

United States Court of Appeals for the Federal Circuit

INTEL CORPORATION,
Appellant

v.

PACT XPP SCHWEIZ AG,
Appellee

2022-1037

Appeal from the United States Patent and Trademark
Office, Patent Trial and Appeal Board in No. IPR2020-
00518.

Decided: March 13, 2023

ROBERT ALAN APPLEBY, Kirkland & Ellis LLP, New
York, NY, argued for appellant. Also represented by DIVA
R. HOLLIS, NATHAN S. MAMMEN, JOHN C. O'QUINN, Wash-
ington, DC.

SANFORD IAN WEISBURST, Quinn Emanuel Urquhart &
Sullivan, LLP, New York, NY, argued for appellee. Also
represented by NIMA HEFAZI, FREDERICK A. LORIG, Los An-
geles, CA; MARK YEH-KAI TUNG, Redwood Shores, CA.

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

PROST, *Circuit Judge*.

The Patent Trial and Appeal Board (“Board”) determined that Intel Corp. (“Intel”) failed to show that claim 5 of U.S. Patent No. 9,250,908 (“the ’908 patent”) was unpatentable as obvious in light of prior art references Kabemoto and Bauman.¹ *Intel Corp. v. PACT XPP Schweiz AG*, No. IPR2020-00518, Paper 34, 2021 WL 3503434 (P.T.A.B. Aug. 9, 2021) (“*Final Written Decision*”). We reverse and remand.

BACKGROUND

I

PACT XPP Schweiz AG (“PACT”) owns the ’908 patent, which relates to multiprocessor systems and how processors in those systems access data. Multiprocessor systems typically store data in several places: there’s a main memory, where all of a system’s data is stored, as well as various cache memories, where smaller pieces of that same data are stored. Cache memories are closer to the processors, allowing the processors quicker access to the data available in a given cache. And a system can use multiple cache levels, where a primary cache is closer to the processor but can store less data than a further-away secondary cache.

The use of multiple cache memories can pose problems for cache coherency, though. Different caches can have local copies of the same data, so inconsistencies may arise if one processor changes its local copy of the data and that change isn’t propagated to the other copies of that data. That’s why multiprocessor systems often require a mechanism to monitor and maintain cache coherency. One way to maintain cache coherency is by “snooping” along a

¹ U.S. Patent No. 5,890,217 (“Kabemoto”); U.S. Patent No. 5,680,571 (“Bauman”).

shared “bus.” *See* J.A. 1911; Kabemoto Fig. 3. Another way to maintain cache coherency is by using a global, segmented secondary cache. *See* Bauman Fig. 6. Both of these mechanisms use a shared entity between processors to detect changes between, and ultimately make changes to, local data copies. *See* Kabemoto col. 17 l. 27–col. 18 l. 6; Bauman col. 5 l. 55–col. 6 l. 40.

II

The ’908 patent claims a multiprocessor system. Intel petitioned for inter partes review of claims 4 and 5 of the ’908 patent. Claim 5 depends from independent claim 4. Before the Board’s Final Written Decision, PACT statutorily disclaimed claim 4. *See Final Written Decision*, 2021 WL 3503434, at *1 n.2, *4. But, to show that claim 5 was unpatentable, Intel still had to demonstrate that the prior art taught all limitations in claim 4 (in addition to the limitations added by claim 5) because claim 5 “includes all of the limitations of that underlying independent claim.” *Id.* at *4; *see, e.g., Vectra Fitness, Inc. v. TNWK Corp.*, 162 F.3d 1379, 1383 (Fed. Cir. 1998).

The claim language at issue in this appeal appears in underlying independent claim 4 and provides:

4. A system, the system comprising:
 - a processing system comprising
 - a plurality of processors; and
 - at least one separated cache not part [of] any processor;
 - ...
 - wherein the at least one separated cache comprises a separated cache segment for at least some of the plurality of processors; the system further comprising:

an interconnect system interconnecting each of the separated cache segments with each of the processors, each of the processors with neighboring processors, and each of the separated cache segments with neighboring separated cache segments; and

an arbiter, the arbiter controlling access of a processor to the interconnect system.


'908 patent claim 4 (emphasis added).

The claimed interconnect system requires three specific interconnections: (1) “each . . . separated cache segment[] with each . . . processor[]”; (2) “each . . . processor[] with neighboring processors”; and (3) “each . . . separated cache segment[] with neighboring separated cache segments.” *Id.* The third of these limitations is relevant to this appeal, and we refer to it as the segment-to-segment limitation.

