UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Intel Corporation,
Petitioner,

v.

Qualcomm Incorporated,
Patent Owner.

Case IPR2018-01334[1]
U.S. Patent No. 8,838,949

PATENT OWNER'S NOTICE OF APPEAL

---

[1] IPR2018-01335 and IPR2018-01336 have been consolidated with the instant proceeding.

Pursuant to 37 C.F.R § 90.2(a), Patent Owner Qualcomm Incorporated ("Qualcomm" or "Patent Owner") hereby appeals to the United States Court of Appeals for the Federal Circuit the Patent Trial and Appeal Board's Final Written Decision on Remand in these consolidated IPRs, entered on March 21, 2023 (Paper 49), and from all underlying orders, decisions, rulings and opinions that are adverse to Qualcomm.

In accordance with 37 C.F.R § 90.2(a)(3)(ii), Patent Owner further indicates that the issues on appeal include, but are not limited to: the Board's claim construction of the term "hardware buffer" as used in claims 1-9 and 12 of U.S. Patent No. 8,838,949; the Board's determination of unpatentability of claims 1-9, 12, 16, and 17 of U.S. Patent No. 8,838,949 under 35 U.S.C. § 103(a); whether the Board followed and complied with the Federal Circuit's mandate; whether the Board followed and complied with the Federal Circuit's directions on remand; whether the Board followed and complied with the Administrative Procedure Act; any finding or determination supporting or related to any of these issues; as well as all other issues decided adversely to Patent Owner in any orders, decisions, rulings, and opinions.

Simultaneous with this submission, a copy of this Notice of Appeal is being filed with the United States Patent and Trademark Office by way of hand delivery to the following address:

Office of the General Counsel
United States Patent and Trademark Office
10B20, Madison Building East
600 Dulany Street
Alexandria, Virginia 22314

In addition, this Notice of Appeal, along with a copy of the Final Written Decision on Remand, is being filed electronically with the United States Court of Appeals for the Federal Circuit along with the required docketing fees.

Respectfully submitted,

Date: March 22, 2023

/ Joshua R. Nightingale /

Joshua R. Nightingale, Reg. No. 67,865
JONES DAY
500 Grant Street, Suite 4500
Pittsburgh, PA 15219

Counsel for Patent Owner

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. § 42.6(e), the undersigned certifies that on March 22, 2023, a complete and entire copy of this **PATENT OWNER'S NOTICE OF APPEAL** has been served in its entirety by e-mail on the following counsel of record for petitioner:

David Cavanaugh
david.cavanaugh@wilmerhale.com

Tom Anderson
tom.anderson@wilmerhale.com

Joseph Haag
joseph.haag@wilmerhale.com

Date: March 22, 2023

/ Joshua R. Nightingale /
Joshua R. Nightingale, Reg. No. 67,865
JONES DAY
500 Grant Street, Suite 4500
Pittsburgh, PA 15219

Counsel for Patent Owner

Trials@uspto.gov
571-272-7822

Paper 49
Date: March 21, 2023

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,
Petitioner,

v.

QUALCOMM INCORPORATED,
Patent Owner.

IPR2018-01334[1]
Patent 8,838,949 B2

Before TREVOR M. JEFFERSON, DANIEL J. GALLIGAN, and AARON W. MOORE, *Administrative Patent Judges.*

Opinion for the Board filed by *Administrative Patent Judge* DANIEL J. GALLIGAN.

Opinion Concurring-in-Part and Dissenting-in-Part filed by *Administrative Patent Judge* AARON W. MOORE.

GALLIGAN, *Administrative Patent Judge*.

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

[1] IPR2018-01335 and IPR2018-01336 have been consolidated with the instant proceeding.

I. INTRODUCTION

This case is on remand from the Court of Appeals for the Federal Circuit to address the patentability of claims 1–9, 12, 16, and 17 of U.S. Patent No. 8,838,949 B2 ("the '949 patent," Ex. 1001). *See Intel Corp. v. Qualcomm Inc.*, 21 F.4th 801, 814 (Fed. Cir. 2021).

For the reasons discussed below, we determine that Petitioner has proven by a preponderance of the evidence that claims 1–9, 12, 16, and 17 are unpatentable. *See* 35 U.S.C. § 316(e) ("In an inter partes review instituted under this chapter, the petitioner shall have the burden of proving a proposition of unpatentability by a preponderance of the evidence.").

A. Background

On July 3, 2018, Intel Corporation ("Petitioner") filed three petitions challenging claims of the '949 patent as follows: IPR2018-01334 (claims 1–9, 22, and 23), IPR2018-01335 (claims 10–17), and IPR2018-01336 (claims 18–21). Petitioner asserts that the claims are unpatentable on the following grounds:

In IPR2018-01334:

Claim(s) Challenged	35 U.S.C. §[2]	Reference(s)/Basis
1–9, 22, 23	103(a)	Bauer,[3] Svensson,[4] Kim[5]

[2] The Leahy-Smith America Invents Act ("AIA") included revisions to 35 U.S.C. §§ 103 and 112 that became effective after the filing of the application for the '949 patent. Therefore, we apply the pre-AIA versions of these sections.
[3] US 2006/0288019, published Dec. 21, 2006 (Ex. 1009).
[4] US 7,356,680 B2, issued Apr. 8, 2008 (Ex. 1010).
[5] Korean Patent Application Publication No. 10-2002-0036354, published May 16, 2002 (Ex. 1011). References to Kim in this Decision are to the English translation provided by Petitioner as Exhibit 1012.

In IPR2018-01335:

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
10–15	103(a)	Bauer, Svensson, Kim
16, 17	103(a)	Bauer, Svensson, Kim, Zhao[6]

In IPR2018-01336:

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
18–21	103(a)	Bauer, Svensson, Kim, Lim[7]

We instituted review in each case on all grounds presented. IPR2018-01334, Paper 10 ("Dec. on Inst."), 29; IPR2018-01335, Paper 10 ("1335 Dec. on Inst."),[8] 38; IPR2018-01336, Paper 10 ("1336 Dec. on Inst."), 32.

After institution, we consolidated IPR2018-01335 and IPR2018-01336 with IPR2018-01334 and terminated IPR2018-01335 and IPR2018-01336. Paper 12.

During the trial, Qualcomm Incorporated ("Patent Owner" or "Qualcomm") filed a Response (Paper 16, "PO Resp."), Petitioner filed a Reply (Paper 21, "Pet. Reply"), and Patent Owner filed a Sur-reply (Paper 25, "PO Sur-reply").

An oral hearing was held on December 12, 2019, a transcript of which appears in the record. Paper 29 ("Tr.").

[6] US 2007/0140199 A1, published June 21, 2007 (Ex. 1013).
[7] US 7,203,829 B2, published Apr. 10, 2007 (Ex. 1014).
[8] We use prefixes "1335" and "1336" to denote papers and exhibits from IPR2018-01335 and IPR2018-01336, respectively. We do not use a prefix for papers and exhibits from IPR2018-01334.

B. Final Written Decision and Federal Circuit Appeal

We issued a Final Written Decision holding that Petitioner had proven claims 10, 11, 13–15, and 18–23 unpatentable but had not proven claims 1–9, 12, 16, and 17 unpatentable. Paper 30 ("Final Decision" or "Final Dec.") at 63–64.

Petitioner filed a Notice of Appeal of the Final Decision with the United States Court of Appeals for the Federal Circuit as to our determination with respect to claims 1–9, 12, 16, and 17. Paper 31. Patent Owner filed a Notice of Cross-Appeal as to our determination with respect to claims 10, 11, 13–15, and 18–23. Papers 32, 33.

On December 28, 2021, the Federal Circuit issued a decision in the appeal vacating our Final Decision as to claims 1–9, 12, 16, and 17 and remanding for further proceedings. *Intel*, 21 F.4th at 814. In particular, the Federal Circuit "vacate[d] the Board's construction of the term 'hardware buffer,' its determination that claims 1–9 and 12 were non-obvious over the prior art, and its conclusion that claims 16–17 lacked sufficient corresponding structure in the specification, and . . . remand[ed] for further proceedings consistent with this opinion." *Id.* At the Federal Circuit, Patent Owner dropped its challenge to our determination of unpatentability as to claims 10, 11, 13–15, and 18–23. *See id.* at 806 ("The Board ruled that Intel had proved the unpatentability of claims 10, 11, 13–15, and 18–23, but Qualcomm, despite filing a cross-appeal to raise the issue, no longer challenges that ruling.").

The Federal Circuit issued its mandate on February 3, 2022.

C. Remand Proceedings

On February 17, 2022, we held a conference call with counsel for the parties to discuss a schedule on remand. *See* Paper 34 at 2–4. Before the call, the parties met and conferred and agreed to a schedule, which we adopted. *See* Paper 34; Ex. 3001. The parties filed the following briefing according to that schedule: Petitioner's Opening Brief on Remand (Paper 35, "Pet. Remand Br."); Patent Owner's Response Brief on Remand (Paper 37, "PO Remand Br."); Petitioner's Reply on Remand (Paper 39, "Pet. Remand Reply"); and Patent Owner's Sur-reply on Remand (Paper 40, "PO Remand Sur-reply").

An oral hearing to discuss the issues on remand was held on August 4, 2022, a transcript of which appears in the record. Paper 46 ("Remand Tr.").

D. Real Parties in Interest

Petitioner identifies itself and Apple Inc. as real parties in interest. Pet. 2. Patent Owner identifies itself as the real party in interest. Paper 4, 2.

E. The '949 Patent and Illustrative Claim

The '949 patent generally relates to loading software from one processor to another in a multi-processor system. Ex. 1001, code (57). One example disclosed in the '949 patent involves loading modem image executable data by first retrieving and processing an image header, which "includes information used to identify where the modem image executable data is to be eventually placed into the system memory of the secondary processor." Ex. 1001, 8:9–21. Figure 3 of the '949 patent is reproduced below.

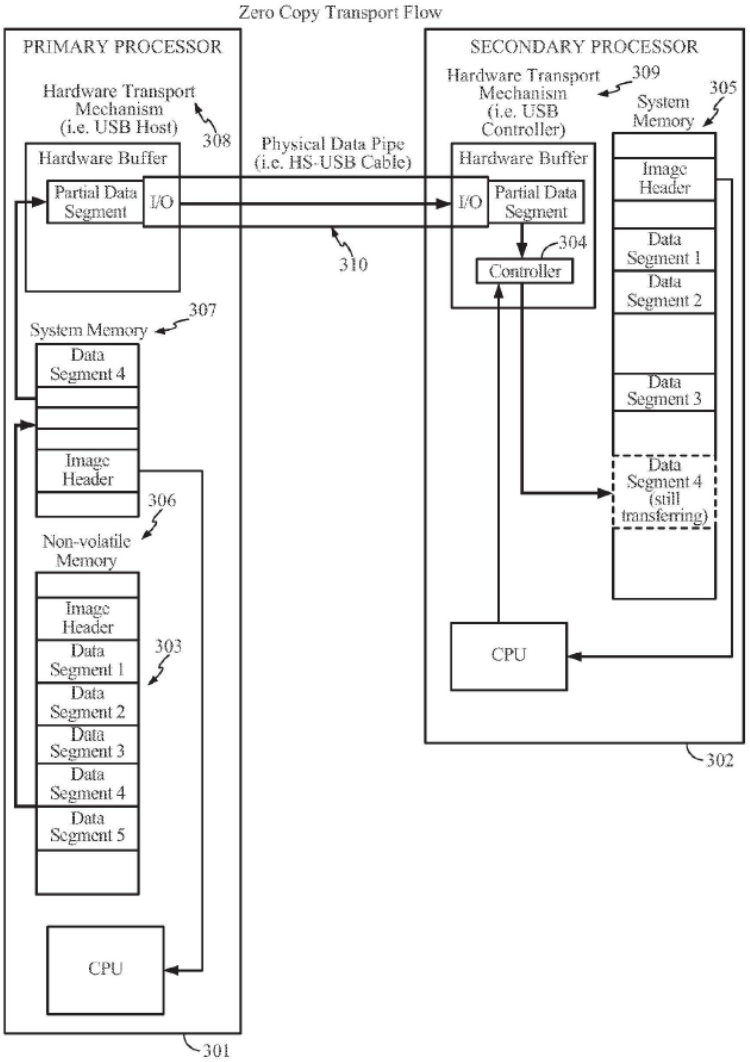

FIG. 3

Figure 3 shows "operational flow for an exemplary loading process for loading an executable image from a primary processor to a secondary processor according to one aspect of the present disclosure." Ex. 1001, 4:10–13. Referring to various components depicted in Figure 3, the '949 patent discloses the following:

> The header information is used by the secondary processor 302 to program the scatter loader/direct memory access controller 304 receive address when receiving the actual executable data. Data segments are then sent from system memory 307 to the primary hardware transport mechanism 308. The segments are

> then sent from the hardware transport mechanism 308 of the primary processor 301 to a hardware transport mechanism 309 of the secondary processor 302 over an inter-chip communication bus 310 (e.g., a HS-USB cable.) The first segment transferred may be the image header, which contains information used by the secondary processor to locate the data segments into target locations in the system memory of the secondary processor 305. The image header may include information used to determine the target location information for the data.

Ex. 1001, 8:21–35. Of particular relevance on remand is that the '949 patent discloses scatter loading data segments "directly" from a "hardware buffer" to their final locations in the secondary processor's "system memory." Ex. 1001, 2:58–63.

As noted above, claims 1–9, 12, 16, and 17 are at issue on remand. Claims 1 and 16 are independent claims, claims 2–9 depend from claim 1, and claim 17 depends from claim 16. Claim 12 depends from independent claim 10, which is unpatentable. *See* Final Dec. 50, 63; *Intel*, 21 F.4th at 806 (noting that Patent Owner dropped its challenge to the Board's unpatentability determination for claim 10).

Claims 1, 10, 12, and 16 are reproduced below.

> 1. A multi-processor system comprising:
> a secondary processor comprising:
> system memory and a hardware buffer for receiving an image header and at least one data segment of an executable software image, the image header and each data segment being received separately, and
> a scatter loader controller configured:
> to load the image header; and
> to scatter load each received data segment based at least in part on the loaded image header, directly from the hardware buffer to the system memory;

a primary processor coupled with a memory, the memory storing the executable software image for the secondary processor; and

an interface communicatively coupling the primary processor and the secondary processor, the executable software image being received by the secondary processor via the interface.

10. A method comprising:

receiving at a secondary processor, from a primary processor via an inter-chip communication bus, an image header for an executable software image for the secondary processor that is stored in memory coupled to the primary processor, the executable software image comprising the image header and at least one data segment, the image header and each data segment being received separately;

processing, by the secondary processor, the image header to determine at least one location within system memory to which the secondary processor is coupled to store each data segment;

receiving at the secondary processor, from the primary processor via the inter-chip communication bus, each data segment; and

scatter loading, by the secondary processor, each data segment [directly[9]] to the determined at least one location within the system memory, and each data segment being scatter loaded based at least in part on the processed image header.

12. The method of claim 10 further comprising loading the executable software image directly from a hardware buffer to the system memory of the secondary processor without copying data between system memory locations.

[9] The issued patent recites "reedy," which appears to be a printing error. The April 30, 2014 claim listing submitted by the applicants during prosecution states "directly."

16. An apparatus comprising:

means for receiving at a secondary processor, from a primary processor via an inter-chip communication bus, an image header for an executable software image for the secondary processor that is stored in memory coupled to the primary processor, the executable software image comprising the image header and at least one data segment, the image header and each data segment being received separately;

means for processing, by the secondary processor, the image header to determine at least one location within system memory to which the secondary processor is coupled to store each data segment;

means for receiving at the secondary processor, from the primary processor via the inter-chip communication bus, each data segment; and

means for scatter loading, by the secondary processor, each data segment directly to the determined at least one location within the system memory, and each data segment being scatter loaded based at least in part on the processed image header.

II. ANALYSIS

A. Level of Ordinary Skill in the Art

We apply the following level of ordinary skill in the art: "a Master's degree in Electrical Engineering, Computer Engineering, or Computer Science plus two years of experience in mobile device architecture and multi-processor systems, or a Bachelor's degree in one of those fields plus four years of experience in mobile device architecture and multiprocessor systems." *See* Final Dec. 8.

B. Claim Interpretation

In this proceeding, we give the claims the broadest reasonable construction in light of the Specification of the '949 patent. 37 C.F.R. § 42.100(b) (2018); *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 (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)); *see also Intel*, 21 F.4th at 808–09 (stating that the broadest reasonable interpretation standard applies).

The claim construction issue on remand for claims 1–9 and 12 concerns the scope of the term "hardware buffer." *See Intel*, 21 F.4th at 810. For claims 16 and 17, constructions for the means-plus-function limitations are at issue. *See Intel*, 21 F.4th at 812–14.

1. Hardware Buffer

a. Pre-Appeal Final Written Decision

The term "hardware buffer" appears in independent claim 1, and claims 2 and 8, which depend from claim 1, and in claim 12, which depends from independent claim 10. Claim 1 recites, in part, "a secondary processor comprising: system memory and a hardware buffer for receiving an image header and at least one data segment of an executable software image" and "a scatter loader controller configured: to load the image header; and to scatter load each received data segment based at least in part on the loaded image header, directly from the hardware buffer to the system memory." Claim 12 recites, "The method of claim 10 further comprising loading the executable software image directly from a hardware buffer to the system memory of the secondary processor without copying data between system memory locations."

In the Final Decision, we analyzed the intrinsic evidence and determined that "the 'hardware buffer' limitations of independent claim 1 and its dependent claims (2–9) and dependent claim 12 'should not be read so broadly as to encompass' the use of a temporary buffer." Final

Dec. 10–17 (quoting *SciMed Life Systems, Inc. v. Advanced Cardiovascular Systems, Inc.*, 242 F.3d 1337, 1343 (Fed. Cir. 2001)). Applying that interpretation, we determined that Petitioner had not proven unpatentability of claims 1–9 and 12. Final Dec. 55–56.

Petitioner appealed our determination as to claims 1–9 and 12. *See Intel*, 21 F.4th at 808–12.

b. Federal Circuit's Decision

The Federal Circuit found our construction "wanting" and "inadequate" and vacated our determination as to claims 1–9 and 12. *Intel*, 21 F.4th at 808–12. The Federal Circuit stated that "it is clear from the claim language that the claim term ['hardware buffer'] has meaning, but it is unclear what that meaning is. There is no definition to be found in the intrinsic evidence." *Id.* at 809. Furthermore, "the determination of that meaning (or range of reasonable meanings) depends on understanding what the intrinsic evidence makes clear is the substance of the invention—what the inventor 'intended to envelop,'" which "in some cases is usefully clarified by expert testimony (as long as that testimony is consistent with the intrinsic evidence)." *Id.* (quoting *Phillips v. AWH Corp.*, 415 F.3d 1303, 1316 (Fed. Cir. 2005) (en banc)). The Federal Circuit concluded that we "did not do enough to reach and articulate that understanding." *Id.* Below, we address the Federal Circuit's decision in more detail in light of the parties' arguments on remand.

c. Determination on Remand

On remand, the parties provide additional briefing and evidence on the proper interpretation of the term "hardware buffer." Pet. Remand Br. 5–13; PO Remand Br. 3–15. Petitioner argues that the broadest reasonable

interpretation of the term "hardware buffer" is "memory that is physically separate from the memory into which the software image is loaded for execution." Pet. Remand Br. 5. Patent Owner counters that the proper interpretation of this term is "a permanent, dedicated buffer that is distinct from system memory." PO Remand Br. 3. For the reasons explained below, we determine that Patent Owner's temporal restriction on the buffer ("permanent") is not required of the broadest reasonable interpretation.

