a stream communication connection allows data to be “transferred” as a “stream.” Because the applicants and examiner do not appear to have ever mentioned a “data path” acting like a “queue” when addressing the “stream communication” limitation, though, we are not persuaded that the prosecution history supports a narrower interpretation including those limitations, as Patent Owner contends.

Based on what information there is in the intrinsic record, in particular the claim language surrounding “stream communication” and the single reference to the “stream communication” limitation in the prosecution history, we conclude that the plain meaning of “establishing

Appendix E

a stream communication connection between functional units” is establishing a connection over which data is streamed between functional units. The interpretation gives meaning to the “stream communication” term and results in claim 15 appropriately having a different scope than claim 1, as claim 15 requires data to be communicated over a particular type of connection that permits streaming, whereas claim 1 would encompass any type of connection. We do not find any support in the intrinsic record for Patent Owner’s proposed interpretation requiring a “data path” acting like a “queue.”

Patent Owner points to a number of sources of extrinsic evidence in support of its narrower interpretation. As an initial matter, none of those references is a technical dictionary, textbook, or similar source that typically are referenced to help determine a term’s ordinary meaning. *See* Tr. 37:15–21, 38:10–17. Instead, Patent Owner refers to unrelated patents and other types of documents. *See* PO Resp. 49–57; Sur-Reply 12–14. For example, Patent Owner cites the ’666 patent, which discloses:

A stream is a data path between a producer and consumer of data, where the producer and consumer run concurrently. The path between the producer and consumer is made up of a data connection, a “valid” signal, and a reverse direction “stall” signal. FIG. 1 shows typical signals used in a stream connection as is well known and will be recognized by one skilled in the relevant art. The use of a First-In-First-Out buffer 110, or “FIFO” buffer, removes the

Appendix E

need for tight synchronization between the producer 120 and consumer 130. The producer 120 will generate data values 125 at its own rate, allowing them to accumulate in the FIFO buffer 110. As the FIFO buffer 110 approaches becoming full, it will issue a stall signal 140 to the producer 120 so that it will suspend the generation of data values 125 until the stall signal is released. The consumer 130 will take 150 values 145 from the FIFO buffer at its own rate and as the values 145 are available.

Ex. 2027, col. 2, ll. 39–54 (emphases added); *see* PO Resp. 49–51; Sur-Reply 12. According to Patent Owner and Dr. Homayoun, this disclosure of using a FIFO buffer to account for different data rates and allow the producer and consumer to run concurrently is “entirely consistent with . . . instantiating reconfigurable [routing] resources to seamlessly communicate computed data between processing elements” in claim 1 and “provid[ing] stream communication between those processing elements” in claim 15. PO Resp. 51; *see* Ex. 2112 ¶¶ 155–156.

Although sharing an assignee, the ’666 patent is unrelated to the ’800 patent, has an effective filing date (July 10, 2006) nearly four years after that of the ’800 patent (October 31, 2002), and has a different named inventor. *See* Reply 28; Ex. 1013, 2–3. We also do not see any relationship between the ’666 patent’s description of streaming in connection with a different invention and the ’800 patent’s implementation. We are not persuaded, therefore, that the ’666 patent supports reading “stream

Appendix E

communication” in claim 15 to require a “data path” acting like a “queue” as Patent Owner contends, particularly when the intrinsic record provides some indication of a broader interpretation. “[E]xtrinsic evidence may be used only to assist in the proper understanding of the disputed limitation; it may not be used to vary, contradict, expand, or limit the claim language from how it is defined, even by implication, in the specification or file history.” *Bell Atlantic Network Servs., Inc. v. Covad Comm’cns Group, Inc.*, 262 F.3d 1258, 1269 (Fed. Cir. 2001).

Likewise, Patent Owner cites its own product documentation, dated 2002–2007, as “describ[ing] a stream as a data structure that allows flexible communication between concurrent producer and consumer loops” using an internal buffer, as well as a 1993 Argonne National Laboratory paper. PO Resp. 51–54 (citing Ex. 2108, 94–98; Ex. 2028, 31; Ex. 2101 ¶ 79; Ex. 2112 ¶ 157). We do not see—and Patent Owner does not explain sufficiently—why such product documentation shows how a person of ordinary skill in the art would have understood the language of the “stream communication” limitation in claim 15. Moreover, whereas Patent Owner’s proposed interpretation is for “stream communication” to be a “data path,” the product documentation and paper indicate that a “stream” is a “data structure.” *See id.* at 45; Ex. 2108, 94; Ex. 2028, 31. They are different concepts. *See* Ex. 2064, 154–155 (defining “data path” as “[t]he route that a signal follows as it travels through a computer network” and “data structure” as “[a]n organizational scheme, such as a record or array, that can be applied to data to facilitate interpreting the data or performing operations on it”).

Appendix E

Finally, Patent Owner quotes a number of other patents and patent application publications as allegedly showing that “stream communication” requires a “data path” acting like a “queue.” PO Resp. 55–57. The references use the words “stream” or “streaming” (not “stream communication connection”) in the context of describing embodiments of their own disclosed inventions. We find that they have little probative value as well.

We are not persuaded that the extrinsic evidence cited by Patent Owner shows that there was a uniform understanding of persons of ordinary skill in the art of the meaning of the term “stream communication,” alone or in the context of a “stream communication connection” between functional units. Indeed, one technical dictionary, submitted but not cited by either party in connection with the term “stream communication,” defines “stream” more broadly as “[a]ny data transmission, such as the movement of a file between disk and memory, that occurs in a continuous flow” (as a noun) or “[t]o transfer data continuously, beginning to end, in a steady flow” (as a verb). Ex. 2064, 509. This is consistent with Petitioner’s view that the word “stream” in claim 15 “indicates some kind of movement, some kind of transfer of data,” as opposed to a particular data structure as Patent Owner contends. *See* Tr. 71:5–15, 72:9–10. Notably, the same technical dictionary that defines “stream” as a “data *transmission*” also defines “queue” as “[a] multi-element data *structure* from which . . . elements can be removed only in the same order in which they were inserted.” Ex. 2064, 443, 509 (emphases added). This indicates that a “stream” and a “queue” are distinct concepts, and

Appendix E

contradicts Dr. Homayoun’s testimony that “[a] person of ordinary skill in the art would understand that the term ‘stream’ refers to a specific type of structure called a queue.” *See* Ex. 2112 ¶ 152.

Applying the broadest reasonable interpretation of the claims in light of the Specification, we give “establishing a stream communication connection between functional units” its plain meaning, namely “establishing a connection over which data is streamed between functional units.” We are not persuaded that the extrinsic evidence relied upon by Patent Owner supports a different interpretation. No further interpretation is necessary to resolve the parties’ disputes over the asserted grounds of unpatentability in this proceeding.¹⁶ *See Nidec Motor Corp. v. Zhongshan*

16. Although our interpretation differs from those proposed by the parties for “stream communication,” we did not preliminarily interpret the term in the Decision on Institution, both parties had the opportunity to present arguments regarding the term in their papers and at the oral hearing, and the issue was discussed extensively during the oral hearing. *See* Tr. 20:21–22:18, 33:19–41:14, 68:19–73:17, 84:3–86:13; *TQ Delta, LLC v. DISH Network LLC*, 929 F.3d 1350, 1355 (Fed. Cir. 2019) (holding that the Board did not improperly “change course” under *SAS Institute, Inc. v. ComplementSoft, LLC*, 825 F.3d 1341 (Fed. Cir. 2016), *rev’d and remanded on other grounds*, *SAS Institute, Inc. v. Iancu*, 138 S. Ct. 1348 (2018), by “construing the limitation in the Final Written Decision because it did not construe the term in its Decision to Institute”); *WesternGeco LLC v. ION Geophysical Corp.*, 889 F.3d 1308, 1328–29 (Fed. Cir. 2018) (holding that “the Board is not bound to adopt either party’s preferred articulated construction of a disputed claim term” and “was permitted to issue a new construction in the final written decision given that

Appendix E

Broad Ocean Motor Co., 868 F.3d 1013, 1017 (Fed. Cir. 2017) (“Because we need only construe terms ‘that are in controversy, and only to the extent necessary to resolve the controversy,’ we need not construe [a particular claim limitation] where the construction is not ‘material to the . . . dispute.’” (citations omitted)).

D. Legal Standards

“Anticipation requires that every limitation of the claim in issue be disclosed, either expressly or under principles of inherency, in a single prior art reference,” *Corning Glass Works v. Sumitomo Elec. U.S.A., Inc.*, 868 F.2d 1251, 1255–56 (Fed. Cir. 1989), and that the claim limitations be “arranged or combined in the same way as recited in the claim,” *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1371 (Fed. Cir. 2008). However, “the reference need not satisfy an *ipsissimis verbis* test.” *In re Gleave*, 560 F.3d 1331, 1334 (Fed. Cir. 2009). “In an anticipation analysis, the dispositive question is whether a skilled artisan would ‘reasonably understand or infer’ from a prior art reference that every claim limitation is disclosed in that single reference.” *Acoustic Tech., Inc. v. Itron Networked Solutions, Inc.*, 949 F.3d 1366, 1373 (Fed. Cir. 2020) (citation omitted). “Expert testimony

claim construction was a disputed issue during the proceedings”); *Intellectual Ventures II LLC v. Ericsson Inc.*, 686 F. App’x 900, 906 (Fed. Cir. 2017) (unpublished) (“The Board is not constrained by the parties’ proposed constructions and is free to adopt its own construction,” but “after the Board adopts a construction, it may not change theories without giving the parties an opportunity to respond.”).

Appendix E

may shed light on what a skilled artisan would reasonably understand or infer from a prior art reference.” *Id.*

A claim is unpatentable for obviousness if, to one of ordinary skill in the pertinent art, “the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007) (quoting 35 U.S.C. § 103(a)). The question of obviousness is resolved on the basis of underlying factual determinations, including “the scope and content of the prior art”; “differences between the prior art and the claims at issue”; and “the level of ordinary skill in the pertinent art.” *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). Additionally, secondary considerations, such as “commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. As indicia of obviousness or nonobviousness, these inquiries may have relevancy.” *Id.*

A patent claim “is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art.” *KSR*, 550 U.S. at 418. An obviousness determination requires finding “both ‘that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.’” *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367–68 (Fed. Cir. 2016) (citation omitted); *see KSR*,

Appendix E

550 U.S. at 418 (for an obviousness analysis, “it can be important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does”). “Although the *KSR* test is flexible, the Board ‘must still be careful not to allow hindsight reconstruction of references . . . without any explanation as to *how* or *why* the references would be combined to produce the claimed invention.’” *TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1066 (Fed. Cir. 2016) (citation omitted). Further, an assertion of obviousness “cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *KSR*, 550 U.S. at 418 (quoting *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)); accord *In re NuVasive, Inc.*, 842 F.3d 1376, 1383 (Fed. Cir. 2016) (stating that “conclusory statements” amount to an “insufficient articulation[] of motivation to combine”; “instead, the finding must be supported by a ‘reasoned explanation’” (citation omitted)); *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016) (“To satisfy its burden of proving obviousness, a petitioner cannot employ mere conclusory statements. The petitioner must instead articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness.”).

Appendix E

E. Anticipation Ground Based on Splash2 (Claims 1, 15, 18, 21, and 22)

1. Splash2

Splash2¹⁷ is a book describing the Splash 2 reconfigurable computer system created by the Supercomputing Research Center in the 1990s. Ex. 1007, xi. “Splash 2 is an attached processor system using Xilinx XC4010 FPGAs as its processing elements.” *Id.* “[T]he XC4010 contains a 20 x 20 array of Configurable Logic Blocks (CLBs)” *Id.* at 11.

Figure 2.3 of Splash2 is reproduced below.




FIGURE 2.3 Splash 2 System Architecture

17. The first 11 pages of Splash2 and a four-page summary article about the Splash 2 system were made of record during prosecution of the '800 patent, but not the remainder of the book, including Chapter 8 on which Petitioner primarily relies. *See* Ex. 1005, code (56); Ex. 1006, 49; Ex. 1060; Pet. 1-2. RaPiD, Jeong, Chunky SLD, and Roccatano were not of record during prosecution of the '800 patent. *See* Ex. 1005, code (56).

Appendix E

Figure 2.3 depicts the system architecture of the Splash 2 system, including a set of array boards connected to a SPARCstation 2 host via an interface board. *Id.* at 12–13. Figure 2.4 of Splash2 is reproduced below.

