

**United States Court of Appeals
for the Federal Circuit**

INTEL CORPORATION,
Appellant

v.

QUALCOMM INCORPORATED,
Cross-Appellant

2020-1828, 2020-1867

Appeals from the United States Patent and Trademark Office, Patent Trial and Appeal Board in Nos. IPR2018-01334, IPR2018-01335, IPR2018-01336.

Decided: December 28, 2021

THOMAS SAUNDERS, Wilmer Cutler Pickering Hale and Dorr LLP, Washington, DC, argued for appellant. Also represented by DAVID LANGDON CAVANAUGH, CLAIRE HYUNGYO CHUNG; JOSEPH F. HAAG, Palo Alto, CA.

JENNIFER L. SWIZE, Jones Day, Washington, DC, argued for cross-appellant. Also represented by ROBERT BREETZ, DAVID B. COCHRAN, DAVID MICHAEL MAIORANA, JOSEPH M. SAUER, Cleveland, OH; KELLY HOLT, New York, NY; MATTHEW JOHNSON, JOSHUA R. NIGHTINGALE, Pittsburgh, PA; ISRAEL SASHA MAYERGOYZ, Chicago, IL.

Before PROST, TARANTO, and HUGHES, *Circuit Judges*.

TARANTO, *Circuit Judge*.

Qualcomm Inc. owns U.S. Patent No. 8,838,949, which addresses multi-processor systems in which software stored in non-volatile memory coupled to a first processor is to be used by a second processor. The patent describes and claims systems, methods, and apparatuses for efficiently retrieving an executable software image from the first processor's non-volatile memory and loading it for use by the second processor. Intel Corp. challenged all claims of the '949 patent as unpatentable for obviousness in three inter partes reviews (IPRs) before the Patent and Trademark Office. The Office's Patent Trial and Appeal Board consolidated the proceedings and issued a final written decision holding that Intel had proved unpatentable claims 10, 11, 13–15, and 18–23, but not claims 1–9, 12, 16, and 17. *Intel Corp. v. Qualcomm Inc.*, IPR2018-01334, 2020 WL 1286306, at *27 (P.T.A.B. Mar. 16, 2020) (Final Written Decision). Intel appeals.

We hold first that Intel has adequately demonstrated Article III standing to press this appeal. On the merits, we hold that in the decision before us, the Board failed to tie its construction of the phrase “hardware buffer” to the actual invention described in the specification. For that reason, we vacate the Board's decision as to claims 1–9 and 12 and remand for a new construction. As to claims 16 and 17, which are in means-plus-function format, we also vacate and remand. We conclude that the Board failed to determine for itself whether there is sufficient corresponding structure in the specification to support those claims and whether it can resolve the patentability challenges despite the (potential) indefiniteness of those claims.

I

A

The patent addresses a system with multiple processors, each of which must execute its own “boot code” to play its operational role in the system. Such code must be stored in non-volatile memory (*e.g.*, flash memory or read-only memory), since volatile memory is cleared when the device powers down; and the boot code generally must be transferred to its corresponding processor’s volatile memory in order to be executed by that processor. ’949 patent, col. 1, lines 39–41. In a multi-processor system, one possible design choice is to store the boot code for each processor in its own separate non-volatile memory. Another choice, to avoid the costs of multiple memories each adequate for such storage, is to store the boot code for one processor in the non-volatile memory of another processor, permitting elimination or shrinkage of the non-volatile memory of the first processor. *Id.*, col. 1, line 60, through col. 2, line 14.

The ’949 patent, titled “Direct Scatter Loading of Executable Software Image from a Primary Processor to One or More Secondary Processor in a Multi-Processor System,” assumes the latter, shared-storage choice. It addresses the problem, inherent in that choice, of loading the boot code for a “secondary” processor (into its volatile memory) from the non-volatile memory of a “primary” processor. *Id.*, col. 2, line 58, through col. 3, line 2. It uses a “direct scatter load” procedure to do so. “Scatter loading” refers to moving a “binary multi-segmented” software image into scattered parts (as opposed to one contiguous block) of the secondary processor’s “system memory” before executing it. *Id.*, col. 2, lines 14–22. The patent discloses a “direct” scatter loading process, through which the segments of the software image are transmitted “directly” from a “hardware buffer” to their final locations in the secondary processor’s “system memory.” *Id.*, col. 2, lines 58–63.

