NOTE: This disposition is nonprecedential.

United States Court of Appeals for the Federal Circuit

QUALCOMM INCORPORATED,
Appellant

v.

INTEL CORPORATION,
Appellee

2023-1710

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

Decided: January 24, 2025

JENNIFER L. SWIZE, Jones Day, Washington, DC, argued for appellant. Also represented by DAVID B. COCHRAN, Cleveland, OH; JOHN MICHAEL GRAVES, Minneapolis, MN; ISRAEL SASHA MAYERGOYZ, Chicago, IL; JOSHUA R. NIGHTINGALE, Pittsburgh, PA.

THOMAS SAUNDERS, Wilmer Cutler Pickering Hale and Dorr LLP, Washington, DC, argued for appellee. Also represented by LAURA E. POWELL; RESHMA C. GOGINENI, New York, NY; JOSEPH F. HAAG, Palo Alto, CA.

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

TARANTO, *Circuit Judge*.

On remand following this court's decision in *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 801 (Fed. Cir. 2021) (*Intel 2021*), a majority of the Patent Trial and Appeal Board panel determined that claims 1–9 and 12 of U.S. Patent No. 8,838,949 were unpatentable for obviousness. *Intel Corp. v. Qualcomm Inc.*, No. IPR2018-01334, 2023 WL 2588220, at *1 (P.T.A.B. Mar. 21, 2023) (*Remand Decision*). This court's previous opinion summarizes the dispute, which centers on an issue of claim construction—the meaning of "hardware buffer." *See Intel 2021*, at 804–12. On remand from our 2021 decision, the Board majority changed its earlier construction, which excluded "a temporary buffer" from the scope of "hardware buffer" (and on that basis rejected Intel's unpatentability challenges), to a broader construction, under which a "hardware buffer" is "not limited to a 'permanent' buffer." *Remand Decision*, at *4, *15. Based on the new construction, the Board held all challenged claims unpatentable. *Id.* at *1, *30.

Qualcomm challenges the Board's new construction. It does not dispute that, if the Board's new construction is correct, the prior art discloses the claimed "hardware buffer." We now reject Qualcomm's arguments against the new construction and therefore affirm the Board's decision.

The parties accept that a buffer generally is a set of memory cells (and accompanying connection lines) characterized by its function of serving as a waystation for data being moved from one place to another. *See*, *e.g.*, Oral Arg. at 3:31–4:22 (Qualcomm counsel accepting such a characterization), 36:05–50 (same for Intel counsel); J.A. 3863–65 (Qualcomm expert, quoting dictionary). In the '949 patent, the claimed buffer is a set of memory cells that function as a waystation for a software image being moved from one

place to another in a multi-processor system. *See* '949 patent, col. 1, lines 24–33, col. 2, lines 58–66. More particularly, the software instructions are moved from one processor's non-volatile memory, via the buffer, to a second processor's "system memory" (RAM or comparable memory, *see*, *e.g.*, '949 patent, col. 2, lines 31–34) from which it will be used by the second processor. *See Intel 2021*, 21 F.4th at 804; Qualcomm's Opening Br. at 28–29; Intel's Response Br. at 35. The parties disagree about what is required for the buffer to be a "hardware buffer."

Intel contends, and the Board majority agreed, that a buffer is a hardware buffer in this patent if the memory cells used by the buffer are "physically separate" from the memory cells into which the software is loaded and from which the processor executes it ("system memory"). Qualcomm's Opening Br. at 33; *Remand Decision*, at *15. In Intel's view, this separation exists even if the cells for buffering use are allocated to that function only when the power is turned on (or runtime) and only for that power-on session. *See* Intel's Response Br. at 57 (arguing that the term "hardware buffer" includes a block of memory that is allocated upon boot-up for intermediate storage and is not "deallocated or used for a different purpose" during the session); *Remand Decision*, at *22–24. Qualcomm urges a narrower view, accepted by the Board before our 2021 decision and by the dissenter on remand. In that view, a buffer is a hardware buffer only if its memory cells are *never* used for system memory, rather than assigned to be used for that function (allocated) upon the turning on of power (or runtime). Qualcomm's Reply Br. at 1 (describing a hardware buffer as "a physical, fixed, always-there hardware structure"); Qualcomm's Opening Br. at 63–66 (arguing that the prior-art reference Svensson does not teach a hardware buffer because the component to which Intel maps "hardware buffer" "does not exist" until allocated at runtime).

Neither party has suggested that "hardware buffer" is a term of art with an established meaning in the relevant area. *See Intel 2021*, at 809 ("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."); Oral Arg. at 15:50–16:06 (Qualcomm counsel agreeing that "hardware buffer" is not such a term of art). Moreover, although Qualcomm asserts that Intel's view renders "hardware" superfluous, Qualcomm's Opening Br. at 5–6, 38, and Intel asserts the opposite, Intel's Response Br. at 35–36, neither party provides an illuminating or persuasive explanation of its position on that seemingly important issue. One reason may be that, in explaining the patent's words, neither party delves beneath the terminology of "buffer" and "memory" in their functional meaning to describe concretely the potential makeup and configuration of the physical components and steps used to perform the functions. *Cf.* PETER J. DENNING & CRAIG H. MARTELL, GREAT PRINCIPLES OF COMPUTING 59 (2015) ("The terminology of abstractions [common in computer science] often obscures *the principle of stuff*: the reality that computational actions are implemented as physical processes controlled by programs."). We turn to the specification and prosecution history to seek guidance.