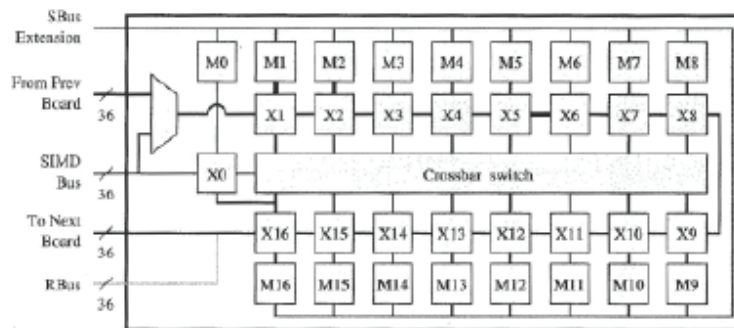


FIGURE 2.4 Array Board Architecture

Figure 2.4 depicts the architecture of one of the array boards. Each array board “contains 17 Xilinx XC4010 FPGA chips as its processing elements. Sixteen of these are connected in a linear array to create a linear data path and the seventeenth provides a broadcast capability to the other 16 chips.” *Id.* at 13 (citations omitted). “Viewed as a machine with a linear data path, the [Single Instruction, Multiple Data (SIMD)] Bus can be used to transmit data from the Interface Board to the first FPGA on the first Array Board. The data can then be moved through the linear data path on that board, then to the first FPGA on the second Array Board, and so on.” *Id.* at 14. The linear data path is bidirectional. *Id.*

Splash2 describes a number of applications programmed on the Splash 2 system. *Id.* at xi. Chapter 8 describes “two systolic array architectures for [genetic]

Appendix E

sequence comparison and their implementations on the Splash 2 programmable logic array.” *Id.* at 97. Splash2 discloses that “[i]n comparing two sequences, it is useful to quantify their similarity in terms of a distance measure,” and one such measure is the “edit distance between two sequences,” which is “the minimum cost of transforming one sequence to the other with a sequence of the following operations: deletion of a character, insertion of a character, and substitution of one character for another.” *Id.* at 98. Each operation has a cost, and the transformation cost is “the sum of the costs of the individual operations.” *Id.* Splash2 discloses a “well-known dynamic programming algorithm” for computing edit distance, and describes how two example nucleotide sequences *TCTAGACC* and *GCATAAGC* would be compared using the equations for that algorithm. *Id.* at 98–99.

Splash2 discloses that the edit distance algorithm has “inherent parallelism” in that “each entry in the distance matrix” for comparing each character of the source sequence to each character of the target sequence “depends on adjacent entries,” and such parallelism can be “exploited to produce systolic algorithms in which communication is limited to adjacent processors.” *Id.* at 98–100, Figs. 8.2–8.3. Specifically, entries of the matrix on the same antidiagonal can be computed in parallel. *Id.* at 100, Fig. 8.4. Splash2 describes two architectures for doing so, one using a bidirectional systolic array and one using a unidirectional systolic array. *Id.* at 100–108. Splash2 discloses that “[b]oth the bidirectional and unidirectional systolic arrays have been implemented on the Splash 2 programmable logic array, with versions for DNA and protein sequences.” *Id.* at 104.

Appendix E

a) Unidirectional Array Implementation

Figure 8.9 of Splash2 is reproduced below.




FIGURE 8.9 Data Flow through the Unidirectional Systolic Array. The source sequence is first loaded into the array. The target sequences are then streamed through the array. The tag acts as a simple instruction telling each PE how to process the incoming data. The SR tag instructs an empty PE to load the source character and distance from the input stream. The PR tag marks the end of the source stream. The TG tag signals a target character. Multiple source and target sequences can be carried on the input stream for uninterrupted pipelined processing.

As shown in Figure 8.9, “data flows through the unidirectional array in one direction. The source sequence is loaded once and stored in the array starting from the leftmost [processing element (PE)]. The target sequences are streamed through the array one at a time, separated by control characters.” Ex. 1007, 103. “In this configuration, each PE computes the distances in one row of the distance matrix. At each time step, the PEs compute the distances along a single antidiagonal in the distance matrix . . .” *Id.* at 104. Splash2 discloses that the “DNA version of the unidirectional array,” for example, has 248 processing elements that would be capable of comparing millions of characters per second. *Id.* at 107.

Appendix E

Figure 8.12 of Splash2 is reproduced below.

```

loop
  if (TAGin = SR) then
    if (SRCch =  $\emptyset$ ) then
      SRCch  $\leftarrow$  CHRin
      CHRout  $\leftarrow$   $\emptyset$ 
      DSTout  $\leftarrow$  PDSTin
    else
      CHRout  $\leftarrow$  CHRin
    endif
    PDSTout  $\leftarrow$  PDSTin
  else-if (TAGin = PR) then
    if (SRCch =  $\emptyset$ ) then
      DSTout  $\leftarrow$  PDSTin
    endif
    PDSTout  $\leftarrow$  DSTin
    CHRout  $\leftarrow$  CHRin
  else-if (TAGin = TG) then
    if (SRCch  $\neq$   $\emptyset$ ) and (CHRin  $\neq$   $\emptyset$ ) then
      DSTout  $\leftarrow$  min  $\begin{cases} \text{PDSTout} + \psi(\text{SRCch}, \text{CHRin}), \\ \text{DSTin} + \psi(\text{SRCch}, \emptyset), \\ \text{DSTout} + \psi(\emptyset, \text{CHRin}) \end{cases}$ 
    else-if (SRCch =  $\emptyset$ ) then
      DSTout  $\leftarrow$  DSTin
    endif
    PDSTout  $\leftarrow$  DSTin
    CHRout  $\leftarrow$  CHRin
  endif
  TAGout  $\leftarrow$  TAGin
endloop

```

FIGURE 8.12 Code executed by each PE in the unidirectional array

Figure 8.12 depicts “[t]he algorithm executed by each PE in the unidirectional array” that “compares one source sequence to a single target sequence.” *Id.* at 104–105. As shown above, SRCch and CHRin represent the character of the source sequence and target sequence being processed, respectively; DSTin and PDSTin represent the distance stream being input; and DSTout and PDSTout represent the computed distance stream being output. *Id.* Splash2 discloses that “[a] unidirectional array of length n can compare a source sequence of length at most n to a target sequence of length m in $O(n + m)$ steps.” *Id.* at 104.

Appendix E

b) Bidirectional Array Implementation

Figure 8.5 of Splash2 is reproduced below.




FIGURE 8.5 Data Flow through the Bidirectional Systolic Array. The source and target sequences are streamed through the array in opposite directions. A comparison is performed when a source character and a target character meet in a PE.

As shown in Figure 8.5, “[t]he source and target sequences enter the array on opposite ends and flow in opposite directions at the same speed.” Ex. 1007, 102. Splash2 discloses that

there is one distance stream associated with each character stream. At each step, the contents of the streams represent the characters to be compared and the distances along one of the antidiagonals of the distance matrix. At the end of the computation, the resulting edit distance is transported out of the array on the distance streams.

Id. (footnote omitted). In the bidirectional array implementation, each processing element “computes the distances along a particular diagonal of the distance

Appendix E

matrix.” *Id.* at 100. Splash2 discloses that the “DNA version of the bidirectional array,” for example, has 384 processing elements that would be capable of comparing millions of characters per second. *Id.* at 107.

Figure 8.7 of Splash2 is reproduced below.

```

loop
  if (SCin ≠ ∅) and (TCin ≠ ∅) then
    PEDist ← min {
      PEDist+ψ(SCin,TCin),
      TDin+ψ(SCin,∅),
      SDin+ψ(∅,TCin)
    }
    else-if (SCin ≠ ∅) then
      PEDist ← SDin
    else-if (TCin ≠ ∅) then
      PEDist ← TDin
    endif
    SCout ← SCin
    TCout ← TCin
    SDout ← PEDist
    TDout ← PEDist
  endif
endloop

```

FIGURE 8.7 Code Executed by Each PE in the Bidirectional Array

Figure 8.7 depicts the algorithm executed by each processing element in the bidirectional array. *Id.* at 101. As shown above, SCin and TCin represent the characters of the source sequence and target sequence being processed, respectively; SCout and TCout represent respective characters being output; SDin and TDin represent the source and target distance stream being input, respectively; and SDout and TDout represent the computed source and target distance stream being output, respectively. *Id.*

Splash2 discloses that “[c]omparing sequences of lengths m and n requires at least $2\max(m + 1, n + 1)$ processors. The number of steps required to compute the edit distance and to transport it out of the array is

Appendix E

proportional to the length of the array.” *Id.* at 103 (footnote omitted). Splash2 explains that the bidirectional array implementation had certain inefficiencies for database search operations, such as the fact that “[a]t each computational step, at most half of the PEs are active” and the source and target sequences were limited to half of the array’s length. *Id.* Those inefficiencies were remedied by the unidirectional array implementation according to Splash2. *Id.*

2. Claim 1

Petitioner argues that Splash2 is prior art under 35 U.S.C. §§ 102(a) and 102(b),¹⁸ and explains in detail how the reference discloses every limitation of claim 1, relying on the testimony of Dr. Stone as support. *See* Pet. 19–46;¹⁹ Ex. 1003 ¶¶ 96–157, 162–164, 174–179, 181–187, 191–204, 206–249. Petitioner argues that both the description

18. Petitioner provides evidence supporting its contention that Splash2, as well as RaPiD, Jeong, Chunky SLD, and Roccatano, are prior art printed publications under 35 U.S.C. §§ 102(a) and 102(b). *See* Pet. 19–20, 56–57, 70 (citing Ex. 1065 ¶¶ 8–12; Ex. 1066 ¶¶ 37–46, 75,77, 80–92; Ex. 1068 ¶¶ 4–6, 9–12); -1606 Pet. 55 (citing Ex. 1066 ¶¶ 47–55; Ex. 1068 ¶¶ 7–12); -1607 Pet. 60 (citing Ex. 1066 ¶¶ 65–73; Ex. 1067); Ex. 1007, Cover 6 (Library of Congress stamp for Splash2). Patent Owner does not assert otherwise in its Response, and we agree that the references are prior art for the reasons stated by Petitioner.

19. Petitioner asserts many of the same arguments, such as those regarding claim 1, in all three of its Petitions. *See* Pet. 19–46; -1606 Pet. 17–55; -1607 Pet. 19–59. Where applicable, we cite the papers in Case IPR2018-01605 for convenience.

Appendix E

of the unidirectional array and the description of the bidirectional array in Splash2 disclose certain limitations of claim 1. Pet. 19–46. We address both implementations below.

*a) Unidirectional Array Implementation**(1) Undisputed Limitations*

Petitioner argues that Splash2 discloses a “method for data processing,” as recited in claim 1, where “the characters of a source sequence of genetic information are compared against a target sequence of genetic information and an edit distance is calculated based on that comparison.” Pet. 28–30. Petitioner contends that Splash2 discloses a reconfigurable computing system comprising “at least one reconfigurable processor” (i.e., the array boards with FPGAs) with “a plurality of functional units” (i.e., the matrix of CLBs configured as processing elements), and “transforming” an “algorithm” (i.e., the edit distance algorithm disclosed in Splash2) into a “calculation” (i.e., calculation of the edit distance for two sequences). *Id.* at 30–32. According to Petitioner, the calculation is “data driven” because a character comparison is triggered by the arrival of data. *Id.* at 31–32. We agree, given our interpretation of the term “data driven” as “the scheduling of operations upon the availability of their operands.” *See supra* Section II.C.1.

Petitioner further contends that Splash2 discloses “forming” at least two functional units to perform the calculation (i.e., loading information into the FPGAs

Appendix E

to create the necessary processing elements) where “only functional units needed to solve the calculation are formed” because the FPGA components “clear all configuration memory before any configuration occurs” and the components instantiated for a particular calculation are tailored to that combination. Pet. 32–33. Petitioner argues that “each formed functional unit . . . interconnects with each other formed functional unit . . . based on reconfigurable routing resources within the at least one reconfigurable processor as established at formation,” as shown by the target sequence streaming through the processing elements in Splash2. *Id.* at 33–35. Finally, “lines of code” of the calculation are “formed as clusters of functional units” (i.e., the processing elements on each FPGA) where “in ‘the DNA version of the unidirectional array, each of the 16 array FPGAs (X1 to X16) holds 14 PEs’” according to Petitioner. *Id.* at 35–36 (quoting Ex. 1007, 107).