III


Intel asserted that the prior art taught a multiprocessor system that used the separated cache and interconnect system as described in claim 4. Three annotated figures help illustrate Intel’s proposed combination: Kabemoto’s Figures 3 and 4 and Bauman’s Figure 6. The annotated versions of Kabemoto’s Figures 3 and 4 are shown below, cropped to focus on the portions relevant to this appeal. Across all three of these annotated figures, processors are shown in yellow, secondary caches in blue, and interconnection systems in gold. *See Final Written Decision*, 2021 WL 3503434, at *7–8; J.A. 932 ¶ 129.

FIG. 3




See *Final Written Decision*, 2021 WL 3503434, at *7 (Kabemoto Fig. 3) (annotations in original).

FIG. 4



See J.A. 932 (Kabemoto Fig. 4) (annotations in original).

The annotated version of Bauman's Figure 6 is shown below, also cropped to focus on the portion relevant to this appeal.



See *Final Written Decision*, 2021 WL 3503434, at *8 (Bauman Fig. 6) (annotations in original).

As is relevant to this appeal, we focus on the yellow processors in Kabemoto, the blue global, segmented secondary cache in Bauman, and the gold interconnection system in both Kabemoto and Bauman.

Intel contended that a person of ordinary skill in the art would combine Kabemoto and Bauman to teach all limitations in claim 4 by “replac[ing] Kabemoto’s secondary caches” with “Bauman’s segmented global [secondary cache],” which is a separated cache. *Final Written Decision*, 2021 WL 3503434, at *8 (cleaned up). A person of ordinary skill, Intel argued, would connect Bauman’s global, segmented secondary cache “to [Kabemoto’s] snoop bus 22 on the outside of [processor] element 14-1” to reach a system with the claimed separated cache and interconnect system. *Id.* (cleaned up).

PACT did not dispute that the combination of Kabemoto and Bauman taught each limitation of claim 4; PACT only argued that Intel failed to demonstrate a motivation to combine Kabemoto and Bauman. *See* J.A. 601. Nevertheless, the Board purported to “agree” with PACT that Intel failed to demonstrate that the prior art disclosed the segment-to-segment limitation. *Final Written Decision*, 2021 WL 3503434, at *7–8; *see* Oral Arg. at 17:20–29, No. 22-1037, https://oralarguments.cafc.uscourts.gov/default.aspx?fl=22-1037_12072022.mp3 (PACT counsel admitting “in all candor” that PACT “did not make this . . . argument” before the Board). The Board also concluded that Intel failed to show that a person of ordinary skill in the art would have been motivated to combine the teachings of Kabemoto and Bauman. *Final Written Decision*, 2021 WL 3503434, at *8–11. And because the Board determined that Intel failed to prove the obviousness of each limitation in claim 4, it upheld the patentability of claim 5. *See id.* at *17.

Intel appeals, and we have jurisdiction under 28 U.S.C. § 1295(a)(4)(A).

DISCUSSION

What the prior art discloses and whether a person of ordinary skill would have been motivated to combine prior art references are both fact questions that we review for substantial evidence. *PAR Pharm., Inc. v. TWI Pharms., Inc.*, 773 F.3d 1186, 1193 (Fed. Cir. 2014). “Substantial evidence is such relevant evidence as a reasonable mind might accept as adequate to support a conclusion.” *Novartis AG v. Torrent Pharms. Ltd.*, 853 F.3d 1316, 1324 (Fed. Cir. 2017) (cleaned up).


Intel argues that two of the Board’s conclusions about underlying independent claim 4 lack substantial evidence. First, Intel asserts that substantial evidence does not support the Board’s determination that the prior art fails to disclose the segment-to-segment limitation. And second,

Intel contends that substantial evidence does not support the Board's determination that there was no motivation to combine Kabemoto and Bauman. We agree on both counts.

I

Intel first asserts that Bauman's Figure 6 teaches the segment-to-segment limitation. Appellant's Br. 46–48. We agree and conclude that the Board's determination that the segment-to-segment limitation wasn't in the prior art lacks substantial evidence.

Intel used a color-coded version of Bauman's Figure 6 to illustrate this teaching, the relevant portion of which is reproduced below:



See *Final Written Decision*, 2021 WL 3503434, at *8 (annotations in original).