As discussed above, the Federal Circuit stated that the proper construction "depends on understanding what the intrinsic evidence makes clear is the substance of the invention—what the inventor 'intended to envelop.'" *Intel*, 21 F.4th at 809 (quoting *Phillips*, 415 F.3d at 1316). More particularly, the court stated that we should provide "a more substance-focused analysis . . . of what the intrinsic evidence shows the asserted advance to be and how, concretely, the 'hardware buffer' relates to that advance." *Id.* at 811. Below, we address the claims, the Specification, and the prosecution history, as well as additional evidence that informs the meaning of the term "hardware buffer."

i. Claim Language

Turning to the claim language first, the Federal Circuit noted that "it is clear from the claim language that the claim term at issue [('hardware buffer')] has meaning, but it is unclear what that meaning is" and that "[t]here is no definition to be found in the intrinsic evidence." *Intel*, 21 F.4th at 809. The Federal Circuit, nevertheless, "reach[ed] three conclusions from the claim language." *Id.* "*First*," the Federal Circuit stated, "because every buffer in our (physical) world is ultimately implemented on a physical device (*i.e.*, hardware), a 'hardware buffer' must mean something more than

just a 'buffer implemented in hardware,' as [Petitioner] urges, or else the word 'hardware' would be erased from the claims." *Id.* "*Second*, because claim 1 requires both a 'system memory' and a 'hardware buffer,' there must be some distinction between those two concepts." *Id.* at 810.

> *Third*, because claim 2 requires loading the executable software image "directly from the hardware buffer to the system memory of the secondary processor without copying data between system memory locations on the secondary processor," the meaning of "hardware buffer" relates to the ability to move the software image "directly" to the second processor's system memory and to avoid "copying data between system memory locations."

Id.

The Federal Circuit then noted that these "conclusions from the claim language advance the claim-construction inquiry only so far" but "do not, on their own, provide a concrete basis for a clarifying definition of 'hardware buffer.'" *Intel*, 21 F.4th at 810. The court stated that "an analysis of the specification" is needed "to arrive at an understanding of what it teaches about what a 'hardware buffer' is, based on both how it uses relevant words and its substantive explanations." *Id.* The Federal Circuit further stated that "[w]hat is needed in this case is a more substance-focused analysis than is yet present, in the Board's opinion or in the present record . . ., of what the intrinsic evidence shows the asserted advance to be and how, concretely, the 'hardware buffer' relates to that advance." *Id.* at 811.

With this guidance in mind, we turn next to the Specification of the '949 patent.

ii. Specification

The written description of the '949 patent uses the term "hardware buffer" only three times, and Figure 3 includes the label "Hardware Buffer"

in each of the primary processor and the secondary processor. Ex. 1001, 2:58–63, 9:37–41, Fig. 3. The first two instances of the term "hardware buffer" are in the following passage: "The system includes a secondary processor having a system memory and a hardware buffer for receiving at a least a portion of an executable software image. The secondary processor includes a scatter loader controller for loading the executable software image directly from the hardware buffer to the system memory." Ex. 1001, 2:58–63. This passage mirrors the claim language and, thus, does not advance the inquiry any further than the claims themselves. The other passage mentioning "hardware buffer" states the following in reference to Figure 3: "In one aspect, the executable software image is loaded into the system memory of the secondary processor without an entire executable software image being stored in the hardware buffer of the secondary processor." Ex. 1001, 9:37–41. This passage sheds some light on a different aspect of the disclosure of the '949 patent, namely, not having to load the entire image to a buffer first before loading the image to its ultimate destination, but the passage does not itself help define the term "hardware buffer."

In the Final Decision, we focused on the '949 patent Specification's apparent differentiation between its "hardware buffer" and the use of a temporary buffer in the prior art, and we concluded that the "hardware buffer" limitations should not be read to encompass the use of a temporary buffer. Final Dec. 15–17 (citing Ex. 1001, 2:23–34, 4:43–47, 5:31–35).

The Federal Circuit stated that, "[a]lthough the Board correctly noted that the specification describes prior art teaching of use of 'temporary' buffers, it did not explain precisely what 'temporary' means or how the

patent-described prior-art use relates to the alleged invention." *Intel*, 21 F.4th at 810. The court further stated that "the Board did not analyze exactly how the use of a hardware buffer, as claimed by Qualcomm, would address the concerns about the prior-art temporary buffers raised in" three passages in the Specification. *Id.* at 810–11 (citing Ex. 1001, 2:23–34, 4:43–47, 5:31–35). We now provide such analysis in light of the parties' arguments on this issue.

Petitioner argues that the Specification of the '949 patent distinguishes other aspects of the prior art rather than the use of temporary buffers per se. Pet. Remand Br. 7–11. For example, the Background section of the '949 patent states the following:

> In a system in [w]hich the software image is loaded onto a target "secondary" processor from a first "primary" processor, one way of performing such loading is to allocate a temporary buffer into which each packet is received, and *each packet would have an associated packet header information along with the payload*. The payload in this case would be the actual image data. From the temporary buffer, some of the processing may be done over the payload, and then the payload would get copied over to the final destination. The temporary buffer would be some place in system memory, such as in internal random-access-memory (RAM) or double data rate (DDR) memory, for example.

Ex. 1001, 2:23–34 (emphasis added). Petitioner argues that this passage is distinguishing between a situation in which each packet has its own header and payload and the situation in the claims requiring that the header and data segments are received separately, rather than distinguishing the use of a temporary buffer. Pet. Remand Br. 10.

Another passage from the '949 patent provides the following:

> In one exemplary aspect a direct scatter load technique is disclosed for loading a segmented image from a primary processor's non-volatile memory to a secondary processor's volatile memory. As discussed further below, the direct scatter load technique avoids use of a temporary buffer. For instance, in one aspect, rather than employing a packet-based communication in which the image is communicated via *packets that each include a respective header*, the raw image data is loaded from the primary processor to the secondary processor.

Ex. 1001, 4:43–52 (emphasis added). Petitioner argues that this passage also shows a distinction between receiving header and data together, as in the '949 patent's characterization of the prior art, and receiving the header and data segments separately, as in the claims. Pet. Remand Br. 10–11.

The '949 patent also mentions the use of a temporary buffer in the following passage:

> Thus, conventional techniques employing a *temporary buffer for the entire image*, and the packet header handling, etc., are bypassed in favor of a more efficient direct loading process. Thus, the exemplary load process of FIG. 3 does not require the intermediate buffer operations traditionally required for loading a software image from a primary processor to a secondary processor. Instead of scatter loading from a *temporary buffer holding the entire image*, the exemplary load process of FIG. 3 allows for direct scatter load the image segments to their respective target destinations directly from the hardware to the system memory.

Ex. 1001, 9:43–54 (emphasis added). Petitioner argues that this passage "distinguishes systems in which the ***entire executable software image*** is copied into a temporary buffer." Pet. Remand Br. 9.

We agree with Petitioner that these passages provide additional context to show that the '949 patent does not necessarily distinguish its invention from the use of a temporary buffer, per se.

The '949 patent also discusses how its data transfer methodology avoids extra memory copy operations, a point the Federal Circuit noted. *See Intel*, 21 F.4th at 811 (noting "other specification passages that support an understanding that use of a 'hardware buffer' relates to one of the key claimed advances of the invention—the elimination of 'extra memory copy operations'" (citing Ex. 1001, 7:16, 7:27–30, 9:42–46)). As mentioned in the preceding section, the Federal Circuit stated that "the meaning of 'hardware buffer' relates to the ability to move the software image 'directly' to the second processor's system memory and to avoid 'copying data between system memory locations,'" as recited in claim 2. *Intel*, 21 F.4th at 810.

There is no question that a focus of the '949 patent is on eliminating extra memory copy operations, but more precisely on extra memory copy operations performed by the secondary processor. For example, claim 2 recites, "The multi-processor system of claim 1 in which the scatter loader controller is configured to load the executable software image directly from the hardware buffer to the system memory of the secondary processor without copying data between system memory locations on the secondary processor." Claim 1 provides that the scatter loader controller is part of the secondary processor. Similarly, the '949 patent states that "no extra memory copy operations occur *in the secondary processor* in the above aspect," referring to the operation described with respect to Figure 3. Ex. 1001, 9:42–43; *see also* PO Remand Br. 6 ("Specifically, the '949 patent states that in performing the direct transfer using the hardware buffer, 'no extra memory copy operations occur in the secondary processor.'" (quoting Ex. 1001, 9:42–43)). To the extent this particular aspect of the '949 patent's

disclosure informs the proper meaning of the term "hardware buffer," we do not see a relevant distinction between the operation of the '949 patent and the asserted prior art. As explained more fully below, both the '949 patent and the asserted prior art describe that the primary processor performs a memory copy operation, followed by a memory copy operation by the secondary processor.

Based on the foregoing, we do not view the Specification as imposing a temporal restriction on the nature of the "hardware buffer." As to "what the intrinsic evidence shows the asserted advance to be and how, concretely, the 'hardware buffer' relates to that advance," *Intel*, 21 F.4th at 811, the prosecution history provides particular insight.

iii. Prosecution History

Petitioner argues that the prosecution history shows that certain language added by amendment, such as the separate receipt of the header and data, reflects the "asserted advance" of the '949 patent. Pet. Remand Br. 11–13. In particular, during prosecution of the '949 patent, the Examiner rejected all then-pending claims as anticipated by International Publication WO 2006/077068 A2 (Ex. 1003), which is the Patent Cooperation Treaty (PCT) application publication that claims priority to Svensson. Ex. 1004 (July 19, 2013 Non-Final Rejection 2–4); *see* Ex. 1003 ("Svensson PCT"), code (30). In rejecting the pending claims, the Examiner found that Svensson PCT's intermediate storage area (ISA) defined within memory 108 describes the recited "hardware buffer." Ex. 1004 at 2; *see* Pet. Remand Br. 11–12. This is the same ISA of Svensson and Bauer that Petitioner contends is a "hardware buffer." *See* Pet. Remand Br. 11–12. As Petitioner notes (Pet. Remand Br. 12), the applicants responded to this

rejection by amending claim 1 to recite "the image header and each data segment being received separately" and a scatter loader controller configured "to load the image header" and "to scatter load each received data segment based at least in part on the loaded image header." Ex. 1005 at 2.

The applicants then provided arguments distinguishing the disclosure of Svensson PCT. According to the applicants, "FIG. 3 of Svensson [PCT] arguably discloses that the software includes a header and a data segment. As disclosed in col. 8 lines 10-19 of Svensson [PCT], *each code and/or data to be transferred includes a header*." Ex. 1005 at 8. Svensson PCT does not have dual-column pages like issued patents, so we understand the applicants to be referring to page 8, lines 10–19 of Svensson PCT, which is the same as column 6, lines 13–25 of Svensson. The applicants continued, stating the following:

> In contrast to Svensson [PCT], claim 1 recites that the image header and each data segment are received separately. Applicants submit that separately receiving the image header and each data segment, as recited in claim 1, is patentably distinguishable from receiving the data and the associated header, as disclosed in Svensson [PCT]. Therefore, because Svensson [PCT] expressly discloses that each code and/or data to be transferred includes a header and fails to disclose that the image header and each data segment are received separately, applicants submit that Svensson [PCT] cannot teach or suggest "the image header and each data segment are received separately," as recited in claim 1.

Ex. 1005 at 8–9. Here, the applicants rely on the separate receipt of the image header and the data as a distinguishing feature.

Furthermore, the applicants stated the following:

> Moreover, the cited portions of Svensson [PCT] disclose that several blocks (e.g., data blocks with headers) are concatenated in the intermediate storage area (ISA). Specifically, the data blocks are concatenated in the intermediate storage area prior to being transferred to the memory (*see* Svensson [PCT], step 220). In contrast, claim 1 recites that each data segment is scatter loaded based at least in part on the loaded image header. That is, the individual data segments of claim 1 are not concatenated with the header files. Rather, the image header file is loaded into memory to scatter load *each data segment directly* from the hardware buffer to the system memory.
>
> Applicants submit that loading *each data segment directly* from the hardware buffer to the system memory, as recited in claim 1, is patentably distinguishable from concatenating the data blocks and headers in the intermediate storage area and then transferring the concatenated data to the memory, as recited in Svensson [PCT]. Thus, because Svensson [PCT] fails to teach or suggest scatter loading *each data segment directly* from the hardware buffer to the system memory, applicants submit that Svensson [PCT] cannot teach or suggest "a scatter loader controller configured to load the image header and to scatter load each received data segment directly from the hardware buffer to the system memory," as recited in claim 1.

Ex. 1005 at 9.

Petitioner argues that the "prosecution amendments and arguments make clear that the 'asserted advance' of the '949 patent, *Intel*, 21 F.4th at 811, was not the 'hardware buffer.'" Pet. Remand Br. 12. Patent Owner responds that "[t]he '949 claims were allowed for reasons unrelated to the Examiner's mistaken belief that Svensson[PCT]'s ISA is a hardware buffer" and, therefore, the applicants did not acquiesce to the Examiner's position that Svensson PCT's ISA is a "hardware buffer." PO Remand Br. 15 (citing *Salazar v. Procter & Gamble Co.*, 414 F.3d 1342, 1344–48 (Fed. Cir. 2005);

3M Innovative Props. Co. v. Avery Dennison Corp., 350 F.3d 1365, 1373–74 (Fed. Cir. 2003)). The cases cited by Patent Owner deal with an applicant's silence in the face of statements by an examiner. *See Salazar*, 414 F.3d at 1347–48 (determining that an "applicant's silence to the examiner's remarks in the Examiner's Statements of Reasons for Allowance" "do[es] not amount to a clear disavowal of claim scope by the applicant"); *3M Innovative*, 350 F.3d at 1373–74 ("An applicant's silence in response to an examiner's characterization of a claim does not reflect the applicant's clear and unmistakable acquiescence to that characterization if the claim is eventually allowed on grounds unrelated to the examiner's unrebutted characterization."). We agree that an applicant's silence should not necessarily be held against it, but for the reasons explained below, we find that the applicant's affirmative statements in this case inform the meaning of the term "hardware buffer."

The claim construction discussion in *University of Massachusetts v. L'Oréal S.A.*, 36 F.4th 1374 (Fed. Cir. 2022), is particularly instructive. The claim at issue in that case recited a "method comprising topically applying to the skin a composition comprising a concentration of adenosine in an amount effective to enhance the condition of the skin without increasing dermal cell proliferation, wherein the adenosine concentration applied to the dermal cells is 10^{-4} M to 10^{-7} M." *L'Oréal*, 36 F.4th at 1377. The disputed issue concerned the language of the "wherein" clause, specifically whether the recited concentration refers to "the concentration as it is applied to the dermal cells," which are in the dermis below the skin surface (epidermis), or "the concentration of adenosine in the composition that is applied to the epidermis." *Id.* at 1378. The court "conclude[d] that the relevant claim

language, especially when viewed in the context of the whole claim, is not plain on its face." *Id.* at 1379. The court considered the specification and determined that it provided "significant evidence that the wherein clause is best read to refer to the concentration of adenosine in the composition applied to the surface of the skin." *Id.* at 1381. But, "[e]ven more strongly than the specification, the prosecution histories of the patents in question resolve the ambiguity surrounding the meaning of 'the adenosine concentration applied to the dermal cells' in the wherein clause." *Id.* at 1382.

The court began its prosecution history analysis by noting that the predecessor to the claim at issue was identical except that it did not recite the "wherein" clause. *L'Oréal*, 36 F.4th at 1382. A dependent claim recited "wherein the adenosine concentration is 10^{-4} M to 10^{-6} M" but did not mention application to the dermal cells. *Id.* The Federal Circuit noted that, "to overcome a prior-art rejection, applicants imported the dependent claim's wherein clause into the independent claim, but altered its wording so that it now read 'wherein the adenosine concentration applied to the dermal cells is 10^{-4} M to 10^{-6} M.'" *Id.* "Importantly," according to the court, "in describing the amendment to the examiner in the office-action response, applicants implied that this new wording changed nothing about where the concentration was to be measured." *Id.* The Federal Circuit then noted various arguments from the applicants that suggested that there was no change in scope from the previously-recited dependent claim, which did not recite "applied to the dermal cells." *Id.* The court then held "that this prosecution history requires that the wherein clause's reference to the recited concentrations being 'applied to the dermal cells' be read as referring to

concentrations of the composition applied to the skin's surface" because "[t]he amendments and comments clearly convey that [the applicant] was continuing the pre-amendment reliance on the concentration in the composition before application to the skin, rather than introducing a materially different, unexplained notion of concentration, no longer assessed before contact with the object of application." *Id.* at 1382–83. In so concluding, the Federal Circuit stated that, "[a]s an initial, legal matter, we reiterate that this is not a case where the prosecution history must meet the standard of clear and unmistakable disclaimer for overcoming an otherwise-plain meaning" because "the meaning of the relevant claim language is not plain." *Id.* at 1383.

We have a similar situation in the present proceeding. The meaning of the term "hardware buffer" is not plain. *See Intel*, 21 F.4th at 809 ("[I]t is clear from the claim language that the claim term at issue has meaning, but it is unclear what that meaning is. There is no definition to be found in the intrinsic evidence."; "We do not discern, and no party has suggested, that 'hardware buffer' has a clear, undisputed meaning in either ordinary English or in relevant technical parlance."). Thus, "this is not a case where the prosecution history must meet the standard of clear and unmistakable disclaimer for overcoming an otherwise-plain meaning." *L'Oréal*, 36 F.4th at 1383. Here, the Federal Circuit directed us to analyze "what the intrinsic evidence shows the asserted advance to be and how, concretely, the 'hardware buffer' relates to that advance." *Intel*, 21 F.4th at 811. We agree with Petitioner that the prosecution history shows that the "asserted advance" was not the use of a hardware buffer. *See* Pet. Remand Br. 12.

The applicants' arguments in prosecution rely on different features to distinguish the asserted prior art (Svensson PCT). *See* Ex. 1005 at 8–9. For example, the applicants argued that "receiving the image header and each data segment, as recited in claim 1, is patentably distinguishable from receiving the data and the associated header, as disclosed in Svensson [PCT]," and that "loading *each data segment directly* from the hardware buffer to the system memory, as recited in claim 1, is patentably distinguishable from concatenating the data blocks and headers in the intermediate storage area and then transferring the concatenated data to the memory, as recited in Svensson [PCT]." Ex. 1005 at 8–9. These statements suggest that features other than the "hardware buffer" were the key distinctions.

But even more probative of the "asserted advance," *Intel*, 21 F.4th at 811, than the applicants' statements highlighted above is the following passage from the same office action response:

> Applicants have demonstrated above that Svensson [PCT] fails to teach or suggest various elements recited in claim 1. Therefore, claim 1 is believed to be allowable over the cited reference. Furthermore, the *independent claims 11, 17, 19, 21, and 23 also recite separately receiving the image header and each data segment and scatter loading each received data segment directly from the hardware buffer to the system memory*. Therefore, the other rejected independent claims are also believed to be allowable. Finally, dependent claims 2, 3, 5-10, 12-16, 18, 20, 22, and 24 are allowable at least by virtue of their dependence on an allowable base claim, in addition to reasons related to their own recitations.