Patent Owner in its Response does not dispute that Splash2 discloses the above limitations of claim 1. Petitioner’s analysis for each of the limitations, supported by the testimony of Dr. Stone, which we credit, is persuasive. *See id.* at 28–36; Ex. 1003 ¶¶ 131–157, 162–164, 174–179, 181–187, 191–204, 206–212.

(2) *Disputed Limitation:*
“Computational Loops”

Claim 1 recites the following “computational loop” limitations:

Appendix E

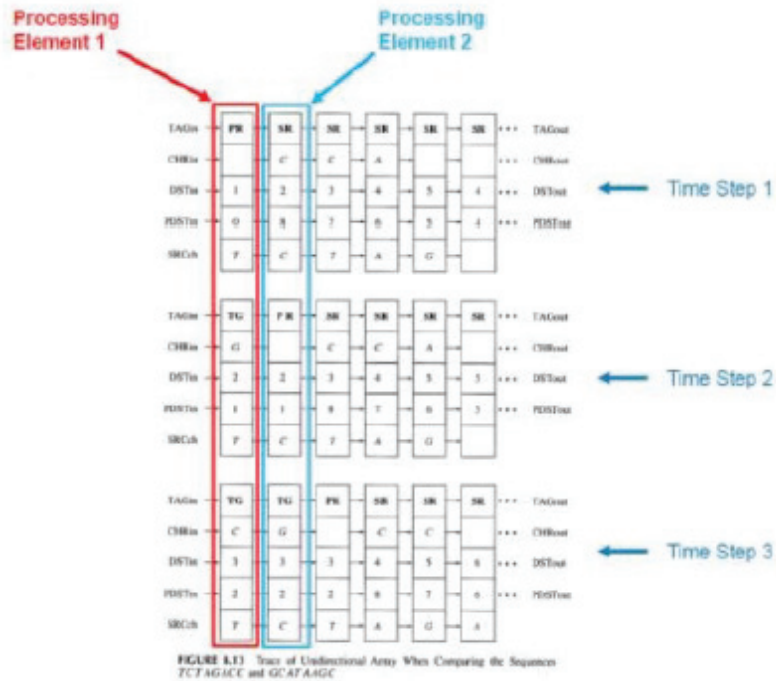
utilizing a first of said formed functional units to operate upon a subsequent data dimension of said calculation forming a first computational loop; and

substantially concurrently utilizing a second of said formed functional units to operate upon a previous data dimension of said calculation generating a second computational loop.

Petitioner argues that a person of ordinary skill in the art would have understood that the loop/endloop code, shown in Figure 8.12 above, “constitutes a time step of a ‘data dimension’ comprising multiple time steps because each such instance of code executes in a single time step, and the systolic computation comprises multiple time steps.” Pet. 36–37 (citing Ex. 1003 ¶¶ 215–216). According to Petitioner, a processing element executes the code “in a single time step,” taking in certain inputs from the preceding processing element (e.g., TAG_{in}, CHR_{in}, distance information DST_{in} and PDST_{in}) and passing along outputs (e.g., TAG_{out}, CHR_{out}, distance information DST_{out} and PDST_{out}) to the downstream neighboring processing element for “each character comparison” performed. *Id.* at 24–25, 27, 38–41, 45–46; *see* Ex. 1003 ¶¶ 113, 117, 145, 217–229.

Appendix E

Petitioner provides the following annotated version of Figure 8.13 of Splash2 (Pet. 38).



Annotated Figure 8.13 depicts the source sequence *TCTAGACC* pre-loaded in the processing elements (row SRCch), and the characters of the target sequence *GCATAAGC* streamed through the processing elements systolically (row CHRin). *Id.* at 37–40 (citing Ex. 1003 ¶¶ 217–229). For example, at Time Step 2, Processing Element 1 compares the first target character *G* to the first source character *T*. *Id.* at 38–39. At Time Step 3, (1) Processing Element 1 compares the second target character *C* to the first source character *T*, and (2) Processing Element 2, which has received the first target

Appendix E

character *G* streamed through the array, compares it to the second source character *C*. *Id.* at 39. Petitioner argues that, as shown in the example above, Processing Element 2 operates on a “subsequent data dimension” (i.e., comparing the second source character to the streamed target characters—the second row of the distance matrix) forming a “first computational loop” (i.e., “Processing Element 2 executing the loop instructions disclosed in Figure 8.12”); and Processing Element 1 operates on a “previous data dimension” (i.e., comparing the first source character to the streamed target characters—the first row of the distance matrix) forming a “second computational loop” (i.e., “Processing Element 1 executing the loop instructions disclosed in Figure 8.12”). *Id.* at 40 (emphases omitted). According to Petitioner, “[t]he same loop instructions are executed in each Processing Element, but they are still ‘first’ and ‘second’ computational loops because they are necessarily different instances of those loop instructions.” *Id.* Petitioner’s analysis for the “computational loop” limitations of claim 1, supported by the testimony of Dr. Stone, which we credit, is persuasive. *See* Ex. 1003 ¶¶ 213–229.

Patent Owner’s arguments are largely premised on its proposed interpretation of “computational loop” as requiring repeated execution “per datum,” which we do not adopt. *See* PO Resp. 76–82 (citing Ex. 2112 ¶¶ 127, 194–209; *supra* Section II.C.2. Patent Owner asserts that “[t]here is no disclosure [in Splash2] of looping or repeating of a computation multiple times *for each data* until a condition is met or a number of repetitions has been satisfied,” as allegedly represented in Figure 4B of

Appendix E

the '800 patent. PO Resp. 76–77, 79–80 (emphasis added). We disagree that the claim requires such a process, and instead interpret “computational loop” to mean “a set of computations that is executed repeatedly, either a fixed number of times or until some condition is true or false.” *See supra* Section II.C.2.

Patent Owner further contends that the code shown in Figure 8.12 of Splash2 does not illustrate a “computational loop.” PO Resp. 76–82. Patent Owner first points to the if/else-if conditional statements in Figure 8.12, which select an execution path based on whether a condition is true or false. *Id.* at 77–78. Petitioner in its analysis, however, relies on the overall loop/endloop aspect of the code, not the if/else-if statements within that code. Pet. 25–27, 36–40. Petitioner first introduces the bidirectional array implementation and argues that Splash2 “discloses in Figure 8.7 the code executed by each processing element, . . . which *includes a loop for computing the edit distance* (i.e., the minimum cost of transforming one sequence to the other) between characters,” reproducing the code of Figure 8.7 (which includes the notations “loop” and “endloop”), then does the same for the unidirectional array code in Figure 8.12 (which also includes “loop” and “endloop”). *Id.* at 24–25, 27 (emphasis added).

What we must determine is how a person of ordinary skill in the art would have understood the code disclosed in Splash2, in particular “loop” and “endloop.” The parties and their respective declarants have different views. *See id.* at 24–27, 36–40; PO Resp. 78–83; Reply 34–36. Patent Owner relies on Dr. Homayoun, who testifies that

Appendix E

“[t]o be defined as a loop, a loop exit condition must be specified. This is generally specified either in the first line of the code (begin loop and then condition) or the last line of [the] loop (loop end and then condition).” Ex. 2112 ¶ 200. According to Dr. Homayoun, the code in Figure 8.12 lacks an “exit condition” at the beginning or end, so if the code was a loop as Petitioner contends, it “would run forever.” *Id.* ¶¶ 201–202. “This means nothing would be calculated by the code because the first piece of data would enter each PE and then run forever. Nothing would be passed to the next PE.” *Id.* ¶ 203. Therefore, “[t]he only reasonable interpretation of this pseudocode is to assume that something replaces the ‘loop-endloop’ syntax so that the pseudocode executes once then passes the data to the next PE.” *Id.* ¶ 204. In that case, the processing element executing the code “would never evaluate the same data more than once” and would not perform a “computational loop” under Patent Owner and Dr. Homayoun’s claim interpretation. *Id.* ¶¶ 205–208.

Dr. Stone testifies that the loop/endloop code in Figure 8.12 represents “a loop for computing the edit distance (i.e., the minimum cost of transforming one sequence to the other) between characters.” Ex. 1003 ¶¶ 113, 117, 144–147. At a particular time step, the processing element “execut[es] the loop instructions disclosed in Figure 8.12” to compare the target character that was streamed in to the source character loaded in that processing element. *Id.* ¶¶ 117, 217–227. The processing element then repeatedly does the same for each input character of the target sequence (i.e., “multiple comparisons”). *Id.* Responding to Dr. Homayoun’s testimony regarding the loop/endloop

Appendix E

code,²⁰ Dr. Stone points out that the repeated comparisons result in the code being executed a fixed number of times because the target and source sequences are of limited length (m and n , respectively) and each processing element executes the loop instructions until it completes computing distances along the diagonal matrix. Ex. 1076 ¶¶ 12–13.

Based on our review of the full trial record, we find that Dr. Stone’s reading, explained in detail in his declarations, is most consistent with the disclosure of Splash2 and that a person of ordinary skill in the art would have read the reference in the manner he describes. Figure 8.12 is labeled as a “loop,” indicating that the code within the “loop” and “endloop” designations is executed repeatedly by the processing element. Splash2 expressly describes the reason *why* it is executed repeatedly: the target sequence is “streamed through the array” so that a single incoming target character can be compared to the single source character loaded in that processing element. Ex. 1007, 103. The target sequence and source sequence both contain a large set of characters of “length m ” and “length n ,” respectively. *Id.* at 104. The target sequence can be any “length,” whereas the “maximum length” of the source sequence is limited by the “length” of the array. *Id.* Naturally, when the end of the m -length target sequence

20. Again, we do not agree with Patent Owner that Petitioner’s Reply and Dr. Stone’s reply testimony are improper, as they respond directly to Patent Owner’s arguments in the Response and Dr. Hodayoun’s testimony regarding how a person of ordinary skill in the art would have understood the loop/endloop code in Splash2. *See* Sur-Reply 8–9; Reply 15–18, 34–36; Ex. 1076 ¶¶ 9–19; 37 C.F.R. § 42.23(b).

Appendix E

is reached, no more characters can be streamed in and the loop would complete for that processing element (i.e., after a fixed number of times executing the code). *See* Ex. 1076 ¶ 13. Thus, a person of ordinary skill in the art would have understood the code in Splash2 to loop and have an exit condition—it executes repeatedly (i.e., once for each target character streamed through the array) a fixed number of times (i.e., until the last target character in the target sequence is streamed through and compared).

Indeed, Splash2 expressly discloses limits on the number of times a processing element executes the code. In the unidirectional array implementation, “each PE computes the distances in one row of the distance matrix,” and “[a]t each time step, the PEs compute the distances along a single antidiagonal in the distance matrix,” indicating that processing would complete when there are no more distances to compute. *See* Ex. 1007, 104; *see also id.* at 100 (disclosing that each processing element “computes the distances along a particular diagonal of the distance matrix” in the bidirectional array implementation), 102 (“At the end of the computation, the resulting edit distance is transported out of the array on the distance streams.”). Dr. Stone’s understanding of the loop/endloop code is consistent with the text of Splash2 and the “loop” and “endloop” designations shown in the code itself. *See* Ex. 1076 ¶¶ 9–13.

Importantly, we also note that Splash2’s description of the unidirectional array implementation would disclose the “computational loop” limitations even under Patent Owner’s proposed interpretation. Because the loop/

Appendix E

endloop code is executed repeatedly for the same source character (e.g., source character *T* for Processing Element 1 and source character *C* for Processing Element 2 in annotated Figure 8.13 above), comparing it to each target character as it is streamed in, computations are executed repeatedly “per datum” (i.e., for the same source character “datum”) a fixed number of times. *See id.* ¶¶ 9–11, 14.

Patent Owner’s next argument is that Splash2 does not disclose “utilizing” the two “formed functional units” on the reconfigurable processor to operate on data dimensions forming “computational loops,” as recited in claim 1, because “even if Splash2 could be read to disclose two computational loops, it does not disclose forming those loops in the FPGAs of Splash2.” PO Resp. 83–85; Sur-Reply 17–19. According to Patent Owner, the external Sun workstation, not the FPGAs, would “handle any looping.” PO Resp. 83 (emphasis omitted). As support, Patent Owner relies on testimony from its declarants (Ex. 2112 ¶ 209; Ex. 2166 ¶¶ 42–43) and one paragraph from the thesis of Richard Peyton Halverson, Jr., Ph.D., “The Functional Memory Approach to the Design of Custom Computing Machines,” Ph.D. diss., University of Hawaii, 1994 (Ex. 2169, “the Halverson thesis”). PO Resp. 83–84.