As Bauman explains, its Figure 6 “illustrates the data path between . . . processors, the second-level cache, and the memory.” Bauman col. 11 ll. 18–20. The second-level cache consists of “Segments 0–3,” *id.* at col. 11 ll. 20–23, each of which is labeled “SEG” and identified in blue in the color-coded figure above, *Final Written Decision*, 2021 WL 3503434, at *8. The data path—i.e., the interconnection

system linking the processors and second-level cache—is shown in gold. *Id.*

The Board, in finding that Intel failed to show the segment-to-segment limitation in the asserted prior art, faulted Intel for failing to explain how Kabemoto’s “snoop bus 22” connected each cache segment to its neighboring segment. *See id.* at *7–8. According to the Board, this was a fatal flaw because, per the Board’s understanding, Intel relied on Kabemoto’s snoop bus 22 to disclose all three limitations of the claimed “interconnect system” in the ’908 patent. *See id.* But Intel relied on Bauman to teach a separated and segmented cache in its petition and throughout the proceedings. *E.g.*, J.A. 216–17 (citing J.A. 933–35 ¶¶ 133–34).

Bauman’s Figure 6 teaches—if not plainly illustrates—the segment-to-segment limitation of the claimed interconnect system: each blue cache segment is connected to its neighboring blue cache segments via the gold data path. That was Intel and PACT’s understanding at the Board, and that’s our understanding from the record on appeal. We can discern no other reasonable understanding of this figure.² Accordingly, the Board’s determination that the

² It could be that the Board reached the opposite conclusion based on claim construction of “interconnect system,” but neither Intel nor PACT raised such an explanation in their briefing. Although the Board contended that it “need not reach the [claim-construction] issue of whether” the three “interconnect system” limitations were limited to “direct” connections, the only way the Board’s segment-to-segment conclusion makes sense is if the Board did, in fact, construe “interconnect system” to require direct connections and exclude indirect connections. *See Final Written Decision*, 2021 WL 3503434, at *3.

The Board’s reasoning as to why Intel failed to show the segment-to-segment limitation appears to reject Intel’s

prior art did not teach the segment-to-segment limitation lacks substantial evidence, and we reverse it.

II

Intel also argues that the Board’s rejection of its “known-technique” rationale for a motivation to combine lacks substantial evidence.³ Appellant’s Br. 57–63 (citing *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398 (2007)); *see also Intel Corp. v. Qualcomm Inc.*, 21 F.4th 784, 797 (Fed. Cir. 2021) (determining that the Board’s reasons for finding a lack of motivation to combine were not supported by substantial evidence “[u]nder applicable legal principles”). We agree and reverse the Board’s contrary finding.

A

The motivation-to-combine analysis is a flexible one. “[A]ny need or problem known in the field of endeavor at the time of invention and addressed by the patent can provide a reason for combining the elements in the manner

proffered construction that the “interconnect system” limitations can be shown via direct or indirect connections. Even if the Board is right that Intel’s proposed combination relied on Kabemoto’s “snoop bus 22” to teach the segment-to-segment limitation, Kabemoto’s snoop bus 22 does so—just with indirect connections. Should there be only one connection from the whole of Bauman’s global secondary cache to Kabemoto’s snoop bus 22, *see id.* at *7–8, each of the segments in that secondary cache would still be connected to each other through their shared singular connection to the snoop bus.

³ Intel offered three motivations to combine Kabemoto and Bauman, one of which was premised on a known-technique rationale and all of which the Board rejected. *See Final Written Decision*, 2021 WL 3503434, at *8–11. We discuss only the known-technique rationale because it is sufficient for reversal.

claimed.” *KSR*, 550 U.S. at 420 (emphasis added). And “[a] person of ordinary skill is also a person of ordinary creativity, not an automaton.” *Id.* at 421. So, “in many cases[,] a person of ordinary skill will be able to fit the teachings of multiple patents together like pieces of a puzzle.” *Id.* at 420. That’s why the motivation-to-combine analysis “need not seek out precise teachings directed to the specific subject matter of the challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ.” *Id.* at 418.

Additionally, “universal” motivations known in a particular field to improve technology provide “a motivation to combine prior art references *even absent any hint of suggestion* in the references themselves.” *Intel*, 21 F.4th at 797–99 (cleaned up) (emphasis in original) (determining that the Board’s rejection of “increasing energy efficiency,” a “generic concern” in electronics, as a motivation to combine lacked substantial evidence (cleaned up)).

Similarly, “if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *KSR*, 550 U.S. at 417. This is the so-called “known-technique” rationale. And if there’s a known technique to address a known problem using “prior art elements according to their established functions,” then there is a motivation to combine. *Intel*, 21 F.4th at 799–800. And we specify *address* a known problem because “[i]t’s not necessary to show that a combination is the *best* option, only that it be a *suitable* option.” *Id.* at 800 (cleaned up) (emphasis in original).