Ex. 1005 at 9 (emphases added). Then-pending independent claims 11, 17, 19, 21, and 23 correspond, respectively, to issued independent claims 10, 16, 18, 20, and 22. *See* Ex. 1005 at 2–7 (amendments to claims in response to

office action). The applicants' representations in this passage are significant because none of the other independent claims recited a "hardware buffer." This is not a situation in which the applicants were silent, as Patent Owner argues. *See* PO Remand Br. 15. Rather, the applicants affirmatively represented to the Office the reasons they thought the claims were allowable, and the fact that five of the six pending independent claims did not recite "hardware buffer" is strong evidence that other features, such as the separate receipt of image header and data segments, constituted the "asserted advance."

In *L'Oréal*, the "applicants implied that this new wording [in the amendment] changed nothing about where the concentration was to be measured" and distinguished prior art references "with no suggestion of a difference in location of concentration measurement," i.e., whether on or below the skin surface. 36 F.4th at 1382. Similarly, the applicants' arguments in the prosecution of the '949 patent strongly suggest that the presence or absence of the term "hardware buffer" was not a distinguishing feature over the asserted prior art in any relevant respect.

This discussion takes us back to the Federal Circuit's instruction: "What is needed in this case is a more substance-focused analysis . . . of what the intrinsic evidence shows the asserted advance to be and how, concretely, the 'hardware buffer' relates to that advance." *Intel*, 21 F.4th at 811. The applicants' representations to the Office during prosecution, "on which skilled artisans are entitled to rely," *L'Oréal*, 36 F.4th at 1383, show that the "asserted advance" relates not to the "hardware buffer" recited in only one of the independent claims but, rather, to other features, such as receiving the image header and the data segments separately and loading the

data based on the image header, which are features reflected in amendments made to all then-pending independent claims and also reflected in the applicants' arguments. Ex. 1005 at 2–7 (amendments to claims), 8 ("Applicants submit that separately receiving the image header and each data segment, as recited in claim 1, is patentably distinguishable from receiving the data and the associated header, as disclosed in Svensson [PCT]."), 9 ("In contrast [to Svensson PCT], claim 1 recites that each data segment is scatter loaded based at least in part on the loaded image header.").

Thus, the applicants' representations to the Office tend to show that there is no relevant distinction between the ISA in Svensson PCT, which is the same as Svensson and Bauer's ISA, and the recited "hardware buffer," which further confirms that the "hardware buffer" is not limited by any temporal restriction on its creation, to the extent the ISA is temporary.

iv. Additional Considerations

The Federal Circuit additionally "note[d] two matters that would benefit from attention on remand." *Intel*, 21 F.4th at 812. We address each in turn.

1. Inventor Testimony

First, the court directed us to district court testimony from a named inventor (Steve Haehnichen) of the '949 patent regarding "the difficulty of 'design[ing] [the] system so the hardware could place the data right where it needed to be,'" which the court said "may bear on the remand inquiry." *Id.* (alterations in original) (citing Ex. 2004 at 222). Patent Owner argues that "Mr. Haehnichen's testimony reflects the fact that the permanent hardware buffer of the '949 patent enables a more efficient, direct transfer that

outperforms less efficient techniques that involve 'copying things around memory'" and that "the permanent, dedicated nature of the hardware buffer enables it to be integrated into an efficient direct data transfer mechanism, such as the USB controller described in the '949 patent." PO Remand Br. 8–9 (citing Ex. 2003 at 213:22–214:12, 215:20–217:2, 221:25–222:10; Ex. 2015 ¶¶ 35–40). For the reasons discussed below, we disagree with Patent Owner.

Mr. Haehnichen testified as an employee for Patent Owner Qualcomm regarding work he did that related to the '949 patent. Ex. 2003 at 201, 203. He explained that Apple, Qualcomm's customer, wanted a phone with no flash memory on the modem, which provided certain benefits such as cost and space savings. Ex. 2003 at 203–04. Mr. Haehnichen further explained that the flash memory is what stores applications when the phone is turned off and that the '949 patent was directed to loading a program from the host to a modem that did not have its own flash memory. Ex. 2003 at 203. He testified that Apple requested a modem boot time of one second, which Qualcomm had never done. Ex. 2003 at 214. This boot time would require transferring a 20 to 25 megabyte (MB) image from the host to the modem, which would require at least a 25 MB per second transfer rate. Ex. 2003 at 214. Mr. Haehnichen testified that the protocol at the time, which was named Gobi, would not meet these requirements, so Qualcomm would have to break compatibility with Gobi and develop something new, which he called Sahara. Ex. 2003 at 216. When asked how Sahara was able to meet Apple's speed requirements, Mr. Haehnichen explained that "the single biggest thing we did to speed it up was to avoid copying things around memory in the modem." Ex. 2003 at 216:16–21.

Mr. Haehnichen explained this difference in more detail as follows:

> So if you look at Gobi, it reads things into memory, figures out where they are, figures out what they are, where they need to go, and it copies them there; whereas, in Sahara, we would receive the data directly from USB and the only place it would go is right where it needs to be. So we compared to like a nonstop flight. Instead of a flight with a stop in the middle and a long layover, you just put it right where it needs to be at the end.

Ex. 2004 at 222:3–10. When asked "How were you able to design your system so the hardware could place the data right where it needed to be," Mr. Haehnichen testified as follows:

> That was pretty hard. We had to bring in the experts on that piece of hardware, which is the USB receiver, and we had to tell them what we were trying to do, and they told us how the hardware works and what the constraints of that hardware and how we would have to organize and set it up so that we could receive it from the very first byte into the address it needs to be.

Ex. 2004 at 222:13–19. Mr. Haehnichen identified the USB experts as two of his co-inventors on the '949 patent. When asked "How were you able to program that USB hardware to put the data directly into its final destination of memory," Mr. Haehnichen provided the following explanation:

> So to get everything set – so before the bytes arrive, they already know where to go. We have to first get the headers of each image, and the headers told us the destination, the size, characteristics of the data so that we could program the USB hardware data directly and have it allowed to receive right away.

Ex. 2004 at 223:2–7.

This testimony supports that any "asserted advance" had more to do with receiving the image header before the data than it had to do with the use of a "hardware buffer" or the particular characteristics of such a buffer. The

following exchange during cross-examination of Mr. Haehnichen underscores this point:

> Q The '949 patent proposes a solution to what it's describing here to more efficiently load software images from the application to the modem processor, correct?
>
> A I agree.
>
> Q And that solution was to transfer the image header first separate from the rest of the image, correct?
>
> A That's necessary, yes.
>
> Q Sir, let's be very, very precise here. Qualcomm's solution was to transfer the image header first separately from the rest of the image, correct?
>
> A That enables the rest of the solution, yes.
>
> Q Well, sir, let's look at your deposition. I just want to make sure we have this 100 percent right. And if you go to your deposition at page 61, lines 14 to 19. "Question: And – And my question is with respect to the first sentence in particular. Qualcomm's solution was to transfer the image header first separately from the rest of the image, correct?["] "Answer: That's correct."
>
> Were you asked that question and did you give that answer?
>
> A I did. The solution was to transfer the image header first.
>
> Q Separately, right?
>
> A That's right.

Ex. 2004 at 243:9–244:7.

Mr. Haehnichen's testimony further shows the importance of receiving the image header separately from the data segments, which is consistent with the applicants' arguments during prosecution of the '949 patent, which we discuss in the preceding section. This is not to say that separately receiving the image header and data segments was the only

difference highlighted in Mr. Haehnichen's testimony. During his deposition in the district court litigation, Mr. Haehnichen was asked "Can you describe for me the specific differences between the Gobi protocol and what you claim to be the alleged invention of the '949 patent?" Ex. 2004 at 246:19–21. The attorney cross-examining Mr. Haehnichen at trial read into the trial record Mr. Haehnichen's complete answer to this question from the deposition, which is as follows:

> I can. I guess the single biggest difference between how Gobi loads images from the host and how '949 loads images from the host are that the flashless boot – we'll call it flashless boot – described in '949, it receives the headers first from the host. The target or modem then looks at the headers to know where[] it's going to be placing images in memory. It configures the USB hardware to transfer directly into that memory, and then it requests those images from the host. So in one pass, the host pushes data down the USB pipe directly into the system memory of the target where it stays there until execution.
>
> Gobi, in contrast, loads the images into a series of ping-pong buffers or holding buffers where it studies the headers and then decides where to place the data, and so the data is copied from modem system memory to modem system memory into its final destination. In our study, we found that the steps of the process was adding time and would hold us back from getting any higher than roughly half the required data rate. So those are the key differences between the two.
>
> There are other differences. The actual protocol commands are different and some of the other, you know, set up and process is different, but I think main reason we were able to get the performance we got was because of copying it directly into system memory.

Ex. 2004 at 248:9–249:25.[10] This testimony again highlights the separate receipt of headers and data, which is recited in the claims of the '949 patent, as well as other features recited in the claims, such as loading directly to system memory and not copying between system memory locations. We find that this testimony does not in itself provide a clarifying definition for "hardware buffer." Mr. Haehnichen mentions "configur[ing] the USB hardware to transfer directly into that memory" (Ex. 2004 at 249:3–5), but, to the extent configuring USB hardware in particular represented some advance over the art, that advance is not reflected in the claims of the '949 patent, which do not specify a particular communication interface between the primary and secondary processor. Indeed, the '949 patent mentions several possible inter-processor communication buses, including "a[n] HSIC bus (USB-based High Speed Inter-Chip), an HSI bus (MIPI High Speed Synchronous Interface), a SDIO bus (Secure Digital I/O interface), a UART bus (Universal Asynchronous Receiver/Transmitter), an SPI bus (Serial Peripheral Interface), [and] an I2C bus (Inter-Integrated Circuit)." Ex. 1001, 5:35–43.

In sum, we find that Mr. Haehnichen's testimony confirms that the separate receipt of image headers and data was a key advance over Qualcomm's previous bootloader techniques, which is consistent with the applicants' representations to the Office during prosecution of the '949 patent, as discussed in the preceding section.

[10] The examining attorney read this into the record in segments and paused to confirm with the witness that this was his testimony. We have reproduced it as one answer that Mr. Haehnichen gave at his deposition.

2. Software/Hardware Distinction

Second, the Federal Circuit noted that its "cases sometimes speak of hardware and software implementations of computer functions, reflecting a distinction that appears in usage in the field." *Intel*, 21 F.4th at 812 (citing *Tomita Techs. USA, LLC v. Nintendo Co.*, 681 F. App'x 967, 972 (Fed. Cir. 2017); *Texas Digital Systems, Inc. v. Telegenix, Inc.*, 308 F.3d 1193, 1214–16 (Fed. Cir. 2002); *Nazomi Commc'ns, Inc. v. ARM Holdings, PLC*, 266 F. App'x 935, 936 (Fed. Cir. 2008)). The court stated that "[n]either the Board nor the parties have explained what if any bearing that usage may have on a proper understanding of 'hardware buffer' in claims 1–9 and 12 here." *Id.*

Citing the testimony of its declarant, Dr. Rinard, Patent Owner argues that the '949 patent distinguishes its "hardware buffer" from prior art buffers that are allocated under software control. PO Remand Br. 4–5 (citing Ex. 2015 ¶¶ 6–15, 24–26). Dr. Rinard refers to the '949 patent's discussion of allocating temporary buffers in the prior art and explains that this is inefficient because it results in extra memory copy operations. Ex. 2015 ¶¶ 23–24 (citing Ex. 1001, 2:17–22, 2:25–26, 2:29–54, 7:20–30). Dr. Rinard explains that "[t]he prior-art systems' use of system memory is also inefficient because it is general-purpose memory that is not specifically configured for the task of buffering data in a multi-processor system." Ex. 2015 ¶ 25.

Although Patent Owner's arguments and Dr. Rinard's testimony attempt to show performance gains by having a buffer that is permanent as opposed to allocated by software, the evidence discussed above shows that the "asserted advance" and associated performance gains related to the

separate receipt of image header information before data, rather than to the temporal nature of the buffer. Thus, whether the buffer is allocated by software does not appear to be the distinguishing feature over the art presented here.

v. Determination of Meaning

In light of the evidence discussed above, we turn back to the parties' dispute. Petitioner argues that the broadest reasonable interpretation of the term "hardware buffer" is "memory that is physically separate from the memory into which the software image is loaded for execution." Pet. Remand Br. 5. Patent Owner counters that the proper interpretation of this term is "a permanent, dedicated buffer that is distinct from system memory." PO Remand Br. 3. For reasons explained below in our analysis of Petitioner's prior art challenges, the relevant dispute is whether a "hardware buffer" must be "permanent" as in Patent Owner's proposed construction. Based on the intrinsic and extrinsic evidence discussed above, we conclude that the term "hardware buffer" is not limited to a "permanent" buffer. Stated in the affirmative, we conclude that the term "hardware buffer" encompasses within its "range of reasonable meanings," *Intel*, 21 F.4th at 809, a "memory that is physically separate from the memory into which the software image is loaded for execution," per Petitioner's proposed construction (Pet. Remand Br. 5), and a "dedicated buffer that is distinct from system memory" (PO Remand Br. 3), per Patent Owner's proposed construction without the temporal limitation that it be "permanent."

As discussed below, we find that Svensson's ISA meets the constructions set forth above, and we find that the asserted combination

addresses the shortcomings of the use of prior art buffers as described in the '949 patent.

2. Means-Plus-Function Limitations

a. Pre-Appeal Final Written Decision

Independent claim 16 is directed to an apparatus that comprises a combination of elements recited in means-plus-function format: "means for receiving . . . an image header," "means for processing," "means for receiving . . . each data segment," and "means for scatter loading." There is no dispute that these elements are governed by 35 U.S.C. § 112, ¶ 6. *See Intel*, 21 F.4th at 812 ("[T]here is no dispute that claim 16 (and hence dependent claim 17) contains terms that are in means-plus-function format governed by 35 U.S.C. § 112(f).").

Under our Rules, for a means-plus-function limitation, a petition "must identify the specific portions of the specification that describe the structure, material, or acts corresponding to each claimed function." 37 C.F.R. § 42.104(b)(3). In its Petition challenging claim 16, Petitioner set forth proposed constructions for the means-plus-function limitations of claim 16, as shown in the table below. IPR2018-01335 Paper 3 ("1335 Pet.") at 17–22.

Limitation	Petitioner's Proposed Function and Structure
means for receiving at a secondary processor, from a primary processor via an inter-chip communication bus, an image header for an executable software image for the secondary processor that is stored in memory coupled to the primary processor, the executable software image comprising the image header and at least one data segment, the image header and each data segment being received separately	**Function:** receiving at a secondary processor, from a primary processor via an inter-chip communication bus, an image header for an executable software image for the secondary processor that is stored in memory coupled to the primary processor **Structure:** secondary processor (*e.g.*, 110, 210, 302) connected to a primary processor (*e.g.*, 104, 204, 301) via an inter-chip communication bus (*e.g.*, 134, 234, 310) for a USB-based High Speed Inter-Chip (HSIC) bus, a MIPI High Speed Synchronous Interface (HSI) bus, a Secure Digital I/O Interface (SDIO) bus, a Universal Asynchronous Receiver/Transmitter (UART) bus, a Serial Peripheral Interface (SPI) bus, or an Inter-Integrated Circuit (I2C) bus. 1335 Pet. 18–19 (citing Ex. 1001, 5:35–43, Fig. 3; Ex. 1007, 17–18; Ex. 1020[11] ¶ 81).

[11] This exhibit was originally filed as Exhibit 1102 in IPR2018-01335. After consolidation, it was filed as Exhibit 1020. *See* Paper 15 (Joint Identification of Exhibits).

Limitation	Petitioner's Proposed Function and Structure
means for processing, by the secondary processor, the image header to determine at least one location within system memory to which the secondary processor is coupled to store each data segment	**Function:** processing, by the secondary processor, the image header to determine at least one location within system memory to which the secondary processor is coupled to store each data segment **Structure:** a modem processor coupled to a system memory. 1335 Pet. 19–20 (citing Ex. 1001, 3:9–12, 4:58–5:43, 5:59–6:39, 7:60–10:44, 8:50–56, 9:27–41, Figs. 1–3; Ex. 1008, 4–5; Ex. 1020 ¶ 83).
means for receiving at the secondary processor, from the primary processor via the inter-chip communication bus, each data segment	**Function:** receiving at the secondary processor, from the primary processor via the inter-chip communication bus, each data segment **Structure:** a secondary processor (*e.g.*, 110, 210, 302) connected to a primary processor (*e.g.*, 104, 204, 301) via an inter-chip communication bus (*e.g.*, 134, 234, 310) for a USB-based High Speed Inter-Chip (HSIC) bus, a MIPI High Speed Synchronous Interface (HSI) bus, a Secure Digital I/O Interface (SDIO) bus, a Universal Asynchronous Receiver/Transmitter (UART) bus, a Serial Peripheral Interface (SPI) bus, or an Inter-Integrated Circuit (I2C) bus. 1335 Pet. 21 (citing Ex. 1001, 5:35–43, Fig. 3; Ex. 1007, 19; Ex. 1020 ¶ 85).

Limitation	Petitioner's Proposed Function and Structure
means for scatter loading, by the secondary processor, each data segment directly to the determined at least one location within the system memory, and each data segment being scatter loaded based at least in part on the processed image header	**Function:** scatter loading, by the secondary processor, each data segment directly to the determined at least one location within the system memory, and each data segment being scatter loaded based at least in part on the processed image header **Structure:** a modem processor coupled to a system memory. 1335 Pet. 22 (citing Ex. 1001, at code (57), 1:24–33, 4:10–15, 4:58–5:43, 5:59–6:39, 7:60–10:44, 8:21–30, 8:62–67, 9:3–8, 9:16–56, 10:13–18, 10:27–32, Figs. 1–3; Ex. 1008, 6; Ex. 1020 ¶ 87).

In the 1335 Decision on Institution, we agreed with Petitioner that the identified limitations of claim 16 are means-plus-function limitations subject to 35 U.S.C. § 112, ¶ 6, and we agreed with Petitioner's identification of the claimed functions. 1335 Dec. on Inst. 13. We stated, however, that "we have questions as to the sufficiency of Petitioner's identified structures," and we discussed the "means for processing" and "means for scatter loading" limitations. 1335 Dec. on Inst. 13–15. Although we identified some potential issues for the trial, we stated that "we need not construe expressly these means-plus-function limitations" to determine whether to institute trial because "Petitioner has met the threshold for institution as to claims 10–15." 1335 Dec. on Inst. 15.

In its Response, Patent Owner agreed with Petitioner's proposed constructions for the "means for processing" and "means for scatter loading" limitations. PO Resp. 17–21. Patent Owner did not address Petitioner's constructions for the "means for receiving . . . an image header" and "means for receiving . . . each data segment." *See* PO Resp. 17–21.

In its Reply during the trial, Petitioner stated that, "[u]pon consideration of the Board's articulated concerns, Petitioner agrees that the '949 specification fails to disclose sufficient structure to perform the recited functions." Pet. Reply 14.