We disagree. Splash2 expressly discloses that “[b]oth the bidirectional and unidirectional systolic arrays have been implemented *on the Splash 2 programmable logic array*, with versions for DNA and protein sequences.” Ex. 1007, 104 (emphasis added). “In the DNA version of the unidirectional array, each of the 16 array FPGAs (X1 to X16) holds 14 PEs.” *Id.* at 107. Splash2’s discussion of

Appendix E

the two implementations for the edit distance algorithm repeatedly refers to the functionality of an individual processing element (PE). *See, e.g., id.* at 100 (“There are several ways to map the edit distance computation onto a linear systolic array. We describe two such mappings.”), 104 (“At each time step, the PEs compute the distances along a single antidiagonal in the distance matrix The algorithm executed by each PE in the unidirectional array is listed in Figure 8.12.”). Thus, the processing elements that perform the looping are instantiated within the FPGAs in Splash2. *See* Ex. 1003 ¶¶ 208, 217–229; Ex. 1076 ¶¶ 18–19.

The Halverson thesis does not support a different reading. It discloses:

Splash 2 contains one or more boards each with an array of 16 well connected XILINX 4010 chips [Gokhale and Minnich, 1993]. The architecture does an excellent job supporting pipelined and SIMD processor configurations. Splash 2, for example, can be programmed in dbC, which is a superset of C used on other SIMD computers. The dbC preprocessor produces C that runs on the Sun and VHDL which define SIMD processors with an instruction set tailored to the application, one or more of which fit into each XILINX chip. When the actual program executes, looping is still handled in the Sun, which transmits SIMD instructions to the Splash 2 board(s).

Appendix E

Ex. 2169, 37–38. This isolated reference is to the Splash 2 system and how it “can be programmed in [Data-parallel Bit-serial C (dbC)].” *See id.* It does not mention or relate to the particular systolic array architectures and edit distance calculations described in Chapter 8 of Splash2. Further, whereas the paragraph above references the Splash 2 system being programmed in dbC, Splash2 describes programming the edit distance calculations in VHDL. *See id.* at 37–38, 208; Ex. 1007, 70, 106; Ex. 1074, 94–95, 97 (“Gokhale and Minnich” paper referenced in the Halverson thesis describing dbC and VHDL); Ex. 1003 ¶ 150; Ex. 1076 ¶¶ 15–17.

Dr. Stone conducted a detailed analysis of Splash2, citing specific figures and discussion in the reference, and explained why a person of ordinary skill in the art would have understood Splash2 to disclose utilizing two functional units to operate on different data dimensions of a calculation forming two “computational loops,” as recited in claim 1. After reviewing the full trial record, we credit that testimony. For the foregoing reasons, we find persuasive the arguments and evidence presented by Petitioner and Dr. Stone and, therefore, find that Splash2, in its description of the unidirectional array implementation, discloses the “computational loop” limitations of claim 1.

(3) *Disputed Limitation: “Seamlessly”*

Claim 1 recites that “said implementation of said calculation enables said first computational loop and said second computational loop execute concurrently and pass

Appendix E

computed data seamlessly between said computational loops” (the “seamlessly” limitation).

Petitioner contends that the computational loops in Splash2 “execute concurrently” and “pass computed data seamlessly between” the computational loops because each processing element “execute[s] an instance of” the loop instructions, shown in Figure 8.12 above, and “each Processing Element directly communicates computed distance information to a neighbor.” Pet. 45–46 (citing Ex. 1003 ¶¶ 244–249). Dr. Stone explains that “[t]he output of a functional unit *passes directly* to the input of the next functional unit with *no intervening gaps or additional interfaces.*” Ex. 1003 ¶ 246 (emphases added). Each functional unit takes in “DSTin from the previous functional unit in the series” and outputs “DSTout for the next functional unit in the series.” *Id.* ¶ 246 (citing Ex. 1007, 104, Fig. 8.10). Dr. Stone provides the following annotated version of a portion of Figure 8.13 of Splash2 (*id.* ¶ 247).



Appendix E

The annotated portion above of a trace of the unidirectional array depicts “the seamless communication of the DST data in the systolic calculation as a sequence of red arrows.” *Id.* Dr. Stone explains that

[t]he red arrows indicate . . . that the computed output DSTout from a functional unit is *directly connected* to the next functional unit input DSTin. Similarly, all output data from one functional unit are seamlessly connected to the corresponding inputs of the functional unit in the sequence. There are *no intermediate interfaces between modules* to translate the output of one module into a form where it can be used as input data to the next module. Such interfaces, if they were to exist, would be “seams” in the communication links between adjacent functional units.

Id. (emphases added). As explained above, we interpret “pass computed data seamlessly between said computational loops” to mean “communicate computed data directly between functional units that are calculating computational loops.” *See supra* Section II.C.3. Petitioner’s analysis for the “seamlessly” limitation of claim 1, consistent with that interpretation and supported by the testimony of Dr. Stone, which we credit, is persuasive. *See* Ex. 1003 ¶¶ 242–247, 249.

Patent Owner makes four arguments in its Response. First, Patent Owner contends that in passing computed data between the computational loops operating on

Appendix E

data dimensions of the calculation, claim 1 “specifically require[s] a sequence of computations *per datum*,” but the loop/endloop code in Splash2 “merely selects an execution path for the processor to execute once per datum.” PO Resp. 91–92. We disagree for the reasons stated above. *See supra* Sections II.C.2, II.E.2.a.2.

Second, Patent Owner argues that “the FPGAs [in Splash2] must communicate with the Sun workstation (which is handling any looping) through the Sbus” and the “boundary between the FPGAs and the workstation . . . clearly constitutes a ‘seam,’” citing the same declarant testimony and Halverson thesis discussed above in connection with the “computational loop” limitations. PO Resp. 91–92, 95. Again, we disagree that the Sun workstation handles the looping discussed in Chapter 8 for the edit distance calculations. *See supra* Section II.E.2.a.2.

Third, Patent Owner asserts that “Splash2 is, at best, ambiguous on whether memory is used to store the results from each processing element after each time step to preserve it for output and later use.” PO Resp. 92–95 (citing Ex. 2112 ¶¶ 210–219). Patent Owner contends that “storage is likely necessary to preserve the values calculated at each timestep,” where the computed data is overwritten at each time step and the overall edit distance is based on the repeated calculations and overwriting (rather than just the final time step). *Id.* at 93; *see* Ex. 2112 ¶ 210 (identifying, for example, SDout, TDout, and PEDist from the bidirectional array implementation). Patent Owner points out that Splash2 has “local memory

Appendix E

at each FPGA for storage purposes,” which “can be used for storage of results.” PO Resp. 93–94 (citing Ex. 1007, 88, 95, 102 n.3; Ex. 1035, 5; Ex. 2157, 205–206). According to Patent Owner, “it is equally (if not more) plausible for a [person of ordinary skill in the art] to interpret Splash2 to use the local memory due to the known timing problems in systolic systems.” *Id.* at 94–95.

We are not persuaded that Splash2 is ambiguous in the manner Patent Owner contends. The disclosure and figures discussed above indicate that computed distance information is communicated directly between processing elements. In the context of the unidirectional array implementation, for example, a processing element takes in DST_{in}, executes the loop/endloop code to compute DST_{out}, and outputs DST_{out} to the next processing element in the series. *See* Ex. 1007, 104 (Fig. 8.10 depicting the inputs and outputs for the unidirectional array PE), 105 (Fig. 8.12 showing how DST_{out} is calculated), 106 (Fig. 8.13 depicting direct connections between processing elements); Ex. 1003 ¶¶ 245–249; *see also* Ex. 1007, 101–102 (Figs. 8.6–8.8 showing inputs of SD_{in} and TD_{in} and output of the computed PED_{ist} as SD_{out} and TD_{out} in either direction for the bidirectional array implementation). Claim 1 recites that the implementation enables the computational loops to “pass computed data seamlessly between said computational loops.” There is no indication in Splash2 itself that there is any memory or other intervening structure *between* the disclosed processing elements. That data may be stored temporarily *within* a processing element (e.g., in a storage register) does not automatically mean that memory between the processing

Appendix E

element and another processing element is used for communicating that data. *See* Ex. 1007, 104 (stating that the unidirectional array PE “stores two distances, DST and PDST,” but never stating that memory between the PE and another PE is used to do so).

Fourth, Patent Owner argues that Splash2 does not disclose the “seamlessly” limitation even under Petitioner’s proposed interpretation because “the Xilinx FPGAs . . . clearly contain structure (such as the buffered switch matrix) within the internal routing resources to connect processing elements.” PO Resp. 95–100 (citing Ex. 1035, 28–31; Ex. 2077, 19–29, 32–34, 37–41, 46–51, 59–65, 70). Again, just because an FPGA has memory does not mean that when functional units are instantiated within the FPGA, memory is necessarily placed *between* functional units or used to transfer data from one functional unit to another. *See* Reply 25, 47; *supra* Section II.C.3.

For the foregoing reasons, we find persuasive the arguments and evidence presented by Petitioner and Dr. Stone and, therefore, find that Splash2, in its description of the unidirectional array implementation, discloses the “seamlessly” limitation of claim 1.

b) Bidirectional Array Implementation

Petitioner’s explanation for why a person of ordinary skill in the art would have understood Splash2’s description of the bidirectional array implementation to disclose the limitations of claim 1 is similar to its explanation for the unidirectional array implementation. With respect to the

Appendix E

“computational loop” limitations, Petitioner argues that a person of ordinary skill in the art would have understood that the loop/endloop code shown in Figure 8.7 above “constitutes a time step of a ‘data dimension’ comprising multiple time steps because each such instance of code executes in a single time step, and the systolic computation comprises multiple time steps.” Pet. 36–37 (citing Ex. 1003 ¶¶ 215–216). Petitioner provides the following annotated version of Figure 8.8 of Splash2 (*id.* at 42).




FIGURE 8.8 Trace of Bifunctional Array When Comparing the Sequences TCTAGACC and GCATAAGC

Annotated Figure 8.8 depicts the characters of the source sequence TCTAGACC streamed left to right through the processing elements (row SCin), and the characters of

Appendix E

the target sequence *GCATAAGC* streamed right to left through the processing elements (row TCin). *Id.* at 42–45 (citing Ex. 1003 ¶¶ 230–241). For example, at Time Step 2, Processing Element 2 compares the first target character *G* to the first source character *T* (shown in yellow). *Id.* at 42. At Time Step 4, (1) Processing Element 2 compares the second target character *C* to the second source character *C* (shown in orange), and (2) Processing Element 1, which has received the first target character *G* streamed through the array, compares it to the third source character *T* (shown in green). *Id.* at 42–43.

Petitioner argues that, as shown in the example above, Processing Element 2 operates on a “subsequent data dimension” (i.e., comparing streamed target and source characters on a particular diagonal of the distance matrix) forming a “first computational loop” (i.e., “Processing Element 2 executing the loop instructions disclosed in Figure 8.7”); and Processing Element 1 operates on a “previous data dimension” (i.e., comparing streamed target and source characters on a particular diagonal of the distance matrix) forming a “second computational loop” (i.e., “Processing Element 1 executing the loop instructions disclosed in Figure 8.7”). *Id.* at 43–45 (emphases omitted). As with the unidirectional array implementation, “the loop instructions executed in each Processing Element are first and second computational loops because they are necessarily different instances of those loop instructions.” *Id.* at 45.

With respect to the “seamlessly” limitation, Petitioner argues that each processing element “execute[s] an

Appendix E

instance of” the loop instructions shown in Figure 8.7 above and “each Processing Element directly communicates computed distance information to a neighbor.” *Id.* at 45–46 (citing Ex. 1003 ¶¶ 244–249). Dr. Stone explains that the same passing of data occurs in the unidirectional and bidirectional array implementations, where “[t]he output of a functional unit *passes directly* to the input of the next functional unit with *no intervening gaps or additional interfaces.*” Ex. 1003 ¶¶ 246, 248 (emphases added). Dr. Stone provides the following annotated version of a portion of Figure 8.8 of Splash2 (*id.* ¶ 248).