Within the specification, the term "hardware buffer" appears three times. '949 patent, col. 2, lines 58–63 ("The system includes a secondary processor having a system memory and a *hardware buffer* for receiving at . . . 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.") (emphases added); *id.*, col. 9, lines 37–41 ("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.") (emphasis added). Those uses of the phrase do not indicate

in any way why either the Intel view or the Qualcomm view of the phrase is the right one.

We likewise receive no truly persuasive clarification from the specification's use of "allocate" in simply describing an aspect of prior art, which, among other things, is not tied to a distinguishing of "hardware buffer." '949 patent, col. 2, lines 25–28. The specification's use of "temporary" or "temporary buffer" gives Qualcomm something more supportive to cite, but the support, in the end, is not clear or especially strong. *See*, *e.g.*, *id.*, col. 2, lines 23–34. The specification does not identify what "temporary" means in a way that distinguishes Intel's view. The specification notes avoidance of use of a "temporary buffer" in noting what is missing in "one exemplary aspect" (*i.e.*, "an example, instance, or illustration") of the invention described in the '949 patent. *Id.*, col. 4, lines 43–47; *id.*, col. 4, lines 22–23. In the one place where "temporary buffer" and "hardware buffer" are used together, what is distinguished is "employing a temporary buffer for the entire image," contrasting "direct scatter load[ing of] the image segments." *Id.*, col. 9, lines 37–56; *see also*, *e.g.*, *id.*, col. 2, lines 22–34 (explaining how conventional techniques may involve loading "actual image data").

Importantly, the patent specification does not explain how *only* a hardware buffer under Qualcomm's proposed construction, in contrast to the broader class of buffers covered by Intel's proposed construction, would serve the efficiency goals that the specification attributes to the invention. The specification overwhelmingly credits the *process* taking place within the hardware buffer—not the hardware buffer itself—for the invention's benefits. *See*, *e.g.*, *id.*, col. 7, lines 17–30 ("Aspects of the present disclosure provide techniques for efficiently loading the executable software images from the primary processor's non-volatile memory to the secondary processor's volatile memory. . . . Thus, aspects of the present disclosure avoid extra memory copy operations, thereby improving

performance . . . ."); *id.*, col. 9, lines 43–46 ("Thus, conventional techniques . . . are bypassed in favor of a more efficient direct loading process."); *id.*, col. 10, lines 16–18 ("In this exemplary aspect, there is no CPU processing done on the actual data segments, thereby improving efficiency of the load process.") The specification does not indicate that a buffer qualifying under Intel's view could not carry out that process with comparable benefits. Nor does it tout benefits from being able to skip the step of runtime allocation of memory cells to the buffer function.

What remains is any insight the prosecution history provides. Prosecution history can illuminate a term's meaning even when the patentee has not explicitly made relevant disclaimers. *Personalized Media Communications, LLC v. Apple Inc.*, 952 F.3d 1336, 1340 (Fed. Cir. 2020) ("Accordingly, even where 'prosecution history statements do not rise to the level of unmistakable disavowal, they do inform the claim construction.'" (quoting *Shire Development, LLC v. Watson Pharmaceuticals, Inc.*, 787 F.3d 1359, 1366 (Fed. Cir. 2015))). Prosecution history can be telling about how a patentee expects a relevant artisan to understand a term, even if what a relevant artisan is expected to infer about a term's scope points in the direction of broadening rather than narrowing. *Honeywell Inc. v. Victor Co. of Japan*, 298 F.3d 1317, 1323–24 (Fed. Cir. 2002).

Here, the prosecution history is illuminating. The Patent and Trademark Office first rejected the patent as anticipated by the Svensson prior art. J.A. 1549. In doing so, the examiner mapped the claim phrase "hardware buffer" onto the intermediate storage area of Svensson—the same component of the same prior art Intel now relies on as teaching the hardware buffer limitation in its obviousness challenge to the patent. J.A. 1549; *Remand Decision*, at *8 (explaining that the examiner found a "hardware buffer" described by "the same [intermediate storage area] of Svensson" that Intel now asserts as prior art). Before us,

Qualcomm argues that Svensson's intermediate storage area does not meet the hardware buffer limitation. Qualcomm's Opening Br. at 63–66. But when Qualcomm responded to the examiner's initial rejection, it did not expressly or implicitly challenge the examiner's characterization of the intermediate storage area as a hardware buffer; it instead modified the claims to make clear that the *process* its buffer performed was different from that of Svensson's intermediate storage area. J.A. 1555; *see also* J.A. 1561–62 (distinguishing its claimed data-loading processes from Svensson's practices). Although silence on a point that arises during patent prosecution may not often be properly given significance, *3M Innovative Properties Co. v. Avery Dennison Corp.*, 350 F.3d 1365, 1373–74 (Fed. Cir. 2003), in this case, the silence is telling. Qualcomm did not even hint, let alone declare, that the Svensson intermediate storage area—undisputedly allocated at runtime—was outside its claimed "hardware buffer."

In this case, there is something to be said for each side's view of how to read the claim phrase "hardware buffer." We need not, however, decide which is the better meaning of the phrase. This case is governed by the broadest-reasonable-interpretation standard, *Intel 2021*, at 808–09, which recognizes the distinctive duty of applicants or patentees to respond to uncertainty of scope by making clarifying changes, *see Cuozzo Speed Technologies, LLC v. Lee*, 579 U.S. 261, 280–81 (2016); *In re Yamamoto*, 740 F.2d 1569, 1571 (Fed. Cir. 1984). We conclude that Qualcomm has not persuasively shown the Board majority's current claim construction to be unreasonable. We therefore affirm the Board's claim construction and its resulting unpatentability determinations.

AFFIRMED