In the Final Written Decision, we stated the following:

> [I]t is Petitioner's burden to identify corresponding structure. Because Petitioner asserts during the trial that the Specification of the '949 patent fails to disclose sufficient corresponding structure for the "means for processing . . . the image header" and "means for scatter loading" limitations (Pet. Reply 14), Petitioner has not met this burden. Furthermore, in the absence of the requisite showing by Petitioner of sufficient corresponding structure for the means-plus-function limitations, we need not further address the construction of these claim terms to determine whether Petitioner has met its burden to prove, by a preponderance of the evidence, unpatentability of independent claim 16 and dependent claim 17.

Final Dec. 18. We further stated the following:

> Because Petitioner has not met its burden under our Rules to show structure corresponding to the claimed function to which we can compare the prior art's disclosure, we determine Petitioner has not shown, by a preponderance of the evidence, that claim 16 and, by virtue of its dependency, claim 17 are unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, Kim, and Zhao.

Final Dec. 62.

Petitioner appealed our determination as to claims 16 and 17. *See Intel*, 21 F.4th at 812–14.

b. Federal Circuit's Decision

The Federal Circuit held that we erred in our determination that Petitioner had not met its burden to demonstrate unpatentability of claims 16 and 17. *Intel*, 21 F.4th at 812. The court noted that, "[u]nder 35 U.S.C. § 318, as construed in *SAS Institute, Inc. v. Iancu*, —— U.S. ——, 138 S. Ct. 1348, 200 L.Ed.2d 695 (2018), the Board is obligated to 'issue a final written decision with respect to the patentability of' every claim challenged by the petitioner." *Intel*, 21 F.4th at 812 (quoting *SAS Inst.*, 138 S.Ct. at 1353). According to the Federal Circuit, however, "that obligation does not mean that the Board must reach a determination of the patentability of a claim on the presented prior-art grounds if such a determination is rendered impossible because of the indefiniteness of an essential claim limitation." *Id.* at 812–13.

The Federal Circuit remanded the case to us to

> decide one or both of two issues. One is whether it can resolve the prior-art challenge to the patentability of claims 16 and 17 despite the potential indefiniteness of the means-plus-function terms. The other is whether those terms are actually indefinite. We express no view on which issue the Board should consider first. If the Board determines both that there is indefiniteness and that such indefiniteness renders it impossible to adjudicate the prior-art challenge on its merits, then the Board should conclude that it is impossible to reach a decision on the merits of the challenge and so state in its decision.

Intel, 21 F.4th at 814.

c. Remand Proceedings

On remand, "we asked the parties to consider whether the prior art challenges to claims 16 and 17 can be resolved based on the means-plus-function constructions proposed in the Petition." Paper 34 at 3.

For the "means for processing" and "means for scatter loading," Petitioner notes that it "presented the same constructions . . . that Patent Owner had previously advanced in a separate ITC proceeding." Pet. Remand Br. 18 (citing Ex. 1008 (ITC Joint Claim Construction Chart), 4–6); *see* 1335 Pet. 19–22 (noting that the constructions proposed in the Petition for "means for processing" and "means for scatter loading" are the constructions Patent Owner proposed in the ITC).

On remand, Patent Owner notes that "[t]he parties have briefed that issue, and [Patent Owner] rests on its briefs." PO Remand Br. 19. As noted above in § II.B.2.a, Patent Owner agreed with Petitioner's proposed constructions for the "means for processing" and "means for scatter loading," which makes sense because these are the constructions that Patent Owner proposed at the ITC. PO Resp. 17–21; Ex. 1008, 4–6. Patent Owner also did not dispute Petitioner's proposed constructions for the "means for receiving . . . an image header" and "means for receiving . . . each data segment." *See* PO Resp. 17–21.

Thus, in compliance with our Rules, the Petition "identif[ies] the specific portions of the specification that describe the structure, material, or acts corresponding to each claimed function" of the means-plus-function limitations in claim 16. 37 C.F.R. § 42.104(b)(3); 1335 Pet. 17–22. Patent Owner does not dispute Petitioner's proposed constructions for the means-plus-function limitations. *See* PO Resp. 17–21; PO Remand Br. 19. Indeed,

Patent Owner stated that "these terms do not need to be construed in order for the Board to reach its Final Written Decision" because "[n]one of the arguments [Patent Owner] makes . . . to distinguish the prior art requires construction of these limitations." PO Resp. 17. In view of Patent Owner's agreement with the constructions proposed in the Petition for the means-plus-function limitations of claim 16, we apply those constructions in our analysis of the Petitioner's unpatentability contentions. Thus, based on the parties' arguments, we "can resolve the prior-art challenge to the patentability of claims 16 and 17 despite the potential indefiniteness of the means-plus-function terms," and we need not determine "whether those terms are actually indefinite." *See Intel*, 21 F.4th at 814; *see also Realtime Data, LLC v. Iancu*, 912 F.3d 1368, 1375 (Fed. Cir. 2019) ("The Board is required to construe 'only those terms . . . that are in controversy, and only to the extent necessary to resolve the controversy.'" (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng'g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999)).

C. Principles of Law

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

considerations, if in evidence.[12] *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

D. Obviousness over Bauer, Svensson, and Kim (Claims 1–9, 12)

Petitioner contends that claims 1–9 and 12 of the '949 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, and Kim. Pet. 23–71; 1335 Pet. 60–62.

1. Svensson

Svensson describes a multi-processor system in which data are sent from a host processor to a client processor. Ex. 1010, code (57). Figure 1 of Svensson is reproduced below.

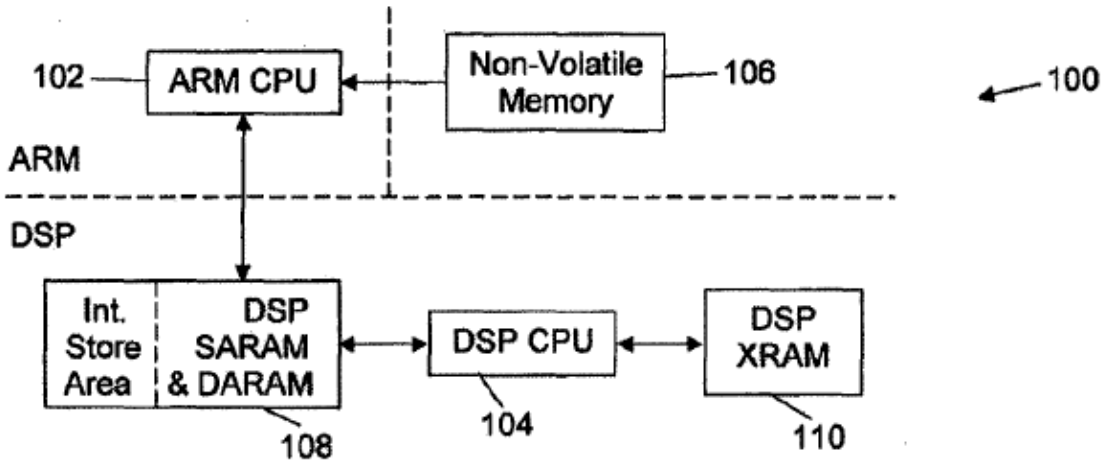

FIG. 1

Figure 1 depicts multi-processor system 100 having host processor 102 and client processor 104. Ex. 1010, 3:49–50. Client processor 104 is the processor for a digital signal processor (DSP) device. Ex. 1010, 3:54–58. As Svensson explains, "[m]ost commercially available DSP devices include

[12] Patent Owner does not present any objective evidence of nonobviousness (i.e., secondary considerations) as to any of the challenged claims.

on-chip memories, and as indicated in FIG. 1, the DSP includes 'internal' single-access RAM (SARAM) and dual-access RAM (DARAM) 108, as well as an 'external' RAM (XRAM) 110." Ex. 1010, 3:64–4:1. Svensson explains that "XRAM 110 is invisible to, i.e., not accessible by, the CPU 102," whereas CPU 102 can access "internal" SARAM and DARAM 108. Ex. 1010, 4:5–8, 4:13–14. DSP processor 104 can access both RAMs 108 and 110. Ex. 1010, 4:7–8.

Because host processor 102 cannot access XRAM 110, Svensson discloses a technique for sending data from host processor 102 to be stored in XRAM 110. Ex. 1010, Fig. 2, 4:15–6:11, 7:7–8. Svensson's Figure 2 is reproduced below.

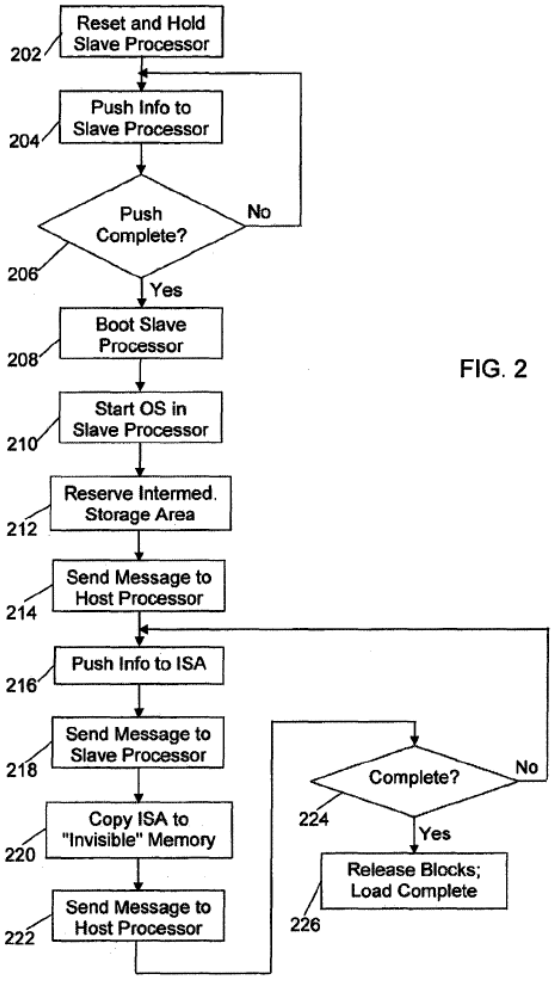

FIG. 2

Figure 2 is a flow chart of Svensson's bootloader operation. Ex. 1010, 3:34, 4:15–19. In step 212, a block of memory in "internal" memory 108 is reserved as an intermediate storage area (ISA) for data that are being sent from the host to the invisible memory of the client processor. Ex. 1010, 5:21–28. After the host transfers data to the ISA (step 216), the host tells the client the ISA has been loaded and indicates whether more data are coming (step 218). Ex. 1010, 5:53–63. The client then copies the data from the ISA to its "invisible" memory (step 220) and responds to the host when copying is finished (step 222). Ex. 1010, 5:63–6:3. "If there is more code and/or data to load (Step 224), this cycle of copying and messaging (Steps 216–224) can be repeated as many times as required." Ex. 1010, 6:4–6.

2. Bauer

Bauer discloses the file format depicted in Figures 1A, 1B, and 1C, which are reproduced below.

FIG. 1A

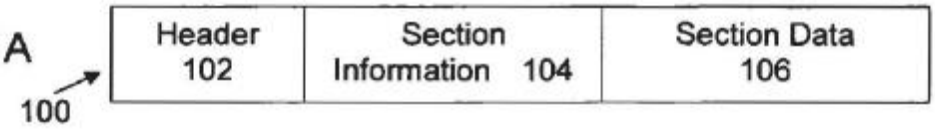

FIG. 1B

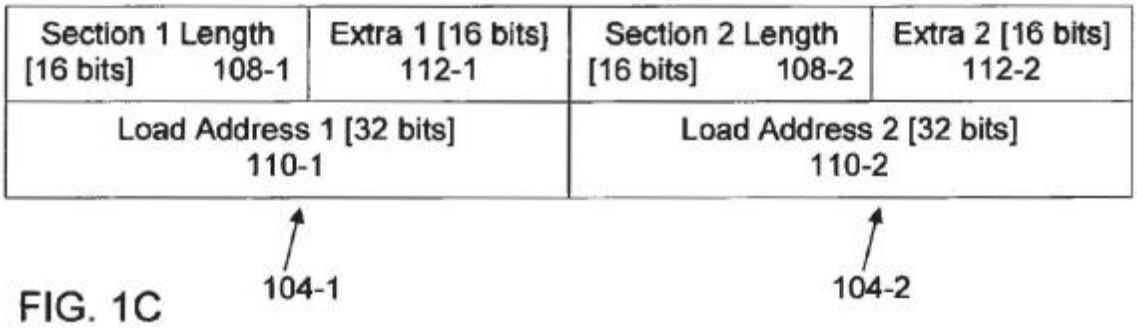

FIG. 1C

Figure 1A shows the format for a data image, Figure 1B shows the header of the data image, and Figure 1C shows the section information of the data image. Ex. 1009 ¶¶ 21–23. As shown in Figure 1A, binary data image 100 has header 102, section information 104, and section data 106. Ex. 1009 ¶ 32. Each section of data in section data 106 has a section information entry in section information 104, two of which are depicted in Figure 1C as entries 104-1 and 104-2. Ex. 1009 ¶ 34. Each section information entry indicates the length (108) and load address (110) for its respective section data. Ex. 1009 ¶ 34. Additional information about a section may be included in extra information element 112. Ex. 1009 ¶ 34.

According to Bauer, "[h]aving all section information entries 104 collected together in the image 100 advantageously simplifies system navigation through the image, and having all section data arranged in a

sequence makes it possible to optimize loading of the sections." Ex. 1009 ¶ 38. Bauer explains that "[t]here are many possible applications of this format and its individually coded sections," including "[o]bject code and data . . . with a program loader reading the stored information and processing stored sections accordingly." Ex. 1009 ¶ 31. "One example of such a program loader is described in U.S. patent application Ser. No. 11/040,798 filed on Jan. 22, 2005, by M. Svensson et al. for 'Operating-System-Friendly Bootloader'." Ex. 1009 ¶ 31. This is the application that issued as Svensson. Svensson's Figure 1 depicts the same multi-processor system as Bauer's Figure 2, which Bauer says "can advantageously use a binary image 100 having the format depicted in FIGS. 1A, 1B, 1C." Ex. 1009 ¶ 35; *compare* Ex. 1010, Fig. 1, *with* Ex. 1009, Fig. 2.

3. Kim

Kim discloses a system in which a system startup loader in a system management processor provides program blocks to multiple other processors in a system. Ex. 1012, 4:8–21, Fig. 1. Figure 3 of Kim is reproduced below.

FIG. 3

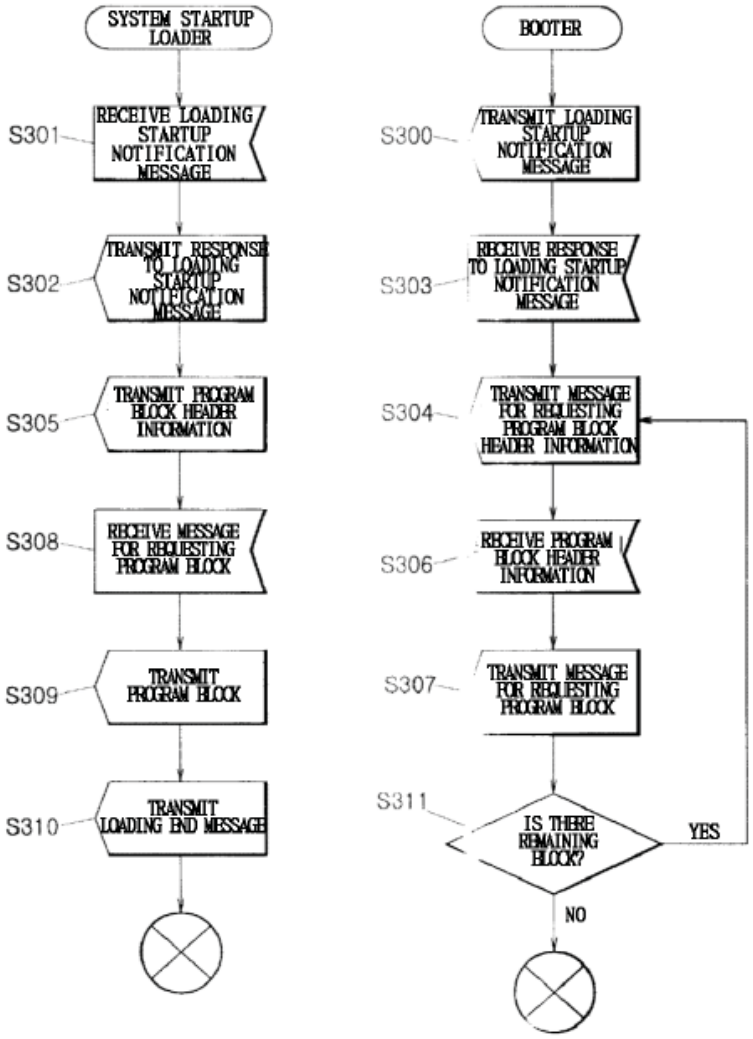

Figure 3 is a flowchart showing a procedure for loading program blocks from the system startup loader to other processors in the system. Ex. 1012, 5:9–11. In step S304, the booter in a processor requests program block header information, which the system startup loader provides in step S305. Ex. 1012, 5:18–21. When the header is received, the booter requests a program block in step S307, which the system startup loader provides in step S309. Ex. 1012, 5:21–24. If there are more blocks to be received, the booter returns to step S304. Ex. 1012, 6:2–4.

4. Claim 1

Claim 1 is directed to a "multi-processor system" and recites much of the same subject matter as claims we held unpatentable in the Final Written Decision. On remand, the dispute between the parties is about whether the combination of Bauer and Svensson teaches "a hardware buffer" that is separate from "system memory." First, we address the undisputed subject matter of claim 1 and then turn to the disputed issues.

a) Undisputed and Previously Addressed Contentions

For the recited components of claim 1, Petitioner identifies particular components in Bauer's Figure 2 (reproduced below) and in Svensson's Figure 1, both of which depict the same multi-processor system, as noted above.

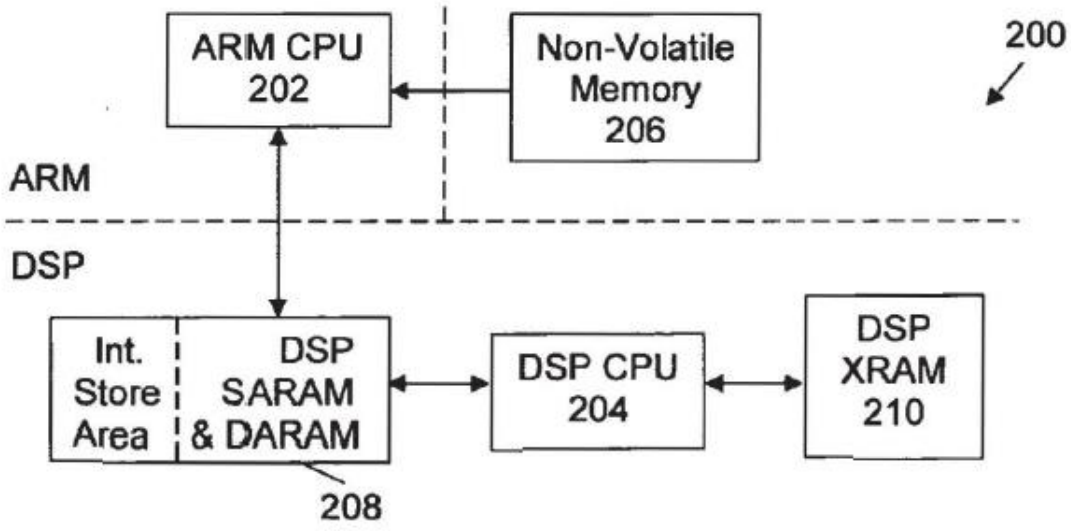

FIG. 2

Figure 2 of Bauer depicts multi-processor system 200 having host processor 202 and client processor 204. Ex. 1009 ¶ 35. In Figure 2, host processor 202 is an advanced RISC (reduced instruction set computer)

machine (ARM) central processing unit (CPU), and client processor 204 is a DSP CPU. Ex. 1009 ¶ 35.