FIGURE 8.8 Trace of Bidirectional Array When Comparing the Sequences TCTAGACC and GCATAAGC

The annotated portion above of a trace of the bidirectional array depicts the seamless communication of distance information. *Id.* Dr. Stone explains that

[t]he blue arrows indicate that the computed output PEDist from a functional unit is directly connected to the next functional unit input PEDist in both directions. The code for the Bidirectional implementation discloses that the computed PEDist is conveyed to the left on output TDout, and to the right on output SDout.

Id. (citing Ex. 1007, 101, Fig. 8.6).

Appendix E

Patent Owner argues in its Response that Splash2 does not disclose the “computational loop” and “seamlessly” limitations of claim 1. PO Resp. 76–85, 91–100. Patent Owner’s arguments apply to both the unidirectional and bidirectional array implementations described in Splash2, and we disagree with those arguments for the reasons explained above.²¹ *See id.*; *supra* Section II.E.2.a. Petitioner’s analysis of how Splash2’s description of the bidirectional array implementation discloses the limitations of claim 1, supported by the testimony of Dr. Stone, which we credit, is persuasive. *See* Ex. 1003 ¶¶ 131–157, 162–164, 174–179, 181–187, 191–204, 206–216, 230–246, 248–249. For the reasons stated above, we find persuasive the arguments and evidence presented by Petitioner and Dr. Stone and, therefore, find that Splash2, in its description of the bidirectional array implementation, discloses the limitations of claim 1.

c) Conclusion

We are persuaded by Petitioner’s explanations and supporting evidence regarding both the unidirectional and bidirectional implementations of Splash2, and find that Splash2 discloses every limitation of claim 1 under

21. The loop/endloop code executed by a processing element operates similarly in both implementations. We note that although the unidirectional array implementation would meet the “computational loop” limitation even under Patent Owner’s proposed interpretation (because the code is executed repeatedly for the same source character), the bidirectional array implementation would not (because the target and source sequences are both streamed through).

Appendix E

both of Petitioner's theories. Petitioner has proven, by a preponderance of the evidence, that claim 1 is anticipated by Splash2 under 35 U.S.C. §§ 102(a) and 102(b).

3. *Claims 18, 21, and 22*

Dependent claims 18, 21, and 22 limit the "calculation" of parent claim 1 to a particular type. Claim 18 recites that the calculation comprises a "search algorithm for data mining," claim 21 recites that it comprises a "genetic pattern matching function," and claim 22 recites that it comprises a "protein folding function." Petitioner argues that a person of ordinary skill in the art would have understood Splash2's edit distance calculation to be each type of calculation. -1607 Pet. 46–50. For example, Petitioner argues that the edit distance calculation is a "genetic pattern matching function" because it "compares characters of two genetic sequences." *Id.* at 48; *see* Ex. 1007, 100–104; Ex. 1003 ¶¶ 84–87, 430–433. Patent Owner does not argue separately dependent claims 18, 21, and 22 in its Response, only disputing Petitioner's contentions with respect to parent claim 1. PO Resp. 73–85, 91–100. We have reviewed Petitioner's contentions regarding claims 18, 21, and 22, which are consistent with the disclosure of Splash2 and supported by the testimony of Dr. Stone, and are persuaded that Petitioner has proven, by a preponderance of the evidence, that dependent claims 18, 21, and 22 are anticipated by Splash2 under 35 U.S.C. §§ 102(a) and 102(b). *See* -1607 Pet. 46–50.

*Appendix E**4. Claim 15*

Claim 15 recites that the “instantiating” step of parent claim 1 (which we read to refer to the “forming” step) includes “establishing a stream communication connection between functional units.” Petitioner argues that Splash2 discloses establishing such a connection “between adjacent Processing Elements in both the unidirectional and bidirectional edit distance calculation because it discloses streaming target and source characters between processing elements both unidirectionally and bidirectionally.” -1606 Pet. 44–45. Patent Owner responds that Splash2 does not disclose the limitation because it does not disclose a “queue” between processing elements or “signaling” for the processing elements to interact with such a queue. PO Resp. 73–76. Patent Owner’s arguments are premised on its proposed interpretation of “stream communication” as “a data path that acts like a queue connecting via the reconfigurable routing resources a producer and a consumer of data that operate concurrently.” *Id.* at 45, 73–76. As explained above, we disagree and instead interpret “establishing a stream communication connection between functional units” to mean “establishing a connection over which data is streamed between functional units.” *See supra* Section II.C.4.

Given our interpretation, Petitioner has made a sufficient showing that Splash2 discloses the limitation. In the unidirectional array implementation, “[t]he source sequence is loaded once” and “target sequences are streamed through the array one at a time, separated

Appendix E

by control characters.” Ex. 1007, 103. “[D]ata flows through the unidirectional array in one direction.” *Id.* In the bidirectional array implementation, “[t]he source and target sequences are streamed through the array in opposite directions.” *Id.* at 101–102. “[T]here is one distance stream associated with each character stream. At each step, the contents of the streams represent the characters to be compared and the distances along one of the antidiagonals of the distance matrix.” *Id.* at 102. Thus, Splash2 discloses in the various implementations establishing a connection over which data (e.g., characters of the source and target sequences, computed distance information) is streamed between functional units. Petitioner has proven, by a preponderance of the evidence, that dependent claim 15 is anticipated by Splash2 under 35 U.S.C. §§ 102(a) and 102(b).

F. Obviousness Ground Based on Splash2 (Claims 1, 15, 18, 21, and 22)

1. Obviousness Contentions

Petitioner argues, in the alternative, that claims 1, 15, 18, 21, and 22 would have been obvious over Splash2. Pet. 46–53; -1606 Pet. 45–51; -1607 Pet. 50–56. Petitioner argues that “[t]o the extent one might argue” that the chapters of Splash2 cannot be considered together for purposes of anticipation, a person of ordinary skill in the art would have found claim 1 to be obvious based on the full disclosure of Splash2. Pet. 47–48. Petitioner further contends that “[t]o the extent one might argue” that Splash2 does not disclose four limitations of claim 1 (i.e.,

Appendix E

the “transforming” step, portions of the “forming” step, the limitation that “only functional units needed to solve the calculation are formed,” and the limitation that “each formed functional unit at the at least one reconfigurable processor interconnects with each other formed functional unit”), a person of ordinary skill in the art would have been motivated to modify Splash2 to include them. *Id.* at 48–53. For example, Petitioner argues that if Splash2 does not disclose the limitation that “only functional units needed to solve the calculation are formed,” doing so would have been obvious because the FPGAs in Splash2 are “cleared of any configuration data before a new configuration can be entered,” the arrays are customized for a particular application, and a person of ordinary skill in the art would have been motivated to form only the processing elements necessary to solve a calculation to “avoid the additional work” of writing code for processing elements that “performed no function [and] served no purpose.” *Id.* at 51. Splash2 is a single reference, and Patent Owner does not dispute that the chapters of Splash2 can be considered together. *See* Ex. 1007, 104 (Chapter 8 indicating that “the bidirectional and unidirectional systolic arrays have been implemented on the Splash 2 programmable logic array,” which was described in earlier chapters). Nor does Patent Owner dispute that Splash2 discloses the four identified limitations. We find that the four limitations are disclosed by Splash2. *See supra* Section II.E.2.a.1. Thus, we need not address the alternative arguments made by Petitioner.

Patent Owner argues that Splash2 fails to disclose certain other limitations of claims 1 and 15. *See supra* Sections II.E.2, II.E.4. We disagree, find that Splash2

Appendix E

expressly teaches all limitations of claims 1, 15, 18, 21, and 22, and conclude that Petitioner has proven anticipation by a preponderance of the evidence. *See id.* Accordingly, absent a persuasive showing of secondary considerations, which we discuss below, the claims also would have been obvious based on Splash2. *See Realtime Data, LLC v. Iancu*, 912 F.3d 1368, 1373 (Fed. Cir. 2019) (“[I]t is well settled that a disclosure that anticipates under § 102 also renders the claim invalid under § 103, for anticipation is the epitome of obviousness.” (citations and internal quotation marks omitted)).

2. *Secondary Considerations of Nonobviousness*

Patent Owner also argues that secondary considerations of nonobviousness demonstrate that claims 1–5, 7–9, 15, 17, 18, and 20–24 would not have been obvious to a person of ordinary skill in the art. PO Resp. 116–120. Although some of the challenged claims are anticipated by—and, therefore, would have been obvious based on—Splash2, because the parties refer to the challenged claims collectively in their arguments regarding secondary considerations, we do so as well and now address those arguments as applied to all of the claims challenged as obvious. *See id.*; Reply 48–51.

“In order to accord substantial weight to secondary considerations in an obviousness analysis, the evidence of secondary considerations must have a nexus to the claims, *i.e.*, there must be a legally and factually sufficient connection between the evidence and the patented invention.” *Fox Factory, Inc. v. SRAM, LLC*, 944 F.3d

Appendix E

1366, 1373 (Fed. Cir. 2019) (citations and internal quotation marks omitted). “The patentee bears the burden of showing that a nexus exists.” *Id.* “To determine whether the patentee has met that burden, we consider the correspondence between the objective evidence and the claim scope.” *Id.* A patentee is entitled to a rebuttable presumption of nexus “when the patentee shows that the asserted objective evidence is tied to a specific product and that product ‘embodies the claimed features, and is coextensive with them.’” *Id.* (citation omitted). However, “[a] finding that a presumption of nexus is inappropriate does not end the inquiry into secondary considerations.” *Id.* “To the contrary, the patent owner is still afforded an opportunity to prove nexus by showing that the evidence of secondary considerations is the ‘direct result of the unique characteristics of the claimed invention.’” *Id.* at 1373–1374 (citation omitted).

“Where the offered secondary consideration actually results from something other than what is both claimed and *novel* in the claim, there is no nexus to the merits of the claimed invention,” meaning that “there must be a nexus to some aspect of the claim not already in the prior art.” *In re Kao*, 639 F.3d 1057, 1068–69 (Fed. Cir. 2011). On the other hand, there is no requirement that “objective evidence must be tied exclusively to claim elements that are not disclosed in a particular prior art reference in order for that evidence to carry substantial weight.” *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1331 (Fed. Cir. 2016). A patent owner may show, for example, “that it is the claimed combination as a whole that serves as a nexus for the objective evidence; proof of nexus is not limited to only

Appendix E

when objective evidence is tied to the supposedly ‘new’ feature(s).” *Id.* at 1330. Ultimately, the fact finder must weigh the secondary considerations evidence presented in the context of whether the claimed invention as a whole would have been obvious to an ordinarily skilled artisan. *Id.* at 1331–32.

As objective evidence of nonobviousness, Patent Owner cites the declarations of Dr. Homayoun (Ex. 2112), Mr. Huppenthal (Ex. 2101), and Dr. El-Ghazawi (Ex. 2166); the deposition transcript of Stephen M. Trimberger, Ph.D. in Cases IPR2018-01599 and IPR2018-01600 (Ex. 2075); the deposition transcript of Dr. Stone in Case IPR2018-01594 (Ex. 2065); and two publications authored by, among others, Dr. El-Ghazawi and one of the authors of Splash2, Duncan Buell, Ph.D.: Duncan Buell *et al.*, “High-Performance Reconfigurable Computing,” *IEEE Computer Society*, pp. 23–27 (Mar. 2007) (Ex. 2168), and Tarek El-Ghazawi *et al.*, “The Promise of High-Performance Reconfigurable Computing,” *IEEE Computer Society*, pp. 69–76 (Feb. 2008) (Ex. 2167). PO Resp. 116–120.