Claims 1 and 2 are representative for the claim-construction issue on appeal. They recite:

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.¹

¹ The indentation of the last two components of the “multi-processor system” (the “primary processor” and the “interface”) has been altered from the original to reflect the

2. 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.*

Id., col. 12, line 60, through col. 13, line 16 (emphases added).

Claim 16 is relevant to the means-plus-function issue on appeal. It recites:

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

fact that they are parts of the multi-processor system, not parts of the secondary processor.

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.

Id., col. 14, lines 17–37 (emphases added).

B

In 2017, Qualcomm initiated actions against Apple Inc. (not party to this appeal) in district court and at the International Trade Commission (ITC), alleging that Apple infringed the '949 patent (and other patents) by making, selling, and using iPhone models that incorporated baseband processors made by Intel. *See* J.A. 6233. Qualcomm's infringement theory identified the Intel baseband processors as the "secondary processors" of the '949 patent and pointed to "Intel design documents that describe how the Intel baseband processors are integrated into the iPhones." *See* J.A. 6262, 6264. In 2019, Qualcomm and Apple settled all litigation worldwide between the two companies, and Qualcomm agreed to license the patents to Apple for six years (and two additional years if Apple wished). J.A. 6084, 6210. Later in 2019, Apple acquired "the majority of Intel's smartphone modem business" (including its production of baseband processors), J.A. 6088, though Intel continues to supply pre-acquisition versions of its baseband processors to Apple and another customer, J.A. 6204.

Meanwhile, in 2018, Intel petitioned the Board for three inter partes reviews, each petition covering various claims of the '949 patent and together covering all. After consolidating the reviews, the Board issued a final written decision. 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 Board also ruled that Intel had failed to prove the unpatentability of the

remaining claims: (a) independent claim 1 and its dependent claims 2–9, plus claim 12 (depending on independent claim 10); and (b) independent claim 16 and its dependent claim 17. See *Final Written Decision*, at *27. Intel appeals those losses.

The Board’s determination upholding claims 1–9 and 12 turned on the construction of the claim phrase “hardware buffer.” No party put forward an explicit construction of that term until Qualcomm’s patent owner response, in which Qualcomm proposed to construe the phrase to mean “a buffer within a hardware transport mechanism that receives data sent from the primary processor to the secondary processor.” J.A. 4224. In reply, Intel argued that “‘hardware buffer’ should be given its ordinary meaning of ‘a buffer implemented in hardware.’” J.A. 4322. In sur-reply, Qualcomm defended its proposed construction but also advanced an alternative construction, “a buffer that is not allocated by the secondary processor.” J.A. 4397. Qualcomm explained: “In the ’949 patent, the hardware buffer is a *permanent* buffer within the hardware transport mechanism, in contrast to a *temporary* buffer in system memory that is allocated by the secondary processor at run time for this purpose.” *Id.* (cleaned up).

The Board rejected both Qualcomm’s original construction, which referred to “a hardware transport mechanism,” and Intel’s “ordinary meaning” construction, “a buffer implemented in hardware.” *Final Written Decision*, at *5–6. Instead, the Board agreed with Qualcomm that “the ’949 patent does differentiate disclosed loading techniques from known prior art techniques that use temporary buffers” and concluded that “the ‘hardware buffer’ limitations . . . ‘should not be read so broadly as to encompass’ the use of a temporary buffer.” *Id.* at *7 (quoting *SciMed Life Systems, Inc. v. Advanced Cardiovascular Systems, Inc.*, 242 F.3d 1337, 1343 (Fed. Cir. 2001)).

That construction, the Board concluded, was dispositive of the prior-art challenges to claims 1–9 and 12. Intel had argued that the required “hardware buffer” was taught by the “intermediate storage area” disclosed in its principal prior-art reference (Svensson, U.S. Patent No. 7,356,680). But because Svensson’s “intermediate storage area” was “reserved at runtime,” the Board determined, it was a temporary buffer and thus not a “hardware buffer.” *Id.* at *23–24.