Petitioner argues that Bauer's multi-processor system 200 is a "multi-processor system," as recited in the preamble of claim 1, and that Bauer's DSP device and ARM CPU are, respectively, "a secondary processor" and "a primary processor." Pet. 25–26, 30. Petitioner argues that Bauer's disclosure in Figure 2 of the ARM processor coupled to non-volatile memory 206 teaches "a primary processor coupled with a memory," as recited in claim 1. Pet. 50 (citing Ex. 1009 ¶¶ 35–36, Fig. 2). For the claimed recitation of "the memory storing the executable software image for the secondary processor," Petitioner argues that, in Bauer and Svensson, the non-volatile memory coupled to the ARM CPU stores executable software for the DSP device. Pet. 50–51 (citing Ex. 1009 ¶¶ 11, 31, 35–36, Fig. 2; Ex. 1010, 4:9–14, 6:12–15, Figs. 1, 3; Ex. 1002 ¶ 152). For example, Svensson discloses that "[t]he SARAM and DARAM 108 [of the DSP device] can be loaded from the non-volatile memory 106 [of the ARM CPU] by the trivial 'push' method." Ex. 1010, 4:9–10.

For the claimed "an interface communicatively coupling the primary processor and the secondary processor, the executable software image being received by the secondary processor via the interface," Petitioner argues Figure 2 of Bauer and Figure 1 of Svensson show an interface coupling the ARM CPU and the DSP device. Pet. 51–52 (citing Ex. 1009 ¶ 36, Fig. 2; Ex. 1010, 4:3–5, 4:9–14, Fig. 1; Ex. 1002 ¶ 154); *see* Ex. 1009 ¶ 36 ("The arrows in FIG. 2 indicate access paths, e.g., busses and direct memory access (DMA) paths, between the CPUs and the memories . . . .").

We are persuaded by Petitioner's contentions for the subject matter addressed above, which is substantially similar to subject matter recited in claim 10. To the extent Patent Owner raised any arguments with respect to the above subject matter of claim 1, we addressed those arguments in the Final Written Decision for similar subject matter in claim 10. Final Dec. 26–50. The determination of unpatentability as to claim 10 is undisturbed because Patent Owner dropped its challenges concerning claim 10 on appeal. *See Intel*, 21 F.4th at 806 ("The Board ruled that Intel had proved the unpatentability of claims 10, 11, 13–15, and 18–23, but Qualcomm, despite filing a cross-appeal to raise the issue, no longer challenges that ruling."). For the reasons stated in our Final Written Decision as to claim 10, we agree with Petitioner that the combination of Bauer, Svensson, and Kim teaches the above-discussed subject matter of claim 1, and Petitioner's reasons to combine the teachings of the references are persuasive. *See* Final Dec. 26–50.

b) Disputed Issues on Remand

Claim 1 recites that the secondary processor comprises

> system memory and a hardware buffer for receiving an image header and at least one data segment of an executable software image, the image header and each data segment being received separately, and
>
> a scatter loader controller configured: to load the image header; and to scatter load each received data segment based at least in part on the loaded image header, directly from the hardware buffer to the system memory.

These limitations are similar to subject matter recited in claim 10. Claim 10 recites "receiving at a secondary processor . . . an image header for an executable software image . . ., the executable software image comprising

the image header and at least one data segment, the image header and each data segment being received separately." Claim 10 further recites "receiving at the secondary processor . . . each data segment" and "scatter loading, by the secondary processor, each data segment [directly[13]] to the determined at least one location within the system memory, and each data segment being scatter loaded based at least in part on the processed image header." We addressed the parties' arguments with respect to claim 10 and determined that claim 10 is unpatentable as obvious based on the same combination asserted against claim 1, the combination of Bauer, Svensson, and Kim. *See* Final Dec. 26–50.

The main differences in language between claim 1 and claim 10 are claim 1's recitation of a "hardware buffer" in the secondary processor and claims 1's recitation of directly loading from that hardware buffer to system memory. Although claim 10 recites "scatter loading . . . each data segment directly," it does not recite a "hardware buffer." During prosecution, however, the applicants did not distinguish between claims that recited a hardware buffer and those that did not. Ex. 1005 at 8–9; *see also* § II.B.c.iii above (discussing prosecution history of the '949 patent). Rather, the applicants distinguished the pending claims by arguing that Svensson PCT discloses receiving headers with the data, not separately, and that Svensson PCT discloses concatenating data blocks and headers in the ISA and then transferring the concatenated data to memory. Ex. 1005 at 8–9. In our Final Written Decision, we concluded that receiving the image header separately

[13] The issued patent recites "reedy," which appears to be a printing error. The April 30, 2014 claim listing submitted by the applicants during prosecution states "directly."

from each data segment would have been obvious based on the combination of Svensson, Bauer, and Kim, relying on Bauer's disclosure of a file format having header and section information collected in one place for all data sections for efficient memory loading and also relying on Kim's disclosures of separate receipt of header and data. *See* Final Dec. 26–42; *see also* § II.D.2 above (overview of Bauer's file format). We also found that the combination of Bauer and Svensson teaches direct loading. *See* Final Dec. 44–46. In particular, we found that loading via the ISA of Svensson and Bauer is as direct as the loading in the '949 patent because "in both the '949 patent and in Svensson, data are received at some location in the secondary processor first before they are directly loaded." Final Dec. 46. Claim 1 expressly recites that the data are received in a "hardware buffer" in the secondary processor, and we found that this language distinguished claim 1 from Bauer and Svensson's ISA, which we found to be temporary and, therefore, excluded from the scope of a "hardware buffer." Final Dec. 55–56. The Federal Circuit vacated that claim construction and our determination as to claim 1. *Intel*, 21 F.4th at 814.

Thus, the parties' dispute on remand centers on whether the prior art combination teaches a "hardware buffer" and limitations involving the hardware buffer. Referring to Bauer's Figure 2, Petitioner argues that DSP XRAM 210 to the right of DSP CPU 204 is "system memory" and that the intermediate storage area (ISA) in DSP SARAM & DARAM 208 is a "separate hardware buffer." Pet. 27 (citing Ex. 1009 ¶¶ 35–36, Fig. 2; Ex. 1010, 3:54–58, 3:64–4:5, Fig. 1; Ex. 1002 ¶ 110).

Patent Owner argues that the ISA is not a permanent and dedicated buffer distinct from the system memory of the DSP and, therefore, is not a

"hardware buffer" under its proposed construction. PO Remand Br. 16–19. For the reasons explained below, we find that the ISA is a "hardware buffer" because it is a dedicated buffer that is physically separate and distinct from the system memory.

As an initial matter, we agree with Patent Owner that Bauer and Svensson's ISA is a temporary buffer in the sense that is does not exist until it is allocated by the program loader. *See* PO Remand Br. 16–17; *see also* Final Dec. 55–56 (finding that the ISA is a temporary buffer); Ex. 1026 ¶ 41 (Petitioner's declarant Dr. Lin testifying that "Svensson discloses that the intermediate storage area is reserved at boot up of the client [DSP] processor. [Ex. 1010] at Fig. 2 (blocks 208-212)."). Petitioner argues that the buffer is permanent because, once it is allocated, it is not deallocated and remains as a buffer that can be used for loading software to the DSP, thereby "function[ing] solely as a permanent buffer." Pet. Remand Br. 17 (citing Ex. 1026 ¶¶ 41, 44–45). Even if the ISA is never deallocated and continues to exist for the time the DSP is operating, it is still not permanent in the sense of existing absent initial allocation. Because we conclude, however, that permanency is not a requirement of a "hardware buffer," as discussed above in § II.B.1, the ISA's initial allocation does not preclude it from being a "hardware buffer."

Below is a version of Bauer's Figure 2 with annotations by Patent Owner.

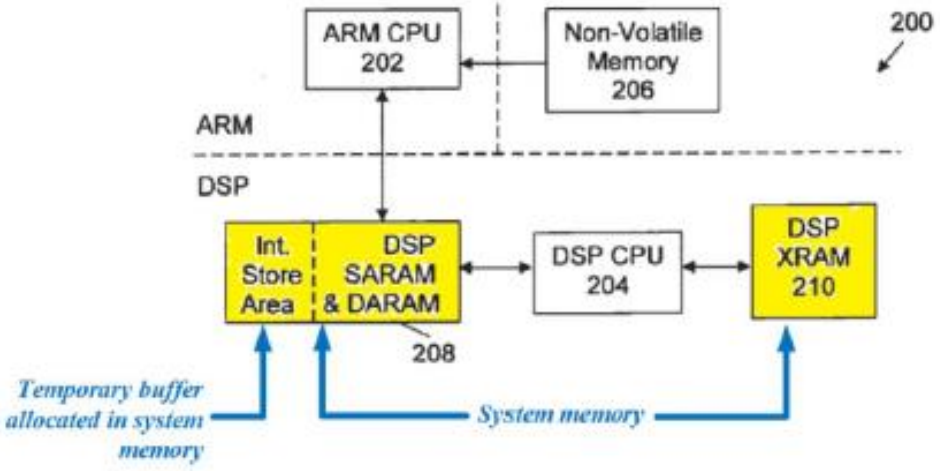

PO Remand Br. 18. In Bauer's Figure 2 above, Patent Owner labels DSP SARAM & DARAM 208 and DSP XRAM 210, on the left and right side of DSP CPU 204, as "System memory," and Patent Owner labels the ISA as "Temporary buffer allocated in system memory." PO Remand Br. 18. As noted above, Svensson's Figure 1 depicts the same multi-processor system as Bauer's Figure 2. *Compare* Ex. 1010, Fig. 1, *with* Ex. 1009, Fig. 2

Patent Owner contends that "the ISA is not distinct from the SARAM & DARAM 108/208, and therefore is not a hardware buffer as properly construed." PO Remand Br. 18–19. Patent Owner's argument, however, is based on grouping memories 208 and 210 together as "system memory." But Bauer and Svensson explain that these two memories are quite different. For example, Bauer describes memory 208 as "internal" to the DSP and memory 210 as "external" DSP memory. Ex. 1009 ¶ 36; *see also* Ex. 1010, 3:64–4:1 (similar disclosure with respect to Svensson's Figure 1). Further, Bauer and Svensson explain that the ARM CPU can access memory 208 but not memory 210, which is only accessible to the DSP CPU. Ex. 1009 ¶ 36; *see also* Ex. 1010, 4:5–8 (similar disclosure with respect to Svensson's

Figure 1). Indeed, Svensson explains, with reference to its Figure 1 (corresponding to Bauer's Figure 2), that "XRAM 110 is invisible to, i.e., not accessible by, the CPU 102 and so boot code cannot be pushed to the XRAM 110." Ex. 1010, 4:11–14. Thus, although memory 208 is memory that can be used by the DSP CPU, it is not system memory in the sense of being the destination for the executable software image being pushed from the ARM CPU. That system memory is external memory 210 in Bauer's Figure 2.

Patent Owner argues that having a temporary buffer (the ISA) in memory 208 results in the "extra memory copy operations" that the '949 patent distinguished in the prior art. PO Remand Br. 4 (citing Ex. 1001, 2:17–22, 2:29–41, 7:16–30; Ex. 2015 ¶ 24). We disagree. As discussed above in § II.B.1, the focus of the '949 patent in this regard is eliminating extra memory copy operations by the secondary processor. For example, the Federal Circuit stated that "the meaning of 'hardware buffer' relates to the ability to move the software image 'directly' to the second processor's system memory and to avoid 'copying data between system memory locations,'" as recited in claim 2. *Intel*, 21 F.4th at 810. Claim 2, however, relates to the configuration of the scatter loader controller, which is a component of the secondary processor as recited in claim 1. The '949 patent also discloses that "[t]he modem processor 110 stores the modem executable image 132 directly into the modem processor RAM (Random Access Memory) 112 to the final destination without copying the data into a temporary buffer in the modem processor RAM 112." Ex. 1001, 5:31–35; *see Intel*, 21 F.4th at 811 (citing this passage). The modem processor is the secondary processor, and, therefore, this passage also shows that the '949

patent's focus is on avoiding extra memory copy operations by the secondary processor. *See also* Ex. 1001, 9:42–43 ("Accordingly, no extra memory copy operations occur in the secondary processor in the above aspect.").

The software loading technique in the combination of Bauer and Svensson has the same number of memory copy operations by the secondary processor as the '949 patent. In the combination of Bauer and Svensson, the ARM CPU (primary processor) pushes the information into the ISA, and the DSP CPU (secondary processor) copies the information to its external memory, which is not accessible to the ARM. *See* Ex. 1010, 5:53–67 ("[T]he host bootloader fills the intermediate storage area with information (code and/or data) to be loaded into the slave's invisible memory (Step 216) . . . . The slave copies the contents of the intermediate storage area to appropriate locations in its slave-private memory (Step 220), thereby implementing its actual loading."), 6:12–15 ("[T]he host fills the intermediate storage area in the memory 108 with code and data that the slave further copies to end destinations in the slave-private memory 110."). Thus, in Svensson's bootloading process, for each piece of data placed in the ISA by the primary processor, the secondary processor performs one transfer to the final destination in system memory. Although the primary processor puts data into the ISA, the secondary processor does not perform multiple copying operations.

Patent Owner also argues that the ISA of Bauer and Svensson is not a "dedicated" buffer. PO Remand Br. 17–19. According to Patent Owner, the "hardware buffer of the '949 patent also enables a more efficient loading process because it is *permanent* and *dedicated* to the single task of loading

data segments directly to system memory." PO Remand Br. 7 (citing Ex. 2015 at ¶¶32–33). Putting aside the "permanent" aspect of Patent Owner's argument, which we determine above is not a requirement of a "hardware buffer," we find that the ISA of Bauer and Svensson is a dedicated buffer. Svensson explains the following:

> The idle process reserves a block of memory in the slave's heap of memory that is located in the memory visible to the host, such as "internal" memory 108 (Step 212). As described in more detail below, this reserved block of memory is used for intermediate storage of information (code and/or data) to be transferred to the slave-private memory, i.e., the memory that is invisible to the host, such as "external" XRAM 110.

Ex. 1010, 5:21–29. Thus, the ISA of Svensson and Bauer is a dedicated buffer because it is a storage area reserved for one purpose, specifically loading data from the primary processor to a memory of the secondary processor not accessible to the primary processor. And the ISA resides in a separate memory from the target destination memory for the transferred data, which underscores the "distinction" between the system memory and the hardware buffer in the combination of Bauer and Svensson. *See Intel*, 21 F.4th at 810 ("[B]ecause claim 1 requires both a 'system memory' and a 'hardware buffer,' there must be some distinction between those two concepts.").

Furthermore, the combination of Bauer, Svensson, and Kim teaches an implementation of an intermediate buffer that is different from the temporary buffer implementations described in the '949 patent. For example, the '949 patent states the following:

> In a system in [w]hich the software image is loaded onto a target "secondary" processor from a first "primary" processor, one way of performing such loading is to *allocate a temporary*

> *buffer into which each packet is received, and each packet would have an associated packet header information along with the payload*. The payload in this case would be the actual image data. From the temporary buffer, some of the processing may be done over the payload, and then the payload would get copied over to the final destination. The temporary buffer would be some place in system memory, such as in internal random-access-memory (RAM) or double data rate (DDR) memory, for example.

Ex. 1001, 2:23–34 (emphasis added). Thus, in this passage, a temporary buffer is used to receive each packet, and each packet has a header with the data payload. This is different from the combination of Bauer, Svensson, and Kim, in which the header is received separately from the data. *See* Final Dec. 30–42.

The '949 patent also describes the use of a temporary buffer in the following passages:

> In one exemplary aspect a direct scatter load technique is disclosed for loading a segmented image from a primary processor's non-volatile memory to a secondary processor's volatile memory. As discussed further below, the direct scatter load technique avoids use of a temporary buffer. For instance, in one aspect, rather than employing a packet-based communication in which the image is communicated via packets that each include a respective header, the raw image data is loaded from the primary processor to the secondary processor. In another aspect, headers are used which include information used to determine the target location information for the data.

Ex. 1001, 4:43–54.

> Accordingly, no extra memory copy operations occur in the secondary processor in the above aspect. Thus, conventional techniques employing a *temporary buffer for the entire image, and the packet header handling, etc.*, are bypassed in favor of a more efficient direct loading process. Thus, the exemplary load process of FIG. 3 does not require the intermediate buffer

> operations traditionally required for loading a software image from a primary processor to a secondary processor. Instead of scatter loading from a *temporary buffer holding the entire image*, the exemplary load process of FIG. 3 allows for direct scatter load the image segments to their respective target destinations directly from the hardware to the system memory. Once the image header is processed, the executable image is directly scatter loaded into target memory, bypassing farther CPU involvement.

Ex. 1001, 9:42–56 (emphasis added).

These passages discuss a prior art temporary buffer used in conjunction with "packets that each include a respective header," which require certain "packet header handling." Such use of a temporary buffer is different from the combination of Bauer, Svensson, and Kim, in which header information is received separately from the data. *See* Final Dec. 30–42. Furthermore, as reproduced above, the '949 patent states that a difference between the prior art and the '949 patent is that, in the latter, "headers are used which include information used to determine the target location information for the data." Ex. 1001, 4:52–54. As we found in the Final Decision, this is precisely what Bauer's file format describes: "Bauer discloses a file format with section information having load addresses for data segments, and Bauer explains that 'an operating system memory manager can load and unload sections of memory according to images in this format.' Ex. 1009 ¶¶ 31, 34." Final Dec. 43; *see* Ex. 1009 Fig. 1C (showing load addresses in the section information).

Another '949 patent passage mentioning a temporary buffer states the following: "The modem processor 110 stores the modem executable image 132 directly into the modem processor RAM (Random Access Memory) 112 to the final destination without copying the data into a temporary buffer in

the modem processor RAM 112." Ex. 1001, 5:31–35. Thus, this passage describes that the secondary processor does not copy data into a temporary buffer, but this is not a distinction over the combination of Bauer, Svensson, and Kim. As discussed above, Svensson discloses that the primary processor pushes data to the ISA, not that the secondary processor copies data to the ISA. *See* Ex. 1010, 5:53–67, 6:12–15.