Patent Owner asserts that there was commercial success and praise by others for “DirectStream’s patented products (SRC-6, SRC-6e),” which involved “instantiating the reconfigurable resources necessary for an algorithm so as to maximize the speed data can be passed between different looping portions of the program” running concurrently “to achieve a particular systolic system that processes data without the associated delays of communications protocols or the Von Neumann bottleneck.” *Id.* at 14–16, 120. Patent Owner, however, does

Appendix E

not provide any explanation or analysis demonstrating that its “SRC-6” or “SRC-6e” products were used to perform the methods recited in any of the challenged claims, or are coextensive with the claimed methods. Mr. Huppenthal discusses “SRC-6 products and the SRC-7 . . . under development” and testifies that “[p]roduction systems would incorporate . . . AI and Machine Learning as described in [the ’800 patent]” and “Army and Air Force applications would make extensive use of program loops and streams as defined in [the ’800 patent].” Ex. 2101 ¶¶ 80, 83. Importantly, though, Mr. Huppenthal never refers to the claims in his testimony and acknowledged that he did not perform “any kind of comparison of the claims [of the ’800 patent] to any particular system.” *See id.* ¶¶ 80–87; Ex. 1073, 106:14–107:14, 108:17–109:5; PO Opp. 2 (acknowledging that Mr. Huppenthal only provides “high-level, general testimony regarding . . . specific hardware features of interest” to “customers who purchased SRC-6 computers” and “does not specifically discuss any applications any customers ran on the SRC-6”). Therefore, we find that a presumption of nexus is inappropriate.

Nor does the other evidence cited by Patent Owner establish a nexus between DirectStream’s products and any of the challenged claims. Patent Owner cites one of the publications and Dr. Homayoun’s declaration mentioning the “SRC-6,” “SRC-6E,” and “SRC-7” products, but does not point to anything in either document explaining how the products embody the claims. *See* PO Resp. 120 (citing Ex. 2167, 2–7, Fig. 4; Ex. 2112 ¶¶ 108–109, 118–119). Absent some explanation of how the cited products worked, there is no basis to say that they were used to practice the

Appendix E

methods recited in the challenged claims. Accordingly, we find that Patent Owner has not established a sufficient nexus between the claims and the alleged commercial success and industry praise.

Patent Owner further asserts that there was a long-felt but unmet need for “very fast processing of large volumes of data, and improvements in conventional systems focused on reducing the delays associated with chip-to-chip communications protocols and also the Von Neumann processing bottleneck.” *Id.* at 118 (citing Ex. 2166 ¶¶ 17–25; Ex. 2168, 3–5; Ex. 2101 ¶¶ 36–41; Ex. 2112 ¶¶ 40–106). According to Patent Owner, there was skepticism in the industry that “reconfigurable processors could adequately satisfy these needs compared to other well-established solutions in the same technology space.” *Id.* (citing Ex. 2075, 129:24–130:20; Ex. 2065, 168:9–169:4, 179:6–13, 197:8–11).

To support a conclusion of nonobviousness, an alleged long-felt need must have been a persistent one that was recognized by those of ordinary skill in the art, must not have been satisfied by another before the challenged patent, and must have been satisfied by the claimed invention. *Perfect Web Techs., Inc. v. InfoUSA, Inc.*, 587 F.3d 1324, 1332–33 (Fed. Cir. 2009). Similarly, evidence showing that “skilled artisans were initially skeptical about the [claimed] invention” is relevant to nonobviousness. *Metabolite Labs., Inc. v. Laboratory Corp. of Am. Holdings*, 370 F.3d 1354, 1368 (Fed. Cir. 2004). We have reviewed all of the cited evidence and do not find it persuasive, as Patent Owner does not provide

Appendix E

any explanation establishing a nexus to the challenged claims. Patent Owner does not explain in its Response, for example, how exactly the claimed methods solved the identified problems of speed and reducing communication delays. Nor does Patent Owner explain how the cited testimony from Dr. Trimberger and Dr. Stone allegedly shows skepticism in the industry that the reconfigurable processor approach *of the challenged claims* would reduce such issues, as opposed to mere views on reconfigurable processors in general. *See* PO Resp. 118.

Patent Owner also contends that others in the industry failed to solve the problems allegedly solved by the '800 patent. PO Resp. 119–120 (citing Ex. 2169, 12–20; Ex. 2176, 4; Ex. 2166 ¶¶ 32–43; Ex. 2101 ¶¶ 30–31, 36–41; Ex. 2112 ¶¶ 179–184; Ex. 2065, 168:9–169:4, 179:6–13, 197:8–11; Ex. 2075, 129:24–130:20). We are not persuaded for similar reasons. Patent Owner does not explain in sufficient detail how the cited evidence demonstrates a “failure of others to do *that which the patent claims*.” *See id.* at 119 (quoting *Para-Ordnance Mfg., Inc. v. SGS Importers Int'l, Inc.*, 73 F.3d 1085, 1094 (Fed. Cir. 1995)) (emphasis added). Without further explanation, we are not persuaded that Patent Owner’s evidence of the failure of others supports the nonobviousness of the challenged claims.

Finally, Patent Owner points to the publications authored by Dr. Buell and Dr. El-Ghazawi in 2007 and 2008 (well after the effective filing date of the '800 patent, October 31, 2002) that “survey[ed] . . . the state of the art” and discussed “DirectStream’s SRC-6 and SRC-6e systems” but “did not include [the Splash 2

Appendix E

system] as an example of a viable [High-Performance Computing Technology (HPRC)] system.” *Id.* at 119–120 (citing Exs. 2167, 2168). Patent Owner does not point to any authority—and we are not aware of any—for the proposition that a reference’s *silence* as to a particular implementation indicates a failure of that implementation to solve a particular problem. *See* Tr. 66:1–67:11. Because the cited documents do not mention the Splash 2 system, we find that they do not support Patent Owner’s contention that the system represented a failure of others.

For the reasons explained above, we conclude that Patent Owner’s evidence purportedly showing commercial success, industry praise, long-felt need, skepticism in the industry, and failure of others does not weigh in favor of nonobviousness of the challenged claims.

3. Conclusion

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, we determine that claims 1, 15, 18, 21, and 22 would have been obvious based on Splash2 under 35 U.S.C. § 103(a).

G. Obviousness Ground Based on Splash2 and RaPiD (Claims 8 and 9)

Petitioner asserts that dependent claims 8 and 9 are unpatentable over Splash2 and RaPiD. Pet. 56–69; *see* Ex. 1003 ¶¶ 303–339. Claim 8 recites that the calculation comprises “a JPEG image compression calculation” and

Appendix E

claim 9 recites that the calculation comprises “an MPEG image compression calculation.” With respect to claim 8, for example, Petitioner argues that RaPiD teaches a discrete cosine transform (DCT) “implemented on a systolic array of a reconfigurable computing system called the Reconfigurable Pipelined Datapath, or ‘RaPiD.’” Pet. 57. RaPiD discloses that “[t]he datapaths constructed in RaPiD are linear arrays of functional units communicating in mostly nearest-neighbor fashion” and “[s]ystolic algorithms, for example, map very well into RaPiD datapaths.” *Id.* (quoting Ex. 1009, 106). RaPiD explains that one application implemented on the disclosed system was a “2-D DCT . . . used in JPEG/MPEG data compression to convert an image from the spatial domain to the frequency domain.” *Id.* at 57–58 (quoting Ex. 1009, 110).

Petitioner asserts that a person of ordinary skill in the art would have combined the teachings of Splash2 and RaPiD such that “the functionality of the cell of RaPiD’s Figure 10 would be instantiated into Processing Elements of the Splash2 Unidirectional Systolic Array,” where “[t]he Processing Elements would perform the matrix multiply calculations described in RaPiD, and then pass the resulting data to the next Processing Element in a systolic fashion.” *Id.* at 59–60, 66–67. Petitioner provides numerous reasons why a person of ordinary skill in the art would have been motivated to do so. *Id.* at 66–69.

Patent Owner does not dispute that RaPiD teaches JPEG and MPEG image compression calculations, as recited in claims 8 and 9, but makes other arguments.

Appendix E

$$z_{mj} = \sum_{n=0}^{N-1} a_{mn} w_{nj}, \quad (5)$$

and thus

$$y_{ji} = \sum_{m=0}^{N-1} z_{mj} w_{mi} \quad (6)$$

As illustrated in the equations above, “both z_{mj} and y_{ji} are equivalent to . . . matrix multiplies,” but because “the z_{mj} values are produced in row-major order but required in column-major order, the results from the z_{mj} DCT must be transposed prior to computing y_{ji} .” *Id.* Figure 8 of RaPiD is reproduced below.




Figure 8: 2-D $N \times N$ DCT

Figure 8 depicts the 2-D DCT as two 1-D DCT calculations with a transpose in between. Dr. Stone explains in detail how the calculation of the dot product equations above requires multiple iterations of (1) multiplications between elements of a first matrix (a or z) with elements of a second matrix (w), and (2) addition of the result of that multiplication to a running sum from previous iterations.

Appendix E

Ex. 1003 ¶¶ 321–329; Ex. 1076 ¶¶ 21–22.²² According to Dr. Stone, the portion of annotated Figure 10 above highlighted in yellow shows “the output [of] the [arithmetic logic unit (ALU) being] looped back to the ALU input.” Ex. 1003 ¶ 327. That path is what “loops the intermediate, or running, sum from the output of the ALU to its input so that it may be used in the next iteration of the DCT calculation.” Ex. 1076 ¶ 25.

With respect to Figure 10, the 2-D DCT partitions an $M \times N$ image into 8×8 sub-images and performs two 1-D DCTs for each, using an 8×8 weight matrix, with the transpose in between. Ex. 1003 ¶¶ 307–311 (citing Ex. 1009, 111, Figs. 8–10). RaPiD discloses that

[s]ince a 2-D DCT performs two multiplies by the same weight matrix, W is loaded only once:

22. Again, we do not agree with Patent Owner that Petitioner’s Reply and Dr. Stone’s reply testimony are improper, as they respond directly to Patent Owner’s argument in the Response that RaPiD teaches a “bypass or forwarding path” rather than a “computational loop.” See Sur-Reply 8–9; Reply 36–40; Ex. 1076 ¶¶ 20–27; 37 C.F.R. § 42.23(b). Petitioner did not change its theory as to how RaPiD teaches a “computational loop,” relying in both its Petition and in its Reply on the portion of Figure 10 highlighted in yellow above. See *Apple Inc. v. Andrea Elecs. Corp.*, 949 F.3d 697, 706–707 (Fed. Cir. 2020) (concluding that a reply was proper where it did not change the “legal ground” asserted in the petition and “relie[d] on the same [teaching] from the same prior art reference to support the same legal argument”); *Chamberlain Grp., Inc. v. One World Techs., Inc.*, 944 F.3d 919, 925 (Fed. Cir. 2019) (“Parties are not barred from elaborating on their arguments on issues previously raised.”). Patent Owner also had the opportunity to cross-examine Dr. Stone about his reply declaration. See Ex. 2178.

Appendix E

one column per cell in both the first 8 cells and last 8 cells. The transpose in between matrix multiplies is performed with two local memories per cell: one to store products of the current sub-image and the other to store the products of the *previous* sub-image. During the computation of the current sub-image, the transpose of the previous sub-image computation is passed to the next 8 cells. The datapath for one RaPiD cell of a 2-D DCT is shown in Figure 10.

Ex. 1009, 111. Thus, as shown in Figure 10, “each cell in the array stores one column of weight data and receives both a stream of row image data and a stream of 1-D DCT data. The cell calculates the DCT for the current sub-image and passes on the DCT data for the previous sub-image along with row image data directly to the next cell.” Ex. 1003 ¶ 313; *see* Ex. 1009, 111, Fig. 10 (depicting the “[c]olumn of matrix W” stored in RAM and the “[r]ow of matrix A” and “[c]olumn of 1-D DCT results” received from the previous cell).

Patent Owner responds that the portion of annotated Figure 10 above highlighted in yellow is not a “computational loop” but rather a “bypass or forwarding path,” which “support[s] back to back execution of operations without stall, by forwarding (or bypassing) the output of an ALU to an input of the same or other ALU.” PO Resp. 89–90 (citing Ex. 2029 ¶¶ 41–48; Ex. 2112 ¶¶ 234–241; Ex. 2043, 301; Ex. 2044, 2, 6–7). According to Patent Owner, Dr. Stone also conceded during cross-examination that Figure 10 shows a bypass or forwarding path. *Id.* at 90 (citing Ex. 2063, 201:21–202:1).

Appendix E

Based on our review of the full record after trial, we find Patent Owner’s arguments unavailing. Dr. Stone’s explanation for how the calculation of the dot product equations requires adding products to a running sum is consistent with the figures and text of RaPiD, and we see no structure in Figure 10 other than the portion highlighted in yellow that would be capable of performing such functionality. Nor does Patent Owner point to any. RaPiD teaches, via the looping of the ALU output in Figure 10, a set of computations that is executed repeatedly a fixed number of times (i.e., until the end of the matrices is reached). *See supra* Section II.C.2.