As to claims 16 and 17, the Board ruled that Intel had failed to show unpatentability because Intel had not met its burden of identifying sufficient structure corresponding to two means-plus-function terms in independent claim 16. Claim 16 recites a “means for processing” and a “means for scatter loading.” *Id.* at *7–8. In IPR2018-01335, Intel’s petition identified those phrases as means-plus-function limitations and offered the same constructions for them that Qualcomm had offered in prior ITC litigation. In its institution decision, the Board wrote that it “ha[d] questions as to the sufficiency of [Intel’s] identified structures” for the two terms. J.A. 5160. But because the petition met the standard for institution on claims 10–15, the Board instituted review on claims 16 and 17 as well, suggesting that the parties “address the constructions of the mean-plus-function limitations in claim 16” and “the impact that a determination that the specification of the ’949 patent does not provide adequate corresponding structure for the recited functions should have on this proceeding and any final written decision.” J.A. 5161.

After institution, Qualcomm argued in its patent owner response that the means-plus-function terms in claim 16 “do not need to be construed in order for the Board to reach its Final Written Decision” because “[n]one of the arguments Qualcomm makes . . . to distinguish the prior art requires construction of these limitations,” J.A. 4226, but also argued that constructions Intel had proposed (the same ones Qualcomm had previously proposed in the ITC

proceeding) were sufficient. In reply, Intel did not defend as correct the structure it had identified in its petition; instead, it said, “Upon consideration of the Board’s articulated concerns, [Intel] agrees that the ’949 specification fails to disclose sufficient structure to perform the recited functions.” J.A. 4325. Intel also agreed with Qualcomm that the Board could address the patentability of claim 16 without construing the means-plus-function limitations. J.A. 4325–26. In sur-reply, Qualcomm pointed out that Intel had “change[d] its position” about the sufficiency of the structure in the specification; and Qualcomm argued, “[S]hould the Board maintain its position that the specification does not disclose corresponding structure for the functions, then this precludes the Board from finding that claim 16 is unpatentable.” Patent Owner Sur-Reply at 14–15, *Intel Corp. v. Qualcomm Inc.*, IPR2018-01334, Paper No. 25 (P.T.A.B. Nov. 8, 2019).

In its final written decision, the Board accepted Intel’s position that “the ’949 specification fails to disclose sufficient structure to perform the recited functions’ for two of the means-plus-function limitations.” *Final Written Decision*, at *7–8, *26. While acknowledging that both parties agreed that the Board should assess the patentability of claims 16 and 17 even if it concluded that there was insufficient corresponding structure for the functions, the Board concluded, “Because [Intel] 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 [Intel] has not shown, by a preponderance of the evidence, that [claims 16 and 17] are unpatentable under 35 U.S.C. § 103(a) as obvious.” *Id.* at *26.

Intel timely appealed. We have statutory jurisdiction under 28 U.S.C. § 1295(a)(4)(A) and 35 U.S.C. §§ 141(c), 319. Qualcomm moved to dismiss the appeal, arguing that Intel lacks Article III standing because Qualcomm had not sued or threatened to sue Intel for infringing the ’949 patent. Apple then moved to intervene in the appeal. We

denied both motions, directing Qualcomm and Intel to address standing in their merits briefs and permitting Apple to move to file an amicus brief. Apple did so, but its proposed brief only repeated its argument that it should be joined as a party, so we denied the motion.

II

We begin by addressing the jurisdictional standing question, and we arrive at the same conclusion this panel has reached today in *Intel Corp. v. Qualcomm Inc.*, No. 20-1664 (Fed. Cir. Dec. 28, 2021). Here, as there, Intel has engaged in activity that has already given rise to an infringement suit by Qualcomm. J.A. 6214–19 (ITC); J.A. 6239–43 (district court); *see also Grit Energy Sols., LLC v. Oren Techs., LLC*, 957 F.3d 1309, 1319 (Fed. Cir. 2020). And Qualcomm has not disputed that, in those proceedings, it identified an Intel product as the “secondary processor” of the ’949 patent. Intel Opening Br. 49–50; *see also* J.A. 6256–57, 6261–62, 6272–73. Thus, for the same reasons as in our companion case, “Intel’s risks transcend mere conjecture or hypothesis.” *Intel*, No. 20-1664, slip op. at 7. That is so even though Intel has only shown that it manufactures the claimed “secondary processor” of the ’949 patent’s claimed inventions, not all the components required by the claims, given the centrality of that component to the claims, the possibility of direct infringement suits based on product testing, and the possibility of indirect infringement suits based on at least inducement. *See JTEKT Corp. v. GKN Auto. Ltd.*, 898 F.3d 1217, 1221 (Fed. Cir. 2018) (“IPR petitioners need not concede infringement to establish standing to appeal.”). We therefore find that Intel has standing and proceed to consider the merits of its appeal.