For all of the reasons discussed above, we find that the ISA of Bauer and Svensson is a "hardware buffer." We also find that the external memory that is the target destination for the software that is pushed to the ISA is a "system memory" of the secondary processor. Thus, we find that the combination of Bauer, Svensson, and Kim teaches a "memory that is physically separate from the memory into which the software image is loaded for execution," per Petitioner's proposed construction (Pet. Remand Br. 5), and also teaches a "dedicated buffer that is distinct from system memory" (PO Remand Br. 3), per Patent Owner's proposed construction without the temporal limitation that it be "permanent."

Thus, for the reasons discussed above and in the Final Decision and based on Petitioner's persuasive contentions and evidence, we conclude that the combination of Bauer, Svensson, and Kim renders obvious a secondary processor that comprises "system memory and a hardware buffer for receiving an image header and at least one data segment of an executable software image, the image header and each data segment being received separately." *See* Final Dec. 26–42 (findings and conclusions as to claim 10's step of "receiving at a secondary processor, from a primary processor via an inter-chip communication bus, an image header for an executable software image for the secondary processor that is stored in memory coupled

to the primary processor, the executable software image comprising the image header and at least one data segment, the image header and each data segment being received separately").

Furthermore, we conclude that the combination of Bauer, Svensson, and Kim renders obvious "a scatter loader controller configured: to load the image header; and to scatter load each received data segment based at least in part on the loaded image header, directly from the hardware buffer to the system memory." This conclusion is based on our finding that Svensson and Bauer's ISA is a "hardware buffer" from which the secondary processor loads data to its XRAM (system memory) based on Bauer's file format with load addresses, as discussed above, and also based on our findings and conclusions for claim 10's step of "scatter loading, by the secondary processor, each data segment [directly] to the determined at least one location within the system memory, and each data segment being scatter loaded based at least in part on the processed image header," which does not require the "hardware buffer" recited in claim 1. *See* Final Dec. 44–50.

c) Conclusion for Claim 1

Having considered the full record developed during trial, on appeal to the Federal Circuit, and on remand, we are persuaded by Petitioner's contentions and evidence, and we determine that Petitioner has proven by a preponderance of the evidence that claim 1 of the '949 patent is unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, and Kim.

5. Dependent Claims 2–9

Petitioner contends that claims 2–9, which depend from claim 1, are unpatentable under 35 U.S.C. § 103(a) as obvious over the combined

teachings of Bauer, Svensson, and Kim. Pet. 53–71. Patent Owner presents additional arguments only against claim 2. PO Resp. 75–77. For the reasons explained below, we determine that claims 2–9 are unpatentable.

Claim 2 recites, "The multi-processor system of claim 1 in which the scatter loader controller is configured to load the executable software image directly from the hardware buffer to the system memory of the secondary processor without copying data between system memory locations on the secondary processor."

Referring to its contentions for claim 1, Petitioner argues that, in Bauer and Svensson, the secondary processor loads data directly from the ISA (hardware buffer) to the final destination in the XRAM (system memory) with "no teaching or suggestion . . . that the executable software image, before it is loaded from the hardware buffer to its final destination in the system memory, is copied to another location in the system memory." Pet. 53–54 (citing Ex. 1009 ¶¶ 16, 27, 31, 36, 43, Fig. 2; Ex. 1010, 4:1–3, 5:21–28, 5:65–67, 6:12–15, Figs. 1, 3; Ex. 1002 ¶¶ 156–159). Patent Owner counters that, in the asserted combination, "data is first loaded into one part of system memory—the intermediate storage area of the DSP SARAM and DARAM—and then copied to the DSP XRAM, which is another part of the system memory." PO Resp. 75–76 (citing Ex. 2007 ¶ 181).

For the reasons explained above in § II.D.4.b, we find that the ISA of Bauer and Svensson is a hardware buffer and that the asserted combination teaches directly loading data from that hardware buffer to the XRAM based on the load addresses in Bauer's file format. We agree with Petitioner that nothing in Bauer suggests additional copy operations in the XRAM (system memory). Rather, Bauer's file format provides load addresses for each data

segment, thus allowing efficient memory loading directly to the target destination. Ex. 1009 ¶¶ 37, 43.

Claim 3 recites, "The multi-processor system of claim 1 in which raw image data of the executable software image is received by the secondary processor via the interface." Petitioner argues that the combination of Bauer, Svensson, and Kim teaches this subject matter because the header is transferred first followed by the data segments such that the data segments are raw image data without packet headers. Pet. 54–56 (citing Ex. 1002 ¶¶ 160–164; Ex. 1009 ¶¶ 30, 32; Ex. 1001, 4:47–52, 7:39–42, 10:1–18).

Claim 4 recites, "The multi-processor system of claim 1 in which the secondary processor is configured to process the image header to determine at least one location within the system memory to store the at least one data segment." Petitioner refers to its contentions for the scatter loading limitations of claim 1 and argues that the secondary processor in Bauer and Svensson processes the image header to determine where to scatter load the data. Pet. 56–57 (citing Ex. 1002 ¶¶ 165–166).

Claim 5 recites, "The multi-processor system of claim 4 in which the secondary processor is configured to determine, based on the received image header, the at least one location within the system memory to store the at least one data segment before receiving the at least one data segment." Referring to its contentions for claim 1 and Bauer's disclosure of retrieving the header and section information before the data segments, Petitioner argues that a person of ordinary skill in the art "would have been motivated to process the image header before receiving the data segments because, upon receiving each data segment, the secondary processor would then more efficiently and quickly load each data segment from the hardware buffer to

the system memory, as taught by Bauer." Pet. 57–59 (citing Ex. 1002 ¶¶ 167–171); *see* Ex. 1009 ¶¶ 28–30 (Bauer's disclosures that header and section information is read and retrieved before image information).

Claim 6 recites, "The multi-processor system of claim 1, in which the secondary processor further comprises a non-volatile memory storing a boot loader that initiates transfer of the executable software image for the secondary processor." Petitioner notes that, in Svensson's disclosed system, the primary processor pushes the bootloader to the secondary processor's volatile memory, but Petitioner argues that it would have been obvious to include the bootloader in a non-volatile memory on the secondary processor. Pet. 59–65. In support of this contention, Petitioner cites Svensson's disclosure of storing a bootloader in a non-volatile memory, and Petitioner argues that there were a finite number of ways to store a bootloader, namely pushing it into RAM or loading it from non-volatile memory. Pet. 59–65 (citing Ex. 1010, 1:17–27, 1:48–56, 2:11–15, 2:47–3:2, 4:9–6:11, 6:44–47, 7:3–15, Fig. 2; Ex. 1009 ¶¶ 31; Ex. 1002 ¶¶ 172–181); *see* Ex. 1010, 1:20–25 ("The bootloader, which is sometimes called a boot loader or a bootstrap loader, is a set of instructions (i.e., program code, sometimes called 'boot code') that can be either 'pushed' into the system's RAM or loaded into the RAM from a non-volatile memory, such as read-only memory (ROM).").

Claim 7 recites, "The multi-processor system of claim 1 in which the primary and secondary processors are located on different chips." Petitioner argues that the ARM and DSP devices in Bauer and Svensson are on different chips based on various disclosures in the references, including that there is a "hardware boundary" between the processors. Pet. 65–66 (citing

Ex. 1009 ¶¶ 35–36, Fig. 2; Ex. 1010, 3:54–60, 3:64–4:1, Fig. 1; Ex. 1002 ¶¶ 182–183). Petitioner argues that, to the extent this is not disclosed in the combination, it would have been obvious because there are limited ways of constructing such a multi-processor system, namely putting the processors on the same chip or different chips. Pet. 66–67 (citing Ex. 1002 ¶¶ 184–185).

Claim 8 recites, "The multi-processor system of claim 1 in which the portion of the executable software image is loaded into the system memory of the secondary processor without an entire executable software image being stored in the hardware buffer." Petitioner argues that the combination of Bauer and Svensson teaches this subject matter because Svensson discloses pushing a portion of data to the ISA, loading from the ISA to system memory, and then pushing more data to the ISA for loading to system memory. Pet. 67–69 (citing Ex. 1009 ¶¶ 33, Ex. 1010, 5:53–611, 6:26–28, 6:37–43, 6:50–54, Fig. 2; Ex. 1002 ¶¶ 186–191).

Claim 9 recites, "The multi-processor system of claim 1 integrated into at least one of a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and a fixed location data unit." Petitioner contends this subject matter would have been obvious to a person of ordinary skill in the art, citing various disclosures in Bauer and Svensson of the multi-processor system in computers and mobile phones. Pet. 70–71 (citing Ex. 1009 ¶¶ 15, 16, 26; Ex. 1010, 7:61–63, 8:26–29; Ex. 1002 ¶¶ 192–193).

Having considered the full record developed during trial, on appeal to the Federal Circuit, and on remand, we are persuaded by Petitioner's

contentions and evidence, and we determine that Petitioner has proven by a preponderance of the evidence that claims 2–9 of the '949 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, and Kim.

6. Claim 12

Claim 12 recites, "The method of claim 10 further comprising loading the executable software image directly from a hardware buffer to the system memory of the secondary processor without copying data between system memory locations." In the Final Written Decision, we held that Petitioner had not proven unpatentability of claim 12 based on our determination at that time that the ISA in Bauer and Svensson is not a hardware buffer. Final Dec. 55–56. As discussed above, we now find that Bauer and Svensson's ISA is a hardware buffer, and we now hold claim 12 unpatentable.

As with claim 2, Petitioner argues that, in Bauer and Svensson, the secondary processor loads data directly from the ISA (hardware buffer) to the final destination in the XRAM (system memory) without copying data between locations in XRAM. 1335 Pet. 60–62. Patent Owner counters, as it does for claim 2, that the ISA is part of system memory and, therefore, that there is copying between system memory locations. For the reasons explained for claim 2, we disagree with Patent Owner's arguments and find that the asserted combination teaches directly loading data from the ISA (hardware buffer) to the XRAM (system memory) based on the load addresses in Bauer's file format. *See* § II.D.5 above. We agree with Petitioner that nothing in Bauer suggests additional copy operations in the XRAM. In the combination of Bauer and Svensson, the XRAM is the "system memory to which the secondary processor is coupled to store each

data segment," as recited in claim 10, because, as discussed in detail above, it is the final destination for the software that is being loaded from the primary processor, and it is coupled to the secondary processor and only accessible by the secondary processor, not by the primary processor. *See* Ex. 1009 ¶ 36 ("The ARM host CPU 202 can access the non-volatile memory 206 and the SARAM and DARAM 208 of the DSP, but not the DSP's XRAM 210 . . . .").

Having considered the full record developed during trial, on appeal to the Federal Circuit, and on remand, we are persuaded by Petitioner's contentions and evidence, and we determine that Petitioner has proven by a preponderance of the evidence that claim 12 of the '949 patent is unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, and Kim.

E. Obviousness over Bauer, Svensson, Kim, and Zhao (Claims 16, 17)

Petitioner contends that claims 16 and 17 of the '949 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, Kim, and Zhao. 1335 Pet. 67–77.

1. Zhao

Zhao discloses a mobile computing device with a dual processor architecture, as depicted in Figure 3, reproduced below. Ex. 1013 ¶ 32.

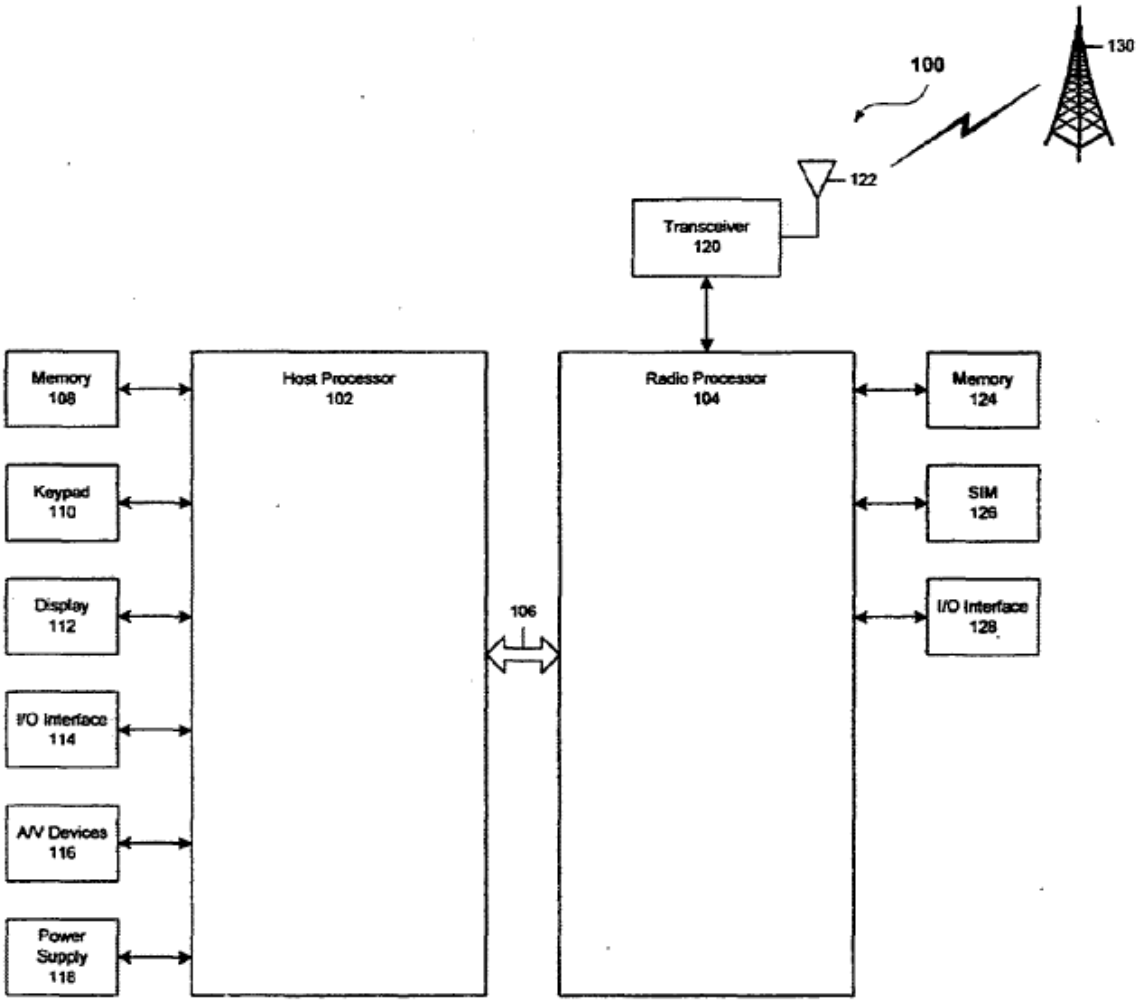

FIG. 3

Zhao's Figure 3, reproduced above, illustrates mobile computing device 100 having two processors, namely host processor 102 and radio processor 104. Ex. 1013 ¶ 32. With reference to Figure 3, Zhao explains the following:

> The host processor 102 and the radio processor 104 may be arranged to communicate with each other using interfaces 106 such as one or more universal serial bus (USB) interfaces, micro-USB interfaces, universal asynchronous receiver-transmitter (UART) interfaces, general purpose input/output (GPIO) interfaces, control/status lines, control/data lines, shared memory, and so forth.

Ex. 1013 ¶ 32.

2. Independent Claim 16

Independent claim 16 is directed to an apparatus that comprises a combination of elements recited in means-plus-function format. Each function in claim 16 is identically recited as a step of the method of claim 10. In essence, claim 16 is an apparatus version of claim 10 with "means for" inserted before each method step. Thus, Petitioner relies on its contentions for claim 10 to show that the combination of Bauer, Svensson, and Kim teaches the recited functions, and Petitioner relies on these references in combination with Zhao to teach the specific structures required of the means-plus-function limitations. 1335 Pet. 67–77.

In the Final Written Decision, we addressed the parties' arguments with respect to claim 10 and determined that Petitioner had proven that claim 10 is unpatentable. Final Dec. 26–50. That determination as to claim 10 is undisturbed because Patent Owner dropped its challenges concerning claim 10 on appeal. *See Intel*, 21 F.4th at 806 ("The Board ruled that Intel had proved the unpatentability of claims 10, 11, 13–15, and 18–23, but Qualcomm, despite filing a cross-appeal to raise the issue, no longer challenges that ruling."). For the reasons stated in our Final Written Decision as to claim 10, we agree with Petitioner that the combination of Bauer, Svensson, and Kim teaches the recited functions of claim 16. *See* Final Dec. 26–50.

For the "means for receiving . . . an image header" and "means for receiving . . . each data segment," Petitioner argues that the combination of Bauer and Svensson "teaches an ***apparatus*** (multi-processor system) having a ***secondary processor*** (DSP device) connected to a ***primary processor*** (ARM device) by an ***inter-chip communication bus*** (buses and/or DMA

paths)" but "does not teach the specific inter-chip communication buses required by the corresponding structure," for which Petitioner relies on Zhao. 1335 Pet. 68–69 (citing Ex. 1009 ¶¶ 35–36, Fig. 2; Ex. 1020 ¶ 183).

Petitioner argues that Zhao teaches various interfaces between two processors in a mobile device, including a universal asynchronous receiver-transmitter (UART) interface. 1335 Pet. 69–70 (citing Ex. 1013 ¶¶ 32, 34, 44, Figs. 3, 5; Ex. 1020 ¶¶ 184–185). As noted above, Zhao explains the following:

> The host processor 102 and the radio processor 104 may be arranged to communicate with each other using interfaces 106 such as one or more universal serial bus (USB) interfaces, micro-USB interfaces, *universal asynchronous receiver-transmitter (UART) interfaces*, general purpose input/output (GPIO) interfaces, control/status lines, control/data lines, shared memory, and so forth.

Ex. 1013 ¶ 32. (emphasis added).

Petitioner argues that a person of ordinary skill in the art "would have found it to be a simple known alternative to select an inter-chip communication bus from amongst the standardized inter-chip communication buses that were known in the prior art, such as those taught in Zhao," and that "[t]he use of standardized inter-chip communication buses . . . simply involves the combination of prior art elements (chips and standardized inter-chip communication buses) according to known methods to achieve a predictable result with a reasonable expectation of success" and, therefore, would have been obvious. 1335 Pet. 70–71 (citing Ex. 1020 ¶ 186).

To teach "a modem processor coupled to system memory," the structure corresponding to the "means for processing" and "means for scatter

loading" limitations, Petitioner argues that, although the combination of Bauer and Svensson teaches a digital signal processor (DSP) coupled to system memory, Bauer and Svensson do not expressly describe the DSP as a modem processor. 1335 Pet. 72. Petitioner argues that Zhao discloses a multi-processor system having a host processor and a radio processor, which can be a modem processor. 1335 Pet. 73 (citing Ex. 1013 ¶¶ 26, 32, 44; Fig. 3; Ex. 1020 ¶ 190). For example, Zhao discloses that "radio processor 104 may be implemented as a communications processor using any suitable processor or logic device, such as a modem processor or baseband processor." Ex. 1013 ¶ 44. Petitioner argues that it would have been obvious to implement the secondary processor in Bauer and Svensson as a modem processor "to allow for the transfer of data over a mobile network, including allowing a user to make and receive calls, as taught by Bauer and Svensson combined." 1335 Pet. 74 (citing Ex. 1009 ¶ 15; Ex. 1010, 7:61–63, 8:26–29). Petitioner argues that a person of ordinary skill in the art "would have understood that this would have been obvious to try and would have involved a simple substitution of one known feature for another." 1335 Pet. 74 (citing Ex. 1020 ¶ 191).