We also are not persuaded that Dr. Stone conceded otherwise during cross-examination. He was asked simply whether “Figure 10 show[s] a bypass path or a forwarding path” and responded “Yes.” Ex. 2063, 201:21–202:2 (emphasis added). He was not asked to—and did not—identify any particular structure shown in Figure 10. Nor was he asked specifically whether the yellow highlighted portion is a bypass or forwarding path. Moreover, there is a bypass or forwarding path shown in Figure 10 that is not the yellow highlighted portion, namely the bottom wire labeled “Column of 1-D DCT results flows out” where DCT data from the previous cell is passed along to the next cell, bypassing the ALU functionality and memory shown in the top portion of the figure. Ex. 1003 ¶ 313; Ex. 1076 ¶ 27. Dr. Stone states that that wire was what he was referring to during cross-examination. Ex. 1076 ¶ 27. We see no contradiction between Dr. Stone’s declaration testimony and his statement that a bypass or forwarding path exists in Figure 10.

Appendix E

Second, Patent Owner argues that RaPiD does not teach passing computed data “seamlessly” between computational loops because Figure 10 “shows storage of results in memory (RAM) before being passed onto the next cell.” PO Resp. 100–101. As an initial matter, we note that Petitioner’s position, supported by the testimony of Dr. Stone, is that a person of ordinary skill in the art would have had reason to combine the references’ teachings such that “each Processing Element of the Splash2 Unidirectional Systolic Array would carry out the calculations of RaPiD Figure 10” and computed data would be “communicated directly from one Processing Element to the next.” Pet. 66 (emphases omitted); *see* Ex. 1003 ¶¶ 314–316, 319, 321, 330, 333. We find, for the reasons explained above, that in the unidirectional array implementation of Splash2, computed data is passed “seamlessly” (i.e., communicated directly, without intervening memory or other structures) between processing elements performing the computational loops. *See supra* Section II.E.2.a.3. Thus, Patent Owner’s argument attacking RaPiD individually with respect to the “seamlessly” limitation is not persuasive. *See In re Mouttet*, 686 F.3d 1322, 1332–33 (Fed. Cir. 2012) (holding that the test for obviousness is “what the combined teachings of the references would have suggested to those having ordinary skill in the art”); *Medichem, S.A. v. Rolabo, S.L.*, 437 F.3d 1157, 1166 (Fed. Cir. 2006) (explaining that in an obviousness analysis, “the prior art must be considered *as a whole* for what it teaches”); *In re Merck & Co., Inc.*, 800 F.2d 1091, 1097 (Fed. Cir. 1986) (“Non-obviousness cannot be established by attacking references individually where the rejection is based upon the teachings of a combination of references.”).

Appendix E

Regardless, though, the RAM that Patent Owner points to in Figure 10 is part of the cell (i.e., functional unit) itself, not between cells.²³ *See* Pet. 61–66; Reply 44; Ex. 1009, 106 (“RaPiD is a linear array of functional units which is configured to form a mostly linear computational pipeline. This array of functional units is divided into identical cells which are replicated to form a complete array.”); Ex. 2063, 201:11–18.

Third, Patent Owner raises various challenges to all of Petitioner’s obviousness grounds collectively, in particular to Petitioner’s reliance on the testimony of Dr. Stone. Patent Owner argues, for example, that Dr. Stone’s testimony is conclusory, grounded in hindsight bias, fails to disclose the underlying facts or data on which his opinions are based under 37 C.F.R. § 42.65(a), fails to “articulate reasons why or how a [person of ordinary skill in the art] would combine the references” or consider whether the asserted combinations were “feasible,” and “assumes only the benefits of his combinations and ignores any problems or drawbacks.” PO Resp. 18–30, 104–110, 114–116 (emphasis omitted). Many of Patent Owner’s arguments are premised on Petitioner failing

23. As explained above, for purposes of determining whether memory is within a processing element or between processing elements, the boundaries of a processing element are not arbitrary, but rather depend on the disclosure of the particular reference and how the reference describes the processing element and the communication of data to and from the processing element. *See supra* Section II.C.3. In this case, Figure 10 of RaPiD shows the “[n]etlist for one cell of [the] 2-D DCT.” Ex. 1009, 111 (emphasis added).

Appendix E

to “acknowledge” purported “shortcomings” of Splash2, such as the “need for looping to be performed by [the] Sun workstation” and Splash2’s lack of seamless passing of data and stream communication. *See, e.g., id.* at 28, 107–110, 115. Because we are not persuaded that Splash2 has any of those alleged deficiencies, those arguments are not persuasive. *See supra* Section II.E.2. Other arguments as to how Dr. Stone allegedly failed to understand the ’800 patent, prior art, and problems and issues in the art are premised on his testimony in a different proceeding challenging the ’687 patent and belied by the detailed explanations provided in his declarations, as explained herein. *See, e.g.,* PO Resp. 18–20, 22, 24–27, 108–109, 114–115 (citing Ex. 2065). We have reviewed Patent Owner’s arguments and cited evidence in the record and disagree with the remainder of Patent Owner’s assertions as well. We address the combination of Splash2 and RaPiD as representative.

Petitioner explains in detail what teachings of the two references it is relying on as teaching the various limitations of claims 8 and 9 and, importantly, explains exactly how a person of ordinary skill in the art would have combined those teachings, i.e., instantiating the cell functionality of Figure 10 in the processing elements of the Splash2 unidirectional array implementation. Pet. 59–60, 66–67. Petitioner asserts that the references are combinable in an obviousness combination because they are both analogous art to the ’800 patent and provides multiple reasons a person of ordinary skill in the art would have had for making the combination. *Id.* at 66–69. For example, RaPiD expressly cites the Splash

Appendix E

2 system as a “very successful example[]” of a field-programmable custom computing machine. Ex. 1009, 106 (citing a reference describing the Splash 2 system), 115; *see* Pet. 66–69. Splash2 also expressly states that the Splash 2 system would be advantageous for image processing applications (of which JPEG and MPEG image compression are examples), suggesting that implementing the RaPiD calculations in the Splash 2 system “would have been successful and efficient.” *See* Ex. 1007, 141–142, 162 (“Even though Splash 2 was not designed specifically for image processing, this platform possesses architectural properties that make it well suited for the computation and data transfer rates that are characteristic of this class of problems. Furthermore, the price/performance of this system makes it a competitive alternative to conventional real-time image processing systems.”); Pet. 68–69. Other documentation indicates that image compression techniques using DCT were being used at the time, such that an ordinarily skilled artisan would have been motivated to find “more efficient and powerful solutions” to do so (e.g., by implementing the RaPiD DCT calculation in the Splash 2 system). *See* Ex. 1049, 73; Pet. 68.

Petitioner’s assertions are supported by the testimony of Dr. Stone. *See* Ex. 1003 ¶¶ 300–339. That testimony is not based on hindsight bias, but rather the content of both asserted references, which Dr. Stone cites extensively in his analysis. *See id.* He explains the exact combination of teachings and the reasons for making the combination. *See id.* The reasons are not conclusory; they are explained in sufficient detail and premised on express statements in Splash2 and RaPiD connecting the two disclosures and

Appendix E

suggesting that their combination would be advantageous for multiple reasons. Patent Owner never addresses those identified reasons to combine in its papers or explains in any way why they are factually incorrect. Finally, we note that Dr. Stone stated during cross-examination that in addition to advantages of the various obviousness combinations, he considered potential disadvantages where relevant to the combination. *See, e.g.*, Ex. 2065, 145:9–146:5. We do not find any evidence in the record that combining RaPiD’s teaching of the DCT calculation with Splash2’s teaching of the unidirectional array implementation would have had any disadvantages that would have outweighed Petitioner’s stated advantages.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claims 8 and 9 would have been obvious based on Splash2 and RaPiD under 35 U.S.C. § 103(a).

H. Obviousness Ground Based on Splash2 and Jeong (Claim 20)

Petitioner asserts that dependent claim 20 is unpatentable over Splash2 and Jeong. Pet. 70–81; *see* Ex. 1003 ¶¶ 393–429. Claim 20 recites that the calculation comprises “an encryption algorithm.” Petitioner relies on Jeong for this limitation, arguing that Jeong teaches “a systolic implementation of a modular multiplication algorithm for encryption suitable for implementation in FGPAs.” Pet. 70–71 (citing Ex. 1061, 214, Fig. 2(a)). Petitioner provides numerous reasons why a person of

Appendix E

ordinary skill in the art would have been motivated to combine the teachings of Jeong and Splash2, including, for example, that Jeong “expressly cites the systolic structures as the platform to which its computing algorithms are mapped, and discloses the intent to make use of FPGAs as a possible means for implementation.” *Id.* at 79–81 (citing Ex. 1061, 211).

Patent Owner does not dispute that Jeong teaches “an encryption algorithm.” Instead, Patent Owner argues that Jeong does not teach two “computational loops,” as recited in parent claim 1, because the reference does not teach repeating a computation multiple times “for each piece of data.” PO Resp. 87–88. Patent Owner’s argument is premised on its proposed interpretation requiring repeated execution “per datum,” which we do not adopt. *See supra* Section II.C.2. Also, as explained above, we find that Splash2 teaches two “computational loops.” *See supra* Section II.E.2.

Likewise, with respect to the “seamlessly” limitation of claim 1, Patent Owner argues that similar to Splash2, “Jeong is ambiguous” regarding where results and inputs are stored. PO Resp. 104. Patent Owner further disputes the combination of references because Splash2 discloses a “linear” system and Jeong discloses a “non-linear” system. *Id.* at 113–114. Petitioner’s position, though, is that a person of ordinary skill in the art would have had reason to implement the “modular multiplication algorithm” of Jeong using the unidirectional systolic array described in Splash2—not the entirety of Jeong’s disclosed system. Pet. 72–73, 79–80; *see* Ex. 1003 ¶¶ 411, 412, 423, 426. That

Appendix E

assertion is supported sufficiently by the testimony of Dr. Stone, which is not conclusory or based on hindsight bias as Patent Owner contends. *See* Ex. 1003 ¶¶ 424–429; *supra* Section II.G.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claim 20 would have been obvious based on Splash2 and Jeong under 35 U.S.C. § 103(a).

I. Obviousness Ground Based on Splash2 and Chunky SLD (Claims 7, 17, and 24)

Petitioner asserts that dependent claims 7, 17, and 24 are unpatentable over Splash2 and Chunky SLD. -1606 Pet. 55–70; *see* Ex. 1003 ¶¶ 343–385. Claims 7, 17, and 24 limit the “calculation” in claim 1 to a particular type of calculation. For example, claim 17 recites that the calculation comprises “a search algorithm for an image search.” Petitioner relies on Chunky SLD for the limitation, arguing that the reference teaches an “automatic target resolution algorithm . . . used in the search for partially obscured images in synthetic aperture radar data” and “implemented on a linear systolic array instantiated in a Splash 2 system.” -1606 Pet. 55 (citing Ex. 1011, 192). Petitioner provides numerous reasons why a person of ordinary skill in the art would have been motivated to combine the teachings of Chunky SLD with Splash2’s teachings regarding the unidirectional array implementation. *Id.* at 64–67. For example, Petitioner notes that Chunky SLD “expressly cites the Splash 2

Appendix E

system as the platform to which its computing algorithms are mapped.” *Id.* at 66 (citing Ex. 1011, 192).

Patent Owner does not dispute that Chunky SLD teaches the limitations of dependent claims 7, 17, and 24 or dispute Petitioner’s asserted reasons to combine Splash2 and Chunky SLD (other than the general arguments regarding motivation addressed above). *See supra* Section II.G. Instead, Patent Owner argues that Chunky SLD does not teach two “computational loops,” as recited in parent claim 1, because the reference does not teach repeating a computation multiple times “for each piece of data.” PO Resp. 86–87. Patent Owner’s argument is premised on its proposed interpretation requiring repeated execution “per datum,” which we do not adopt. *See supra* Section II.C.2. Also, as explained above, we find that Splash2 teaches two “computational loops.” *See supra* Section II.E.2.