III

A

We begin with the claim-construction issue presented respecting claims 1–9 and 12. Claim construction is ultimately a question of law, decided de novo on review, as are the intrinsic-evidence aspects of a claim-construction analysis. *See, e.g., Data Engine Techs. LLC v. Google LLC*, 10 F.4th 1375, 1380 (Fed. Cir. 2021). But we review any underlying fact findings about extrinsic evidence, such as extra-patent usage, for substantial-evidence support when the appeal comes from the Board. *Compare Monsanto Tech. LLC v. E.I. DuPont de Nemours & Co.*, 878 F.3d 1336, 1341 (Fed. Cir. 2018) (substantial-evidence review for Board factual findings) *with Teva Pharms. USA, Inc. v. Sandoz, Inc.*, 574 U.S. 318, 331–32 (2015) (clear-error review for district court factual findings related to claim construction). Here, though there was some expert testimony on the meaning of “hardware buffer,” the Board’s claim-construction reasoning involved only intrinsic evidence from the specification. *See Final Written Decision*, at *5–7. We therefore review it de novo.

In the inter partes review proceedings before us, brought before November 13, 2018, the Board’s claim-construction task was to determine the “broadest reasonable interpretation consistent with the specification.” *PPC Broadband, Inc. v. Corning Optical Commc’ns RF, LLC*, 815 F.3d 747, 751 (Fed. Cir. 2016); *see also Personalized Media Commc’ns, LLC v. Apple Inc.*, 952 F.3d 1336, 1340 n.2 (Fed. Cir. 2020) (noting the change in Board regulations beginning November 13, 2018). When applying that standard, the Board is not limited to choosing the single best interpretation when more than one is reasonable. But it requires that any adopted interpretation be reasonable in light of “general claim construction principles,” which govern in district court. *Microsoft Corp. v. Proxycorr, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015) (noting that the Board

may not “construe claims during IPR so broadly that its constructions are *unreasonable* under general claim construction principles”), *overruled in another respect, Aqua Prods., Inc. v. Matal*, 872 F.3d 1290 (Fed. Cir. 2017) (en banc); *see also In re Smith Int’l, Inc.*, 871 F.3d 1375, 1382 (Fed. Cir. 2017).

“[T]here is no magic formula or catechism for conducting claim construction.” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1324 (Fed. Cir. 2005) (en banc). Claim language and the specification (written description) are the dominant sources of interpretation, and prosecution history can matter to a lesser degree (though arguments based on prosecution history have not been made here). *Id.* at 1312–17; *see also World Class Tech. Corp. v. Ormco Corp.*, 769 F.3d 1120, 1123 (Fed. Cir. 2014); *Proxyconn*, 789 F.3d at 1298 (noting that even under the broadest-reasonable-interpretation standard, prosecution history can matter). Of central importance here, we have explained that “[u]ltimately, the interpretation to be given a term can only be determined and confirmed with a full understanding of what the inventors actually invented and intended to envelop with the claim.” *Phillips*, 415 F.3d at 1316 (quoting *Renishaw plc v. Marposs Societa’ per Azioni*, 158 F.3d 1243, 1250 (Fed. Cir. 1998)). Understanding the concrete substance to which the words used in the claim and the intrinsic evidence refer is necessary because, “[i]n reviewing the intrinsic record to construe the claims, we strive to capture the scope of the actual invention, rather than strictly limit the scope of claims to disclosed embodiments or allow the claim language to become divorced from what the specification conveys is the invention.” *Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 653 F.3d 1296, 1305 (Fed. Cir. 2011).