Assessing whether claimed subject matter involves the “application of a known technique” will “[o]ften” require “a court to look to interrelated teachings of multiple patents; the effects of demands known to the design community or present in the marketplace; and the background knowledge

possessed by a person having ordinary skill in the art.” *KSR*, 550 U.S. at 417–18.

B

Before the Board, Intel asserted that a person of ordinary skill would have been motivated to combine Kabemoto and Bauman because they “relate to the same field of multiprocessor . . . systems” and “address the same problem: maintaining cache coherency.” *Intel Corp. v. PACT XPP Schweiz AG*, No. IPR2020-00518, Paper 1, at 48 (P.T.A.B. Feb. 7, 2020) (“*Petition*”); see *Final Written Decision*, 2021 WL 3503434, at *11 (citing *Petition*, at 48). So, Intel reasoned, a person of ordinary skill “would have naturally turned to Bauman’s segmented [global secondary] cache to use . . . in Kabemoto” since Bauman’s separated cache was known to address the same cache-coherency issue that Kabemoto also sought to address, just through a different mechanism—a shared snoop bus. *Petition*, at 48–49.

The Board rejected Intel’s known-technique rationale. The Board stated that “[i]f . . . Kabemoto already addresses [the] problem [of cache coherency] through the use of a known technique similar to that of Bauman’s, [it] fail[ed] to see why one of ordinary skill in the art would regard Bauman’s technique as an obvious improvement to Kabemoto.” *Final Written Decision*, 2021 WL 3503434, at *11. But the Board’s reasoning belies its conclusion. That Kabemoto and Bauman address the same problem and that Bauman’s cache was a known way to address that problem is precisely the reason that there’s a motivation to combine under *KSR* and our precedent.

There is a motivation to combine when a known technique “has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way,” *KSR*, 550 U.S. at 417, using the “prior art elements according to their established functions,” *Intel*, 21 F.4th at 799–800. And here, there’s no dispute that using a global, segmented

secondary cache “has been used to improve” cache coherency in multiprocessor systems—as in Bauman—and a person of ordinary skill would “recognize that” such a cache “would improve similar” multiprocessor systems—like the one in Kabemoto—by addressing that same cache coherency problem. *See KSR*, 550 U.S. at 417. Bauman itself explains that its global, segmented secondary cache is one of two “primary mechanisms” by which its claimed system “accomplishe[s]” cache coherency. Bauman col. 5 ll. 55–59. There’s accordingly also no dispute that such a combination would constitute a use of Bauman’s secondary cache “according to [its] established function[.]” *Intel*, 21 F.4th at 799–800.

And contrary to the Board’s suggestion, Intel never had to show that replacing Kabemoto’s secondary cache with Bauman’s secondary cache was an “improvement” in a categorical sense. *See Final Written Decision*, 2021 WL 3503434, at *11. Intel just had to show that Bauman’s secondary cache was a “suitable option” to replace Kabemoto’s secondary cache. *Intel*, 21 F.4th at 800 (emphasis omitted).

It’s enough for Intel to show that there was a known problem of cache coherency in the art, that Bauman’s secondary cache helped address that issue, and that combining the teachings of Kabemoto and Bauman wasn’t beyond the skill of an ordinary artisan. Nothing more is required to show a motivation to combine under *KSR*, so we reverse the Board’s finding to the contrary.

III

Although we reverse the two factual findings discussed above, we must remand for the Board to address any remaining dispute about the patentability of claim 5.

Claim 5 recites: “The bus system of claim 4 where the arbiter is operable to allow processor access in chronological sequence.” ’908 patent claim 5. Intel relied on prior art

Chaney to teach this limitation.⁴ See J.A. 221–26, 259–64. PACT did not dispute that Chaney taught this added “chronological sequence” limitation. See *Intel Corp. v. PACT XPP Schweiz AG*, No. IPR2020-00518, Paper 21, at 32–33, 42–43 (P.T.A.B. Nov. 16, 2020) (PACT Response); *Intel Corp. v. PACT XPP Schweiz AG*, No. IPR2020-00518, Paper 26, at 18, 28 (P.T.A.B. Feb. 9, 2021) (Intel Reply). But because the Board failed to analyze the invention of claim 5 as a whole, we must remand.

CONCLUSION

We have considered PACT’s remaining arguments and find them unpersuasive. For the foregoing reasons, we reverse the Board’s findings related to the above-discussed limitations (which appear in claim 5 by virtue of its dependence from claim 4) and remand for further proceedings consistent with this opinion.

REVERSED AND REMANDED

COSTS

No costs.

⁴ U.S. Patent No. 5,930,822 (“Chaney”).