Patent Owner's arguments for claim 16 all concern the functional recitations, which we addressed and rejected in our discussion of claim 10 at pages 26–50 of the Final Decision. Patent Owner does not raise any additional arguments specific to claim 16, and Patent Owner even states that the means-plus-function limitations of claim 16 do not need to be construed because "[n]one of the arguments [Patent Owner] makes . . . to distinguish the prior art requires construction of these limitations." PO Resp. 17; *see*

PO Remand Br. 19 (not raising additional arguments for claim 16 and 17 but, instead "rest[ing] on its briefs").

Petitioner's arguments as to how Zhao teaches specific structures required of the means-plus-function limitations are persuasive. In particular, Zhao expressly discloses a UART interface for communicating between processors in a multi-processor system, and it discloses using a modem processor in such a system. Ex. 1013 ¶¶ 32, 44. We also are persuaded by Petitioner's contentions (summarized above) that a person of ordinary skill in the art would have had reason to combine these teachings of Zhao with the teachings of Bauer, Svensson, and Kim (as discussed for claim 10) with a reasonable expectation of success.

We have considered the full record developed during trial and on remand. For the reasons discussed above, the reasons given in the 1335 Petition, and the reasons discussed in the Final Decision for claim 10, Petitioner has demonstrated by a preponderance of the evidence that claim 16 of the '949 patent is unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, Kim, and Zhao.

3. Claim 17

Claim 17 recites, "The apparatus of claim 16 integrated into at least one of a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and a fixed location data unit." Petitioner refers to its contentions for claim 15, which depends from independent claim 10 and recites "performing the receiving, processing, and loading, in at least one of a mobile phone, a set top box, a music player, a video player, an entertainment unit, a navigation device, a

computer, a hand-held personal communication systems (PCS) unit, a portable data unit, and a fixed location data unit." 1335 Pet. 65–67 (citing Ex. 1009 ¶¶ 15, 16, 26; Ex. 1010, 7:61–63, 8:26–29; Ex. 1020 ¶ 179), 77 (citing Ex. 1020 ¶ 198). Patent Owner does not separately address the additional subject matter of claim 17. *See* PO Resp. We agree with Petitioner that the asserted combination teaches that the apparatus is at least part of a "computer" or a "mobile phone." *See* Ex. 1010, 7:61–63 ("It is expected that this invention can be implemented in a wide variety of environments, including for example mobile communication devices.").

Having considered the full record developed during trial and on remand, we are persuaded by Petitioner's contentions and evidence, and we determine that Petitioner has demonstrated by a preponderance of the evidence that claim 17 of the '949 patent is unpatentable under 35 U.S.C. § 103(a) as obvious over the combined teachings of Bauer, Svensson, Kim, and Zhao.

III. CONCLUSION[14]

For the reasons discussed above, we determine Petitioner has proven by a preponderance of the evidence that the claims at issue on remand are unpatentable, as summarized in the following table:

Claims	35 U.S.C. §	References	Claims Shown Unpatentable	Claims Not Shown Unpatentable
1–9, 12	103(a)	Bauer, Svensson, Kim	1–9, 12	
16, 17	103(a)	Bauer, Svensson, Kim, Zhao	16, 17	
Overall Outcome			1–9, 12, 16, 17	

[14] 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. 16654 (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), (b)(2).

IV. ORDER

Accordingly, it is:

ORDERED that claims 1–9, 12, 16, and 17 of the '949 patent have been shown to be unpatentable; and

FURTHERED 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.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,
Petitioner

v.

QUALCOMM INCORPORATED,
Patent Owner.

IPR2018-01334
Patent 8,838,949 B2

Before TREVOR M. JEFFERSON, DANIEL J. GALLIGAN, and AARON W. MOORE, *Administrative Patent Judges.*

MOORE, *Administrative Patent Judge*, dissenting-in-part.

I respectfully dissent from the majority's construction of "hardware buffer" and its determination that claims 1–9 and 12 have been proven unpatentable.

Essentially, I conclude that the majority fails to appreciate or acknowledge that the central invention, or "asserted advance," of U.S. Patent No. 8,838,949 B2 ("the '949 patent") was the inventors' determination that they could substantially increase boot times by scatter loading *directly from the USB controller of the secondary processor*—i.e., from its "hardware buffer"—instead of passing the data from the USB controller to a temporary buffer and then scatter loading from there to the final location.

With that understanding, it is clear to me that the "hardware buffer" the applicants subsequently described and claimed is a dedicated, permanent buffer, such as that found in a USB controller or the like, and that no such buffer exists in Petitioner's prior art combination.

I. BACKGROUND

A. The '949 Patent

The '949 patent is directed to a multi-processor system in which a primary processor is coupled to non-volatile memory storing an executable software image of a secondary processor. The secondary processor is coupled to a dedicated volatile memory, and the executable software image is "efficiently communicated" from the primary processor to the secondary processor in a "segmented format" and "using a direct scatter load process." Ex. 1001, 1:24–33.

The patent explains that "[p]rocessors may require some software code, commonly referred to as boot code, to be executed for booting up" but that "[a] problem exists on a significant number of devices (such as smart phones)" that incorporate multiple processors. Ex. 1001, 1:38–39, 45–46.

The patent further explains that "[a] processor's boot code may be stored to [each] processor's respective non-volatile memory . . . and upon power-up the boot code software [may be] loaded for execution by the processor from its respective non-volatile memory." Ex. 1001, 1:51–56. In that scenario, "the executable software . . . is not required to be loaded to the processor from another processor in the system." *Id*. at 1:56–59. However, "[a]dding dedicated non-volatile memory to each processor . . . occupies more circuit board space, thereby increasing the circuit board size." *Id*. at 1:60–62.

To deal with the space problem, some multi-processor systems would load the required software from one processor to another processor. *See* Ex. 1001, 2:1–3. A first processor would be responsible for storing boot code for other processors in the system, and "upon power-up the first processor [would be] tasked with loading the respective boot code to the other processors." *Id*. at 2:3–8.

The patent notes that one way of loading a software image from a primary processor to a secondary processor was "to allocate a temporary buffer into which each packet is received," where "each packet would have an associated packet header information" and a "payload [which] would be the actual image data." Ex. 1001, 2:25–28. "[T]he buffer [would be] used to receive part of the image data from the primary processor, and from the buffer the image data [would be] be scattered into the memory (e.g., volatile memory) of the secondary processor." *Id*. at 2:37–41. According to the patent, "[t]he temporary buffer would be some place in system memory, such as in internal random-access-memory (RAM) or double data rate (DDR) memory." *Id*. at 2:32–34.

The patent then summarizes the invention as a system that "includes a secondary processor having a system memory and a hardware buffer for receiving at a least a portion of an executable software image," where "[t]he secondary processor includes a scatter loader controller for loading the executable software image *directly from the hardware buffer* to the system memory." Ex. 1001, 2:58–63 (emphasis added). The system further includes "a primary processor coupled with a memory [that] stores the executable software image for the secondary processor" and "an interface communicatively coupling the primary processor and the secondary

processor via which the executable software image is received by the secondary processor." *Id*. at 2:64–3:2.

The patent explains that "[a]spects of [its] disclosure provide techniques that alleviate the intermediate step of buffering required in traditional loading processes," specifically "avoid[ing] extra memory copy operations, thereby improving performance (e.g., reducing the time required to boot secondary processors in a multi-processor system)." Ex. 1001, 7:24–30.

This "Zero Copy Transport Flow" is described in connection with Figure 3, a portion of which is reproduced below.

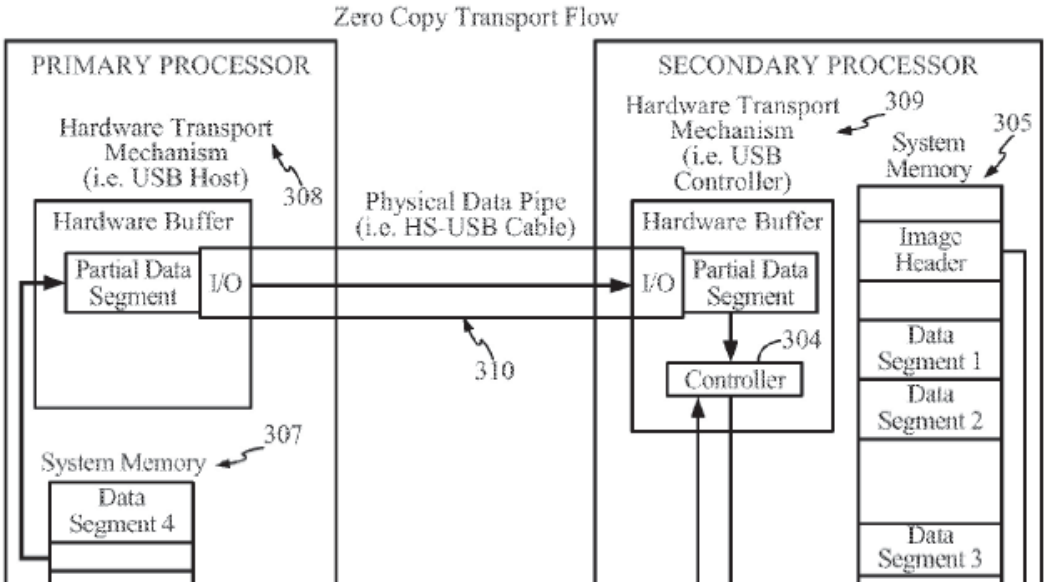

Excerpt from Figure 3 of the '949 Patent

"The primary processor 301 parses the image header to load individual image segments from non-volatile memory of the primary processor 306 to system memory of the primary processor 307." Ex. 1001, 8:13–17. That header information is then "used by the secondary processor

302 to program the scatter loader/direct memory access controller 304 receive address when receiving the actual executable data." *Id*. at 8:20–23.

Data segments are "sent from system memory 307 to the primary hardware transport mechanism 308" and "then sent from the hardware transport mechanism 308 of the primary processor 301 to a hardware transport mechanism 309 of the secondary processor 302 over an inter-chip communication bus 310 (e.g., a HS-USB cable.)." Ex. 1001, 8:24–30.

The image header, which "contains information used by the secondary processor to locate the data segments into target locations in the system memory of the secondary processor 305," may be transferred first. Ex. 1001, 8:31–33. The patent explains that "[t]he image header generally includes a list of segment start addresses and sizes defining where each of the segments should be loaded in the secondary processor's system memory 305." *Id*. at 8:57–60.

"Secondary processor 302 includes a hardware transport mechanism 309 (e.g., a USB controller) that includes a scatter loader controller 304" that is programed "for the transfer of the entire segment directly into the target memory location (within system memory 305) with minimum software intervention by the secondary processor 302." *Id*. at 8:60–61, 9:6–8.

The result is that "no extra memory copy operations occur in the secondary processor"[1] and "conventional techniques employing a temporary buffer for the entire image, and the packet header handling, etc., are

[1] "Zero copy" is something of a misnomer because the data *is* copied, at least to the hardware buffer and then again to the final location. It might be more accurate to describe this as a "zero *extra* copy" or "zero *intermediate* copy" arrangement.

bypassed in favor of a more efficient direct loading process." Ex. 1001, 9:41–45.

B. Procedural History

The Petition alleged that "Bauer and Svensson combined discloses 'a secondary processor comprising . . . a system memory and a hardware buffer'" because "Figure 2 of Bauer shows a secondary processor (DSP device)" and "[t]he secondary processor comprises a DSP CPU 204 coupled to system memory (DSP XRAM 210), and a separate hardware buffer (intermediate storage area (Int. Store Area) within the DSP SARAM & DARAM memory 208)." Pet. 26–27 (citations omitted).

According to Petitioner, "the disclosure of an intermediate storage area that is separate from the system memory (DSP XRAM) is the same alleged inventive feature of the '949 patent," as the patent "makes a distinction between prior art systems that used a 'temporary buffer' that was part of the system memory . . . , and the alleged invention that uses a 'hardware buffer' separate from the system memory." *Id.* at 27 (citations omitted).

In the Final Written Decision, we concluded that the "hardware buffer" limitations "should not be read so broadly as to encompass" the use of a temporary buffer and that "[n]o further interpretation of 'hardware buffer' [was] necessary to resolve the obviousness inquiry." Paper 30, 17. We accordingly found claims 1–9 and 12 patentable over the combination.

The Federal Circuit reversed and remanded, finding that we "did not do enough to reach and articulate" an "understanding what the intrinsic evidence makes clear is the substance of the invention—what the inventor 'intended to envelop.'" *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 801, 809

(2021). The Court concluded that our construction of "hardware buffer" was not necessarily incorrect, but "wanting." *Id*.

To help guide a renewed inquiry, the Court observed that (1) "because every buffer in our (physical) world is ultimately implemented on a physical device (i.e., hardware), a 'hardware buffer' must mean something more than just a 'buffer implemented in hardware,' . . . or else the word 'hardware' would be erased from the claims," (2) "because claim 1 requires both a 'system memory' and a 'hardware buffer,' there must be some distinction between those two concepts," and (3) "the meaning of 'hardware buffer' relates to the ability to move the software image 'directly' to the second processor's system memory and to avoid 'copying data between system memory locations.'" 21 F.4th at 809–810.

On remand, Petitioner argues that "hardware buffer" should be construed as "memory that is physically separate from the memory into which the software image is loaded for execution." Paper 35, 5. Patent Owner contends that it should be interpreted to mean "a permanent, dedicated buffer that is distinct from system memory." Paper 37, 3.

II. CLAIM CONSTRUCTION

The challenged claims must be given the "broadest reasonable interpretation in light of the patent specification" because the *inter partes* review petitions at issue were filed before November 13, 2018. *See* 83 Fed. Reg. 51,340 (Oct. 11, 2018).

I believe that Patent Owner's proposed construction fairly captures the essence of the invention, and that Petitioner's proposal is broader than would be reasonable in view of the intrinsic evidence.

A. The Claim Language

I start with the language of claim 1, which recites "a secondary processor" comprising "system memory and a hardware buffer for receiving an image header and at least one data segment of an executable software image." As the Federal Circuit observed, "hardware" must mean more than just that the buffer is realized in or on hardware "because every buffer in our (physical) world is ultimately implemented on a physical device (i.e., hardware)." 21 F.4th 801, 809. Notably, the other components recited in the claim, e.g., the "system memory" and the "scatter loader controller," are not similarly modified by "hardware," which seems to confirm that "hardware" in "hardware buffer" must mean something *other* than that the buffer exists on or in hardware.

The majority's discussion of the claim language describes the faults found by the Federal Circuit, but does not otherwise explain how the language of the claims themselves might bear on the constructions proposed by the parties. *See* Final Written Decision on Remand ("RFWD") 12–13.

I find that the claim language itself weighs *against* Petitioner's proposed construction because that construction does not give any specific meaning to "hardware." For example, Petitioner does not explain why "physically separate" memory would be a "hardware buffer," but the same memory would not be a "hardware buffer" if not physically separate. Petitioner is essentially arguing that "hardware buffer" means "separate buffer,"[2] but Petitioner does not explain why "hardware" should be interpreted to mean "separate."

[2] *See* Remand Hearing Transcript, Paper 46, 14 (Petitioner's Counsel: "So the meaning given to hardware is that it's physically – the hardware buffer is

Patent Owner's construction, on the other hand, would require that the "hardware buffer" be "*permanent*" and "*dedicated*," thus distinguishing a "hardware buffer" from a "buffer" in general, and giving that claim language meaning consistent with the invention, as explained below.

B. The Written Description

The majority's discussion of the specification concludes that the portions of the specification referring to the "hardware buffer" do not help define the term. *See* RFWD 14. I strongly disagree.

The patent describes and depicts the "hardware buffer" of the secondary processor as a buffer within a "Hardware Transport Mechanism (i.e., USB Controller)." *See* Ex. 1001, Fig. 3. It also describes and depicts a "hardware buffer" of the primary processor as "Hardware Transport Mechanism (i.e., USB Host)." *Id*. These are shown in the annotated excerpt from Figure 3 below.

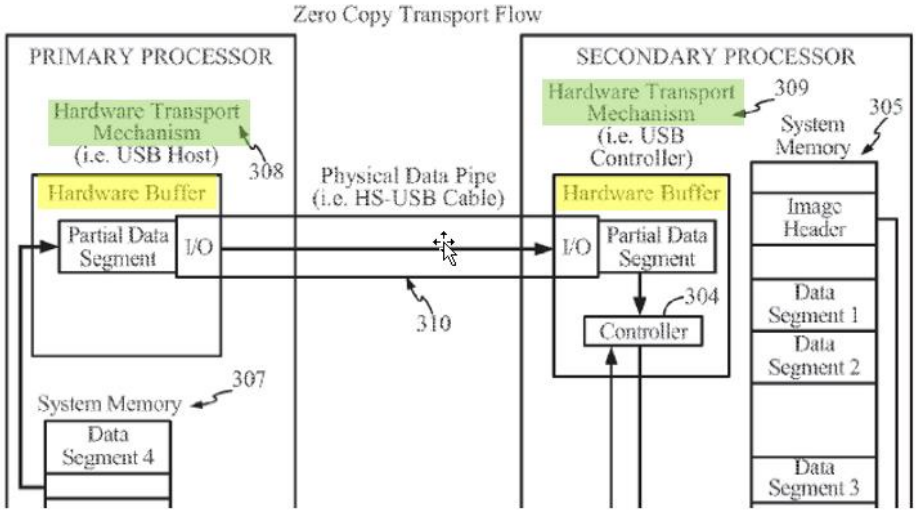

Except from Figure 3 of the '949 Patent

physically separate from the claimed system memory. There are separate pieces of hardware. That's the meaning given to the term hardware and hardware buffer in Petitioner's construction.").

The patent explains that "[d]ata segments are . . . sent from system memory 307 to the primary hardware transport mechanism 308" and then "sent from the hardware transport mechanism 308 of the primary processor 301 to a hardware transport mechanism 309 of the secondary processor 302 over an inter-chip communication bus 310 (e.g., a HS-USB cable.)" Ex. 1001, 8:24–30. Then, a scatter loader controller 304 in the "hardware transport mechanism 309 (e.g., a USB controller)" of the secondary processor 302 is programed "for the transfer of the entire segment directly into the target memory location (within system memory 305) with minimum software intervention by the secondary processor 302." *Id.* at 8:60–9:8.

Thus, instead of sending the data from the hardware buffer of the USB Host, to the hardware buffer of the USB controller, to a temporary buffer in system memory, and then scatter loading it to the final destination, the inventive system *omits the temporary buffer in system memory*, and the scatter loading is done directly from the hardware buffer of the USB controller to the final destination.