Here, it 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. And 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,” *Phillips*, 415 F.3d at 1316—an understanding that in some cases is usefully clarified by expert testimony (as long as that testimony is consistent with the intrinsic evidence), *see id.* at 1318. But, we conclude, the Board did not do enough to reach and articulate that understanding, and its claim construction is therefore wanting.

The phrase “hardware buffer” appears in claims 1, 2, 8, and 12 of the ’949 patent. 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. Nevertheless, we reach three conclusions from the claim language.

First, 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 Intel urges, or else the word “hardware” would be erased from the claims. That consequence, while not inevitably disqualifying a construction in every patent, is counter to an important principle of interpretation, for patent claims as for statutes: “It is highly disfavored to construe terms in a way that renders them void, meaningless, or superfluous.” *Wasica Finance GmbH v. Continental Automotive Systems, Inc.*, 853 F.3d 1272, 1288 n.10 (Fed. Cir. 2017); *see SimpleAir, Inc. v. Sony Ericsson Mobile Commc’ns AB*, 820 F.3d 419, 429 (Fed. Cir. 2016); *Stumbo v. Eastman Outdoors, Inc.*, 508 F.3d 1358, 1362 (Fed. Cir. 2007); *Phillips*, 415 F.3d at 1314; *Duncan v. Walker*, 533 U.S. 167, 174 (2001) (refusing to adopt a statutory construction that “would render the word ‘State’ insignificant, if not wholly superfluous”); *Salman Ranch Ltd. v. United States*, 573 F.3d 1362, 1374 (Fed. Cir. 2009) (“A cardinal rule of statutory construction is that courts should construe statutes so as to avoid rendering superfluous any statutory language.” (cleaned up)). Here, the striking fact that Qualcomm, in its claim language, did not just

say “buffer,” but instead said “hardware buffer,” provides a strong reason to avoid the disfavored result of rendering the word “hardware” superfluous.² *Second*, because claim 1 requires both a “system memory” and a “hardware buffer,” there must be some distinction between those two concepts. *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.”

Those conclusions from the claim language advance the claim-construction inquiry only so far. And they do not, on their own, provide a concrete basis for a clarifying definition of “hardware buffer.” What is needed, then, is an analysis of the specification 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.³ In this crucial respect, the Board fell short in its analysis here, and we think the Board is better positioned than we are to correct the deficiencies so as to arrive

² Though Intel asserts that the specification sometimes uses the word “hardware” in a redundant manner, Intel Opening Br. at 31, the uses Intel identifies are not instances of modifying the word “buffer,” and Intel does not persuasively show why such uses are actually redundant or, in any event, overcome the fact that the *claims* (in which relevant readers expect more precise, less discursive drafting than in the written description) recite “hardware buffer,” not just “buffer.”

³ To date, the parties have not made substantial claim-construction arguments about the prosecution history.

at the understanding needed to draw a confident conclusion about the proper claim construction. We identify some of the deficiencies without suggesting how, if at all, a proper construction will be substantively different from the limited, negative one (excluding “temporary” buffers) that the Board adopted.

Although 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. *Final Written Decision*, at *7. The Board cited three passages from the ’949 specification (column 2, lines 23–34; column 4, lines 43–47; and column 5, lines 31–35) that mention the use of a temporary buffer in prior art assertedly being improved by the invention here. But 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 those passages.

Other seemingly significant characteristics of the prior-art buffers (as described in the ’949 patent’s specification) appear unrelated to their temporal character. Discussing prior art, the specification says that the prior-art buffer “would be some place in system memory.” ’949 patent, col. 2, lines 31–34. That statement aligns with what the claim language already makes clear—Qualcomm’s “hardware buffer” is somehow different from “system memory”—but does not clarify what exactly the difference is. The specification also says that use of a referred-to prior-art buffer would require “copying the data into a temporary buffer in the modem processor RAM.” *Id.*, col. 5, lines 31–35. That statement advances the inquiry into the proper claim construction, because it aligns with 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.” *See id.*, col. 7, line 16 (“*Zero Copy*

Transport flow” (emphasis added)); *id.*, col. 7, lines 27–30 (“Thus, aspects of the present disclosure avoid extra memory copy operations, thereby improving performance (e.g., reducing the time required to boot secondary processors in a multi-processor system).”); *id.*, col. 9, lines 42–46 (“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.”).