Likewise, with respect to the “seamlessly” limitation of claim 1, Patent Owner argues that “Chunky SLD is simply an algorithm deployed on Splash2, and therefore the same ambiguity present in Splash2 [regarding where data is stored] is also present.” PO Resp. 103–104. We disagree that Splash2 is ambiguous. *See supra* Section II.E.2.a.3. Further, Petitioner’s arguments regarding the combined teachings of Splash2 and Chunky SLD are supported sufficiently by the testimony of Dr. Stone, which is not conclusory or based on hindsight bias as Patent Owner contends. *See* -1606 Pet. 60–61, 64–67; Ex. 1003 ¶¶ 377–385; *supra* Section II.G.

Appendix E

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claims 7, 17, and 24 would have been obvious based on Splash2 and Chunky SLD under 35 U.S.C. § 103(a).

J. Obviousness Ground Based on Splash2 and Roccatano (Claims 2–5, 22, and 23)

Petitioner asserts that dependent claims 2–5, 22, and 23 are unpatentable over Splash2 and Roccatano. -1607 Pet. 60–78; *see* Ex. 1003 ¶¶ 250–299, 434–488. Claims 22 and 23 limit the “calculation” in claim 1 to a particular type of calculation, and claims 2–5 further limit the “subsequent and previous data dimensions” of the calculation in claim 1. For example, claim 2 recites that “said subsequent and previous data dimensions of said calculation comprise multiple vectors in said calculation.” Petitioner relies on Roccatano for the limitation, arguing that Roccatano teaches “a program for simulating molecular dynamics interactions of molecules and compounds” using “a systolic loop of processors as laid out on a 3-D grid” shown in Figure 2 of Roccatano. -1607 Pet. 60–62. Petitioner provides numerous reasons why a person of ordinary skill in the art would have been motivated to combine the teachings of Roccatano with Splash2’s teaching of the unidirectional array implementation, including, for example, that “Splash2 discloses that it is one of two reconfigurable systems that have achieved ‘supercomputer performance’ on applications that include molecular biology, which is the underlying application for Roccatano.” *Id.* at 68–70 (citing Ex. 1007, 6).

Appendix E

Patent Owner does not dispute that Roccatano teaches the limitations of dependent claims 2–5, 22,²⁴ and 23 or dispute Petitioner’s asserted reasons to combine Splash2 and Roccatano (other than the general arguments regarding motivation addressed above). *See supra* Section II.G. Instead, Patent Owner argues that Roccatano does not teach two “computational loops,” as recited in parent claim 1, because the reference does not teach repeating a computation multiple times “for each piece of data.” PO Resp. 87. Patent Owner’s argument is premised on its proposed interpretation requiring repeated execution “per datum,” which we do not adopt. *See supra* Section II.C.2. Also, as explained above, we find that Splash2 teaches two “computational loops.” *See supra* Section II.E.2.

Likewise, with respect to the “seamlessly” limitation of claim 1, Patent Owner argues that Roccatano does not teach passing computed data “seamlessly” because it uses multiple processors and, even if not, it is “ambiguous” where the processors store intermediate results, similar to the alleged teachings of Splash2. PO Resp. 101–102. Petitioner’s position, though, is that a person of ordinary skill in the art would have had reason to implement the calculations of Roccatano using the unidirectional systolic array described in Splash2—not using the multiple processors of Roccatano. -1607 Pet. 63–64, 67–73, 77–78; *see* Ex. 1003 ¶¶ 261, 264, 285, 461, 463–465, 467, 470, 483. Those assertions are supported sufficiently by the

24. We also conclude that Petitioner has proven that claim 22 is anticipated by Splash2. *See supra* Section II.E.3.

Appendix E

testimony of Dr. Stone, which is not conclusory or based on hindsight bias as Patent Owner contends. *See* Ex. 1003 ¶¶ 269–288, 480–486; *supra* Section II.G. Patent Owner’s arguments attacking Roccatano individually are not persuasive.

Based on all of the evidence of record, including evidence of secondary considerations of nonobviousness submitted by Patent Owner, *see supra* Section II.F.2, we determine that claims 2–5, 22, and 23 would have been obvious based on Splash2 and Roccatano under 35 U.S.C. § 103(a).

K. Obviousness Grounds Based on Combinations with Gaudiot (Claims 1–5, 7–9, 15, 17, 18, and 20–24)

Petitioner contends that claim 1 is unpatentable over Splash2 and Gaudiot, relying on Gaudiot in the alternative for one limitation of claim 1. *See* Pet. 53–56; -1606 Pet. 51–55; -1607 Pet. 56–59. Specifically, Petitioner states that “[t]o the extent one might argue that Splash2 does not sufficiently disclose that the edit distance calculation is ‘a data driven calculation that is implemented by said reconfigurable computing system at the at least one reconfigurable processor,’” as recited in claim 1, it would have been obvious to modify Splash2 to do so based on the teachings of Gaudiot. Pet. 53–56 (emphasis omitted). Petitioner does not rely on Gaudiot for any of the limitations of claim 1 that are disputed by Patent Owner, namely “computational loops” and passing computed data “seamlessly” between the computational loops. *See id.*; Tr. 51:4–18. As explained above, we conclude that

Appendix E

Splash2 discloses the data driven calculation limitation of claim 1 (which is not disputed by Patent Owner in its Response) and that Splash2 anticipates claim 1. *See supra* Section II.E.2. As such, we need not address Petitioner’s alternative ground of unpatentability of claim 1 based on the combination of Splash2 and Gaudiot. Likewise, for all of Petitioner’s other asserted obviousness grounds, Petitioner relies on certain references “with or without” Gaudiot. *See* Pet. 5; -1606 Pet. 5; -1607 Pet. 5. We address the grounds “without” Gaudiot above and need not address Petitioner’s alternative grounds based on combinations “with” Gaudiot. *See supra* Sections II.F–II.J.

L. Constitutionality Argument

Patent Owner argues that this proceeding should be dismissed because the instant panel “was unconstitutionally appointed in violation of the Appointments Clause of the U.S. Constitution, Article II, Section 2, Clause 2,” and “lacks the constitutional authority to enter a final decision.” Sur-Reply 21–23. We decline to consider Patent Owner’s constitutional challenge, as the issue has been addressed in *Arthrex, Inc. v. Smith & Nephew, Inc.*, 941 F.3d 1320 (Fed. Cir. 2019), *reh’g denied*, 933 F.3d 760 (Fed. Cir. 2020).

*Appendix E*III. CONCLUSION²⁵

Petitioner has demonstrated, by a preponderance of the evidence, that claims 1–5, 7–9, 15, 17, 18, and 20–24 are unpatentable. In summary:

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not shown Unpatentable
1, 15, 18, 21, 22	102(a), 102(b)	Splash2	1, 15, 18, 21, 22	
1, 15, 18, 21, 22	103(a)	Splash2	1, 15, 18, 21, 22	
1, 15, 18, 21, 22	103(a)	Splash2, Gaudiot ^{<?>}		
8, 9	103(a)	Splash2, RaPiD	8, 9	

25. Should Patent Owner wish to pursue amendment of the challenged claims in a reissue or reexamination proceeding subsequent to the issuance of this Decision, we draw Patent Owner's attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. See 84 Fed. Reg. 16,654 (Apr. 22, 2019). If Patent Owner chooses to file a reissue application or a request for reexamination of the challenged patent, we remind Patent Owner of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. See 37 C.F.R. §§ 42.8(a)(3), 42.8(b)(2).

Appendix E

8, 9	103(a)	Splash2, RaPiD, Gaudiot		
20	103(a)	Splash2, Jeong	20	
20	103(a)	Splash2, Jeong, Gaudiot		
7, 17, 24	103(a)	Splash2, Chunky SLD	7, 17, 24	
7, 17, 24	103(a)	Splash2, Chunky SLD, Gaudiot		
2-5, 22, 23	103(a)	Splash2, Roccatano	2-5, 22, 23	
2-5, 22, 23	103(a)	Splash2, Roccatano, Gaudiot		
Overall Outcome			1-5, 7-9, 15, 17, 18, 20-24	

IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that claims 1-5, 7-9, 15, 17, 18, and 20-24 of the '800 patent have been shown to be unpatentable;

FURTHER ORDERED that Petitioner's Motion to

312a

Appendix E

Exclude (Paper 60) is *granted-in-part*, *denied-in-part*, and *dismissed-in-part*;

FURTHER ORDERED that Patent Owner's Motion to Exclude (Paper 61) is *denied-in-part* and *dismissed-in-part*; and

FURTHER ORDERED that Exhibit 2170 is expunged from the record of this proceeding.

This is a final decision. Parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2018-01605
Patent 7,620,800 B2

FOR PETITIONER:

Joseph A. Micallef
Jason Greenhut
Scott Border
SIDLEY AUSTIN LLP
jmicallef@sidley.com
jgreenhut@sidley.com
sborder@sidley.com

FOR PATENT OWNER:

Alfonso Chan
Joseph DePumpo
SHORE CHAN DePUMPO LLP
achan@shorechan.com
jdepumpo@shorechan.com

313a

**APPENDIX F — DENIAL OF REHEARING OF
THE UNITED STATES COURT OF APPEALS FOR
THE FEDERAL CIRCUIT, FILED AUGUST 6, 2021**

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

2020-1993

FG SRC, LLC,

Appellant,

v.

MICROSOFT CORPORATION,

Appellee.

Appeal from the United States Patent and Trademark
Office, Patent Trial and Appeal Board in No. IPR2018-
01604.

**ON MOTION AND ON COMBINED PETITION
FOR PANEL REHEARING AND
REHEARING EN BANC**

Before MOORE, *Chief Judge*^{*}, NEWMAN, LOURIE, DYK,
PROST^{**}, O'MALLEY, REYNA, TARANTO, CHEN, HUGHES,
and STOLL, *Circuit Judges*.

* Chief Judge Kimberly A. Moore assumed the position of
Chief Judge on May 22, 2021.

** Circuit Judge Sharon Prost vacated the position of Chief
Judge on May 21, 2021.

314a

Appendix F

PER CURIAM.

ORDER

FG SRC LLC filed a combined petition for panel rehearing and rehearing en banc. The petition was referred to the panel that heard the appeal, and thereafter the petition for rehearing en banc was referred to the circuit judges who are in regular active service.

FG SRC LLC separately moved to dismiss the above-captioned appeal and remand the appeal to the United States Patent and Trademark Office. Microsoft Corporation responded in opposition, and FG SRC LLC replied.

Upon consideration thereof,

IT IS ORDERED THAT:

The petition for panel rehearing is denied.

The petition for rehearing en banc is denied.

The motion is denied.

The mandate of the court will issue on August 13, 2021.

FOR THE COURT

August 6, 2021

Date

/s/Peter R. Marksteiner
Peter R. Marksteiner
Clerk of Court

**APPENDIX G — DENIAL OF REHEARING OF
THE UNITED STATES COURT OF APPEALS FOR
THE FEDERAL CIRCUIT, FILED AUGUST 6, 2021**

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

2020-1925, 2020-1926

FG SRC, LLC,

Appellant,

v.

MICROSOFT CORPORATION,

Appellee.

Appeals from the United States Patent and Trade-
mark Office, Patent Trial and Appeal Board in Nos.
IPR2018-01601, IPR2018-01602, IPR2018-01603,
IPR2018-01605, IPR2018-01606, IPR2018-01607.

**ON MOTION AND ON COMBINED PETITION
FOR PANEL REHEARING AND
REHEARING EN BANC**

Before MOORE, *Chief Judge*^{*}, NEWMAN, LOURIE, DYK,
PROST^{**}, O'MALLEY, REYNA, TARANTO, CHEN, HUGHES,
and STOLL, *Circuit Judges*.

* Chief Judge Kimberly A. Moore assumed the position of
Chief Judge on May 22, 2021.

** Circuit Judge Sharon Prost vacated the position of Chief
Judge on May 21, 2021.

316a

Appendix G

PER CURIAM.

ORDER

FG SRC LLC filed a combined petition for panel rehearing and rehearing en banc. The petition was referred to the panel that heard the appeal, and thereafter the petition for rehearing en banc was referred to the circuit judges who are in regular active service.

FG SRC LLC separately moved to dismiss the above-captioned appeal and remand the appeal to the United States Patent and Trademark Office. Microsoft Corporation responded in opposition, and FG SRC LLC replied.

Upon consideration thereof,

IT IS ORDERED THAT:

The petition for panel rehearing is denied.

The petition for rehearing en banc is denied.

The motion is denied.

The mandate of the court will issue on August 13, 2021.

FOR THE COURT

August 6, 2021

Date

/s/Peter R. Marksteiner

Peter R. Marksteiner
Clerk of Court