The difference can be illustrated like this:

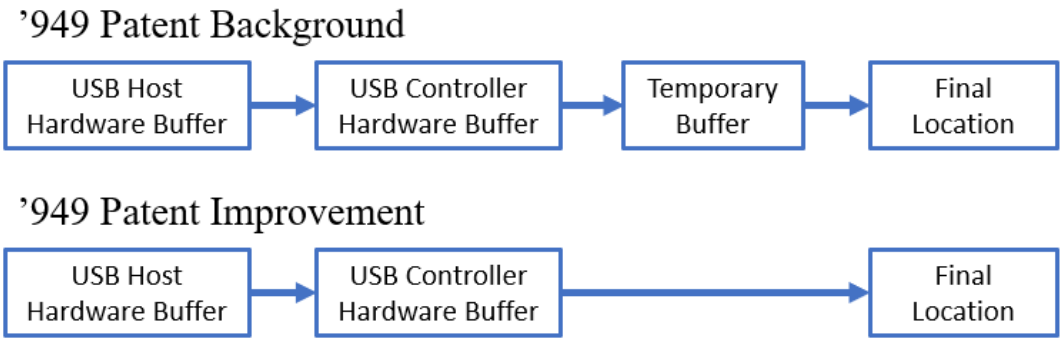

Block Diagram Showing the Improvement of the '949 Patent

Thus, Petitioner's assertion that the patent "makes a distinction between prior art systems that used a 'temporary buffer' that was part of the system memory . . . , and the alleged invention that uses a 'hardware buffer' separate from the system memory," Paper 3, 27, is not really accurate. The true distinction is that the prior system used ***both*** a hardware buffer and a temporary buffer, while the inventive system used ***only*** the hardware buffer.

With that understanding of the invention, I have no trouble concluding that one of ordinary skill in the art would have understood that the '949 patent is using "hardware buffer" to refer to a dedicated, permanent buffer, such as that found in a USB host or controller, as shown—*and actually labeled "hardware buffer"*—in Figure 3. *See* Ex. 2007 ¶ 70; Ex. 2015 ¶ 21.

Because the core of the invention was the concept of scatter loading directly from the USB controller to its final location, it's little wonder that the applicant wrote claims specifically directed to that embodiment.

Petitioner argues that "[b]eyond being separate from the claimed system memory, the hardware buffer is not described as being, and a POSITA would not understand it as being, located in any particular area within the secondary processor." Paper 35, 8. For the reasons stated above, I disagree: the hardware buffer is clearly described (and pictured) as being located within the USB controller, and nowhere else.

Petitioner next argues that "[t]he specification . . . aims to distinguish what is claimed from other aspects of the prior art, not the use of temporary buffers per se." Paper 35, 9. Specifically, Petitioner argues that the specification (a) "distinguishes systems in which the ***entire executable software image*** is copied into a temporary buffer" and (b) distinguishes "between the buffer receiving a header separately from data segments (as

claimed) and receiving a ***packet that includes both a header and data segments***." *See id.* at 9–11.

Petitioner's assessment ignores the forest to focus on the trees. As explained above, the inventors were dealing with a prior art system in which the data was received by the USB controller, then copied to a temporary location, then copied to the final location. They were trying to determine how to make *that* system, *which included the USB controller*, faster. And they did it by *eliminating the temporary buffer and using the USB controller itself as the scatter loader*. The fact that the image was written as segments instead of an entire image is just a consequence of the core idea of using the USB controller as the scatter loader because, presumably, the hardware buffer would not have been large enough to hold the entire image. Similarly, the use of a single header was not itself the invention, but rather one of the techniques the inventors developed *in order to use the USB controller as a scatter loader*.

Following Petitioner's lead, the majority focuses on the details (i.e., loading the image as segments, using a separate header) that the inventors worked out to allow the USB controller to scatter load the data directly. Those details were necessary to implement the idea of scatter loading from the USB controller, but the bigger picture is that, as the specification tells us, the point of the "direct scatter load technique" was to "avoid[] use of a temporary buffer." Ex. 1001, 4:43–52. That was accomplished in the system of the patent by scatter loading directly from the "hardware buffer" of the USB controller to the final memory locations.

The majority also asserts that it does "not see a relevant distinction between the operation of the '949 patent and the asserted prior art." RFWD 18. I find this point misplaced, for two reasons.

First, our task is to construe the claims in view of the intrinsic evidence and only *then* compare the construed claims to the prior art. *See Enzo Biochem, Inc. v. Applera Corp.*, 599 F.3d 1325, 1332 (Fed. Cir. 2010). From *the claim construction perspective*, the asserted prior art is just extrinsic evidence that does not appear to aid in the search for the meaning of "hardware buffer," as neither reference even uses that term.

Second, the distinction between the invention and the prior art combination is that the invention uses the hardware buffer of the secondary processor's USB controller to write the data directly to its final location, where the combination has the data written into the system memory of the secondary processor (the ISA) and then copied to the final location. There is only "no distinction" when one first concludes that "hardware buffer" would cover the ISA. This point is thus inherently and improperly circular. The majority cannot reasonably support a construction using a conclusion that arises from the construction itself.

For these reasons, I would find that the specification requires the claimed "hardware buffer" to be dedicated and permanent, as in the USB controllers described and shown in the patent, and that it would be unreasonable for it to be interpreted to encompass "temporary" buffers, which are not described or even suggested as a substitute for the hardware USB controller.

To be perfectly clear, there are *two* "buffers" discussed in the patent. The first is the buffer in the USB controller (what I find to be the "hardware

buffer"), and the second is the temporary buffer in system memory. The point of the invention was to eliminate the second, temporary buffer, and scatter load directly from the first, permanent buffer. There is simply no suggestion in the patent that the first buffer might by anything other than a permanent, dedicated buffer, such as that of the disclosed USB controller.

C. The Prosecution History

Regarding the prosecution history, the majority finds that the applicant's silence when the Examiner mapped the "hardware buffer" to Svensson's ISA, which the applicant avoided with a different amendment, "should not necessarily be held against it." RFWD 21. I agree. *See, e.g.*, *Biogen Idec, Inc. v. GlaxoSmithKline LLC*, 713 F.3d 1090, 1099–100 (Fed. Cir. 2013) ("[I]t is the applicant, not the examiner, who must give up or disclaim subject matter that would otherwise fall within the scope of the claims."); *Salazar v. Procter & Gamble Co.*, 414 F.3d 1342, 1344–48 (Fed. Cir. 2005); *3M Innovative Properties Co. v. Avery Dennison Corp.*, 350 F.3d 1365, 1373–74 (Fed. Cir. 2003).

However, the majority goes on to "find that the applicant's affirmative statements[3] in this case inform the meaning of the term 'hardware buffer.'" RFWD 21. They point to *University of Massachusetts v. L'Oréal S.A.*, 36 F.4th 1374 (Fed. Cir. 2022), and its explanation that '"the prosecution history [need not] meet the standard of clear and unmistakable disclaimer for overcoming an otherwise plain meaning' because 'the meaning of the relevant claim language is not plain.'" RFWD 23. The majority then finds

[3] As explained below, the "affirmative statements" to which the majority refers were not about the "hardware buffer" limitation, but instead related to how other limitations distinguished the claims over the cited art.

that "[t]he meaning of the term 'hardware buffer' is not plain," so "this is not a case where the prosecution history must meet the standard of clear and unmistakable disclaimer for overcoming an otherwise-plain meaning," and "the prosecution history shows that the 'asserted advance' was not the use of a hardware buffer." *Id*. at 23. The majority then says "[t]he applicants' arguments in prosecution rely on different features to distinguish the asserted prior art" and, thus, "suggest that features other than the 'hardware buffer' were the key distinctions." *Id*. at 24.

This makes no sense to me. First, I find the majority's reliance on *L'Oréal* misplaced. In that case, the Federal Circuit noted that the general rule requiring a "clear and unmistakable disclaimer" did not apply because "the meaning of the relevant claim language [was] not plain." 36 F.4th 1383. The court then considered prosecution history that *specifically involved the disputed term*, finding that the applicant's affirmative statements during prosecution *about the disputed wherein clause*, bore on the construction of that clause.

But here the majority is not pointing to the prosecution history about the term "hardware buffer," *because there is none*. Instead, the majority is faulting the applicant for choosing to focus its patentability arguments on aspects *other than* the "hardware buffer." This is simply not consistent with *Biogen*, *Salazar*, and *3M Innovative Properties*, with which the majority appears to agree. *L'Oréal* does not override those cases; instead, they stand for different propositions—*L'Oréal* for the idea that disclaimer does not always need to be "clear and unmistakable," and *Biogen*, *Salazar*, and *3M Innovative Properties* for the idea that disclaimer will not be found in an applicant's silence on an issue.

The inventors wanted to eliminate the use of the temporary buffer and determined that they could instead copy the image data directly from the USB controller to its final location. In order to do that, they sent a single header first, so the controller could use the header information to place the segments into the correct locations in system memory as they were received. So, the invention, really, was (a) direct copying from the relatively fast USB controller to the final location, and (b) doing that by sending a single header first to allow direct loading of the image data.[4] The applicant was free to traverse the Examiner's combination on either basis, and silence in response to the examiner's characterization of a different aspect of the claim is not acquiescence to that characterization. *See Biogen*, 713 F.3d at 1099–1100.

The majority's approach, that silence will work against an applicant if claim language is later deemed not "plain," does not appear to be the law and is not consistent with conventional prosecution practice.[5]

The majority also points to an argument in which the applicant stated that "the independent claims 11, 17, 19, 21, and 23 also recite separately receiving the image header and each data segment and scatter loading each received data segment directly from the hardware buffer to the system memory," and finds this is "significant because none of the other independent claims recited a 'hardware buffer.'" RFWD 24. According to

[4] I understand that the individual segments did not need headers because all of the information that would have been in the individual headers was being sent first.

[5] Holding that *Biogen*, *Salazar*, and *3M Innovative Properties* would not apply should an issued claim term be found to not have a "plain" meaning would be very problematic from a prosecution perspective. Patent prosecutors would need to rebut every argument (perhaps unnecessarily limiting their claims) for fear of silence later being used against the patentee.

the majority, "the applicants affirmatively represented to the Office the reasons they thought the claims were allowable," and "the fact that five of the six pending independent claims did not recite 'hardware buffer' is strong evidence that other features, such as the separate receipt of image header and data segments, constituted the 'asserted advance.'" *Id*. at 25.[6]

This seems to me an untenable reading of the situation. It appears that what actually happened is that the prosecutor mistakenly thought all of the claims recited "hardware buffer." If anything, it suggests the applicant viewed "hardware buffer" as important to the invention. And it seems a stretch to say that an argument that *included* "hardware buffer" is "strong evidence" that something *else* was the "asserted advance," as the majority contends.

I see nothing in the file history to indicate that "hardware buffer" was meant to encompass "temporary" buffers. As noted, in the patent, the "hardware buffer" was the first (USB controller) buffer and the "temporary buffer" was the second (system memory) buffer that the inventors were eliminating.

D. Extrinsic Evidence

The majority considers the trial testimony of inventor Steve Haehnichen and finds it to show that "any 'asserted advance' had more to do with receiving the image header before the data than it had to do with the use of a 'hardware buffer' or the particular characteristics of such a buffer." RFWD 28–29. I think that conclusion is incorrect.

[6] The majority finds this "even more probative of the 'asserted advance'" than the applicant's argument that image header limitations distinguished the claims over the prior art. RFWD 24.

Mr. Haehnichen explained that, while trying to achieve a faster boot time, the engineers determined that "USB speed is fast, but then the processing inside [i.e., the copying from the temporary buffer to the final location] was slowing things down," so they "really focused on getting rid of all of those stages of copying[,] so things would come *right in the pipe from USB and land right in memory where they needed to be*." Ex. 2003, 216:18–217:2 (emphasis added). He compared the new process "to like a nonstop flight" where "instead of a flight with a stop in the middle and a long layover, you just put it right where it needs to be at the end." Ex. 2004, 222:8–10. In that analogy, the temporary buffer in system memory was the long layover. The core idea was to use the USB controller to write the data to the final location, instead of to a temporary buffer that would then require another copy operation.

In order to determine how to write directly from the USB controller to the final destination, the engineers on the team brought in "the experts on that piece of hardware, which is the USB receiver and [told] them what [the team was] trying to do." Ex. 2004, 222:13–15. The experts (also named inventors) told the team "how the hardware works and what the constraints are of that hardware" and "how [they] would have to organize and set it up so that [they] could receive it from the very first byte into the address it needs to be." *Id*. at 222:15–19.

Mr. Haehnichen explained that, in order to make it work, they had "to first get the headers of each image," where "the headers told [them] the destination, the size, characteristics of the data so that [they] could program the USB hardware data directly and have it allowed to receive right away." Ex. 2002, 223:3–7. In other words, they found that they needed to send the

header information first so that the USB controller would know where to put the data as it was received.

In my view, Mr. Haehnichen's testimony confirms that the invention was the use of the USB controller as the scatter loader, and the various other things, including the separate header, needed to make that work. Thus, the "hardware buffer" was part of—if not the heart of—the inventors' "asserted advance." This extrinsic evidence thus firmly supports the conclusion that the "hardware buffer" is a permanent, dedicated buffer, such as that of the USB controller the inventors were seeking to use as a direct scatter loader in order to avoid the need for the slow, temporary, intermediate buffer.[7]

E. Claim Construction Conclusion

Because I find Patent Owner's proposed construction sufficient to capture the essence of "hardware buffer" as that term is used in the '949 patent, I would construe it to mean "a permanent, dedicated buffer that is distinct from system memory." I find that Petitioner's proposed construction is not consistent with the specification or the extrinsic evidence discussed above, and that it appears transparently formulated to support Petitioner's prior art argument.

The majority does not pick either party's proposed construction, concluding instead that "the term 'hardware buffer' is not limited to a

[7] The majority asserts that "Mr. Haehnichen mentions 'configur[ing] the USB hardware to transfer directly into that memory' . . . , but, to the extent configuring USB hardware in particular represented some advance over the art, that advance is not reflected in the claims of the '949 patent." RFWD 30. I disagree. The advance over the art *is* reflected in these claims, as they specifically recite "scatter load[ing] . . . directly from the hardware buffer to the system memory."

'permanent' buffer.'" RFWD 33. The majority then asserts that "in the affirmative . . . the term 'hardware buffer' encompasses within its 'range of reasonable meanings,'" the constructions advanced by *both* parties—except that it removes "permanent" from Patent Owner's construction. *See id.*

I find the removal of "permanent" from Patent Owner's construction incorrect, for the reasons explained above, but I also find the majority's construction inappropriately vague, as it is entirely unclear what, exactly, falls within the "range of meanings." Must the hardware buffer be "physically separate," which is required by Petitioner's construction, but not Patent Owner's?[8] Need it be "dedicated," as in Patent Owner's construction but not Petitioner's?

Instead of determining what a "hardware buffer" *is*, the majority really just tells us one thing it *is not*. I believe the appeals court was, rightly, expecting more. *See* 21 F.4th at 812 (explaining that an "additional, substantive understanding . . . seems likely to support an affirmative construction in place of the Board's purely negative one").

Finally, I turn back to the three observations the Federal Circuit offered to guide our analysis.

The first was that "hardware buffer" must mean something more than just a "buffer implemented in hardware." Patent Owner's construction satisfies that condition because it requires that the buffer be "permanent" and "dedicated," like a buffer in a device such as a USB controller, rather than one that was dynamically allocated. The majority's construction, which

[8] Like Petitioner, the majority does not explain why "physically separate" should be included in the construction of "hardware buffer."

essentially only requires that "hardware buffer" is "not limited to a permanent buffer," does not appear to address this concern.

The second was that there must be some distinction between "system memory" and a "hardware buffer." Patent Owner's construction also satisfies that concern because it requires that the hardware buffer "is distinct from system memory," like the buffer of the USB controller described in the specification. The majority's construction also does not appear to account for this concern, as I see no reason why a buffer that was simply "not permanent" could not exist in system memory.

The third point was that the construction should relate to the ability to move the software image directly to the system memory and avoid copying data between system memory locations. Patent Owner's construction is narrowed to a "permanent, dedicated" buffer like the buffer in the USB controller that the inventors used to avoid the extra copying. It is not clear to me how the majority's construction, "not limited to a permanent buffer," might relate to avoiding the extra copy.

III. PATENTABILITY

Petitioner argues that "even if 'hardware buffer' is construed more narrowly than proposed by Petitioner, the Board should still find the challenged claims obvious" because "even under a construction that excludes temporary buffers, the ISA would satisfy the 'hardware buffer' limitation, since the ISA is never deallocated and instead functions solely as a permanent buffer." Paper 35, 17 (citing Ex. 1026 ¶¶ 41, 44–45).

I do not agree that the ISA is "a permanent, dedicated buffer that is distinct from system memory" for the simple reason that, as Petitioner's expert admits, it "is reserved at boot up of the client [DSP] processor." Ex. 1026 ¶ 41. The ISA would not exist until allocated, and then would exist at that location in the memory only until the system was shut down or rebooted, after which, upon rebooting, another ISA would be established. *See* Ex. 2007 ¶¶ 112–113, 135–144; Ex. 2015 ¶¶ 64–67. In fact, the panel previously determined that "the intermediate storage area of Bauer and Svensson is a temporary buffer." Paper 30, 56. Moreover, the construction I would adopt also requires that the "hardware buffer" be "distinct from system memory," which would appear to exclude the ISA.

In the alternative, Petitioner argues that "it is undisputed that 'hardware buffers' existed in the prior art to the '949 patent" and that "a POSITA would have understood a 'hardware buffer' to have been obvious based on Bauer and Svensson, regardless of how the term 'hardware buffer' is construed." Paper 35, 17 (citing Ex. 1026 ¶ 46). I find this argument untimely, as it was not raised in the Petition. *See* Paper 3, 26–27 (arguing only that the ISA "is . . . a separate hardware buffer"); *SAS Inst., Inc. v. Iancu*, 138 S. Ct. 1348, 1356 (2018) (explaining that the Board does not

have "license to depart from the petition"). But I also find the argument far too superficial to support a finding of obviousness, at least because Petitioner makes no effort to explain how or why a hardware buffer would have fit into the Bauer/Svensson combination.

The Bauer/Svensson combination is based on the idea that the ARM device can write directly to the DSP's ISA. But what role would a hardware buffer have in that system? Adding a USB controller, or the like, to the DSP to receive the data and then write it to the ISA for later copying to the DSP XRAM would result in the extra copy the inventors were trying to avoid. Thus, the modified combination would need to completely replace the ISA with a hardware buffer. Petitioner has neither made that argument, nor explained and supported why it would have been obvious to one of ordinary skill in the art at the time of the invention to do that.

For these reasons, I would find that Petitioner has not shown that the combination teaches or suggests the claimed "hardware buffer," and that Petitioner thus has not shown that claims 1–9 and 12 would have been unpatentable over the combination of Bauer and Svensson.

IV. CONCLUSION

I would construe "hardware buffer" to mean "a permanent, dedicated buffer that is distinct from system memory," which I find to be the only construction before us that is reasonable in view of the the intrinsic (and extrinsic) evidence. And, based on that construction, I would find that Petitioner has not proven claims 1–9 and 12 unpatentable.

For PETITIONER:

David Cavanaugh
Thomas Anderson
Joseph Haag
WILMER CUTLER PICKERING HALE & DORR LLP
david.cavanaugh@wilmerhale.com
tom.anderson@wilmerhale.com
joseph.haag@wilmerhale.com

For PATENT OWNER:

David Cochran
Matthew Johnson
Joseph Sauer
David Maiorana
Richard Graham
Joshua Nightingale
JONES DAY
dcochran@jonesday.com
mwjohnson@jonesday.com
jmsauer@jonesday.com
dmaiorana@jonesday.com
ragraham@jonesday.com
jrnightingale@jonesday.com