As to what the seemingly important relationship between a “hardware buffer” and those characteristics is, we have before us no adequate explanation. The Board’s explanation does not clarify the contemplated concrete operations of a “hardware buffer” in this patent. It therefore fails to clarify how a “hardware buffer,” as contemplated in the specification, produces improved efficiency through “direct loading” and avoiding “extra memory copy operations.” *Id.*, col. 9, lines 37–46. Nor does the Board’s analysis mention the distinctions between “system memory” and “hardware buffer” that are drawn both in the claim language and in the specification.

What 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 (at least in the excerpts drawn to our attention), of what the intrinsic evidence shows the asserted advance to be and how, concretely, the “hardware buffer” relates to that advance. We do not exclude the possibility that the record should be expanded in order to arrive at an adequate understanding at the substantive level. The Board’s construction was based on what Qualcomm proposed only in its sur-reply, so the Board did not benefit from expert explanation of technical operations that might bear on the merits of that construction and produce the needed understanding.

The Board's construction was entirely a negative one—excluding “temporary” buffers. “Although there is no *per se* rule against negative constructions,” *Medicines Co. v. Mylan, Inc.*, 853 F.3d 1296, 1308 (Fed. Cir. 2017) (citation omitted), which in some cases can be enough to resolve the relevant dispute, the Board's construction in this case is inadequate. It is not clear what precisely constitutes a “temporary buffer” as recited in the Board's construction. *Compare* Intel Opening Br. at 40 (positing that a temporary buffer must be both “allocated or reserved at runtime” and “deallocated to be used for another purpose”), *with* Qualcomm Resp. Br. at 48 (arguing that a buffer that is “newly allocated each time the system is booted” is temporary). To resolve even that uncertainty requires the kind of additional, substantive understanding discussed above, which seems likely to support an affirmative construction in place of the Board's purely negative one.

Finally, we note two matters that would benefit from attention on remand. In a trial transcript from the Qualcomm-Apple litigation (a transcript that was before the Board here), a named inventor of the '949 patent testified in some detail about the difficulty of “design[ing] [the] system so the hardware could place the data right where it needed to be.” Transcript of Jury Trial Day 2 (Vol. 2B) at 222, *Intel Corp. v. Qualcomm Inc.*, IPR2018-01334, Exhibit No. 2004 (P.T.A.B. Dec. 12, 2019). That testimony may bear on the remand inquiry. Additionally, our cases sometimes speak of hardware and software implementations of computer functions, reflecting a distinction that appears in usage in the field. *See Tomita Techs. USA, LLC v. Nintendo Co.*, 681 F. App'x 967, 972 (Fed. Cir. 2017) (finding that a “software implementation” of an image processing function was not equivalent to a patent's “hardware implementation” of that function); *Texas Digital Systems, Inc. v. Telegenix, Inc.*, 308 F.3d 1193, 1214–16 (Fed. Cir. 2002) (recognizing the broad range encompassed by a construction as “any firmware, software and/or hardware” that

performs a certain function); *Nazomi Commc'ns, Inc. v. ARM Holdings, PLC*, 266 F. App'x 935, 936 (Fed. Cir. 2008) (considering a patent on a “hardware accelerator” that “significantly speed[s] up the processing of Java bytecodes over prior art systems that used software to perform the conversion from stack-based instructions to register-based instructions”). Neither 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.

The Board’s rejection of Intel’s challenge to those claims rested on the determination that the Svensson reference did not disclose the claimed “hardware buffer,” which in turn depended on the Board’s claim construction of that phrase. Because that construction is inadequate and must be reconsidered, we vacate the rejection of Intel’s challenges to claims 1–9 and 12, and we remand for reconsideration of claim construction as discussed in this opinion.

B

With respect to claims 16 and 17, there 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). In the IPR2018-01335 proceedings here, after Intel agreed with the Board’s suggestion in the institution decision that two of the means-plus-function terms in claim 16 were indefinite for lack of supporting structure, the Board concluded that Intel’s statement necessarily meant that Intel, as the petitioner, had not met its burden to demonstrate the unpatentability of those claims. *See Final Written Decision*, at *26. Intel challenges that merits conclusion. We hold that the Board’s conclusion was error and that a remand is required, because the Board did not decide for itself whether required structure is present in the specification or whether, even if it is not, the absence of

such structure precludes resolution of Intel's prior-art challenges.

Under 35 U.S.C. § 318, as construed in *SAS Institute, Inc. v. Iancu*, 138 S. Ct. 1348 (2018), the Board is obligated to “issue a final written decision with respect to the patentability of” every claim challenged by the petitioner. *Id.* at 1353. But 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. *See Samsung Electronics America, Inc. v. Prisia Eng'g Corp.*, 948 F.3d 1342, 1353 (Fed. Cir. 2020); *cf. Anniston Mfg. Co. v. Davis*, 301 U.S. 337, 355 (1937) (“We do not think that Congress was attempting to require the impossible.”). In such a case, the statutory estoppel provision of 35 U.S.C. § 315(e) does not apply, because the problem of indefiniteness is one of the patentee's own making, not attributable to the challenger. *Cochlear Bone Anchored Solutions AB v. Oticon Med. AB*, 958 F.3d 1348, 1359 (Fed. Cir. 2020) (citing *Samsung*, 948 F.3d at 1353 n.3). To avoid confusion going forward, the Board should, in IPRs where that principle applies, clearly state that the final written decision does not include a determination of patentability of any claim that falls within the impossibility category.

Importantly, it is not always impossible to adjudicate a prior-art challenge, one way or the other, just because some aspect of a claim renders the claim indefinite. *See Samsung*, 948 F.3d at 1355 (noting that indefiniteness “does not necessarily preclude the Board from addressing the patentability of the claims on section 102 and 103 grounds”). For example, the indefiniteness of one limitation may not preclude the Board from rejecting a challenge by finding that another limitation is missing from the argued prior art and its argued combinations and modifications. In the other direction, if a claim limitation requires alternative limitations A or B, and A is indefinite, but B is not, the Board

may well be able to determine that the argued prior art and its argued combinations or modifications cover the B option, thus satisfying the A or B limitation. *See Cochlear*, 958 F.3d at 1359–60. The indefiniteness of a limitation (here, a means-plus-function limitation) precludes a patentability determination only when the indefiniteness renders it logically impossible for the Board to reach such a decision.

The inter partes review scheme as a whole confirms the limited character of the impossibility qualifier to the SAS obligation. That scheme protects the interests not only of the petitioner in securing a determination on the patentability of a claim, but other interests as well. It protects interests of the patentee, as well as of the judicial system and the agency, in diminishing duplication of adjudication burdens and risks in specified respects. *See, e.g.*, 35 U.S.C. § 315(b) (timing limit on seeking IPR once district court action filed), § 315(e) (estoppel provision against other proceedings). And it reflects a public interest in the Board's answering the patentability questions, reflected in the statutory authorization for the Board to decide an inter partes review even after petitioners have withdrawn. 35 U.S.C. § 317.

The Board's treatment of claims 16 and 17 here was contrary to those principles. The Board simply accepted Intel's concession that the required structure for the means-plus-function claim was missing. The Board did not itself conclude that the prior-art analysis task was impossible, and it could not so conclude here without determining for itself that such structure was missing, a legal question that is part of claim construction. *See Noah Systems, Inc. v. Intuit Inc.*, 675 F.3d 1302, 1311 (Fed. Cir. 2012) ("Whether a claim complies with the definiteness requirement of 35 U.S.C. § 112 is a matter of claim construction . . ."). Nor did the Board decide whether it could resolve the patentability dispute even if claim 16 is indefinite. Contrary to Intel's argument in this court, Intel Opening

Br. at 45, there is no proper basis at present for not reaching the merits of the prior-art challenge to claims 16 and 17 or for declaring inapplicability of estoppel. At the same time, contrary to Qualcomm's argument in this court, Qualcomm Resp. Br. at 50–51, the Board could not properly enter a final written decision on the merits of the prior-art challenges with estoppel effect, as the Board did not assess those merits.

On remand, the Board must 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. *See Samsung*, 948 F.3d at 1358.

IV

For the foregoing reasons, we vacate 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 we remand for further proceedings consistent with this opinion.

The parties shall bear their own costs.

VACATED AND REMANDED