

NOTE: This disposition is nonprecedential.

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

MARK H. SANDSTROM,
Appellant

v.

MICROSOFT CORPORATION,
Appellee

2024-1040, 2024-1179

Appeals from the United States Patent and Trade-
mark Office, Patent Trial and Appeal Board in Nos.
IPR2022-00527, IPR2022-00528.

Decided: January 7, 2025

MARK H. SANDSTROM, Alexandria, VA, pro se.

LAUREN ANN DEGNAN, Fish & Richardson P.C., Wash-
ington, DC, for appellee. Also represented by BENJAMIN
JOSEPH CHRISTOFF, CHRISTOPHER DRYER, WALTER KARL
RENNER.

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

PER CURIAM.

Appellant Mark Sandstrom is the named inventor, and now the assignee, of U.S. Patent No. 9,632,833 and its parent, No. 9,424,090, which describe and claim related methods and systems for assigning processor cores among applications to improve processing efficiency. Microsoft Corporation petitioned the Patent and Trademark Office (PTO) to institute inter partes reviews of multiple claims of the two patents, under 35 U.S.C. §§ 311–19, alleging unpatentability for obviousness. The PTO’s Patent Trial and Appeal Board, acting for the PTO Director, instituted the requested reviews, and the Board, after conducting the reviews, concluded that all the challenged claims were unpatentable for obviousness. *Microsoft Corp. v. ThroughPuter, Inc.*, No. IPR2022-00527, 2023 WL 6115987 (P.T.A.B. Sept. 18, 2023) (*’833 Decision*); *Microsoft Corp. v. ThroughPuter, Inc.*, No. IPR2022-00528, 2023 WL 6122631 (P.T.A.B. Sept. 18, 2023) (*’090 Decision*). ThroughPuter, Inc., the assignee at the time, appealed, and Mr. Sandstrom, upon thereafter becoming the assignee, was substituted as appellant. We now affirm.

I

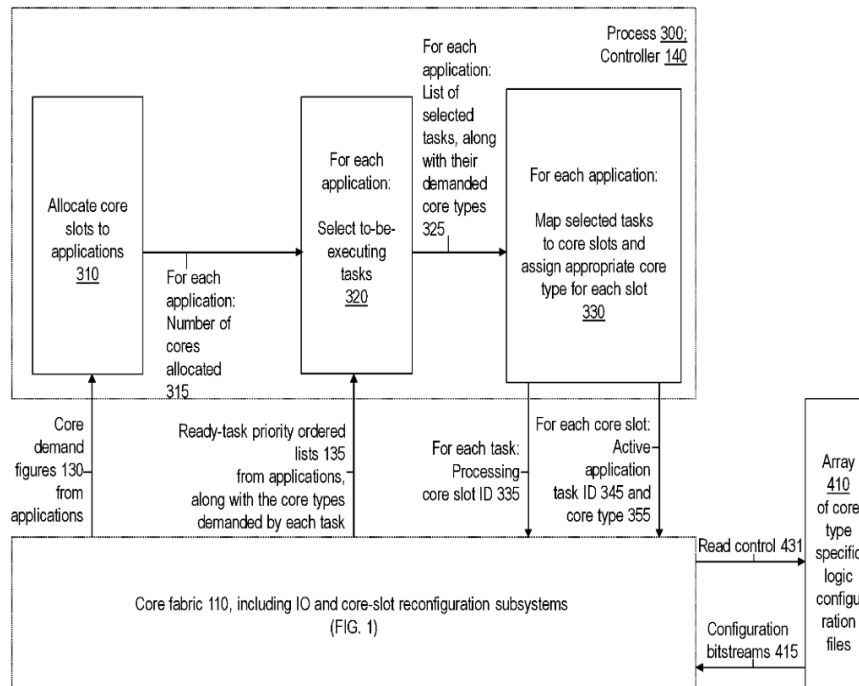
A

The ’833 patent, titled “Scheduling Application Instances to Processor Cores over Consecutive Allocation Periods Based on Application Requirements,” was issued on April 25, 2017, and the ’090 patent, its parent, titled “Scheduling Tasks to Configurable Processing Cores Based on Task Requirements and Specification,” was issued on August 23, 2016. The patents share a specification, so we cite only to the ’833 specification for simplicity.

The patents recite as a background fact that a computer may be assigned different types of tasks (or jobs) and that different processors may be best suited for

carrying out different assignments. '833 patent, col. 1, line 43, through col. 2, line 3. Including different processors (or processor cores) in the computer and matching task types to processor types, the patents add, can result in savings of time, energy, and other costs. *See id.*, col. 2, lines 23–28. The patents describe and claim methods and systems for doing just that. *Id.*, col. 1, lines 38–41.

The patents assert, however, that it is difficult to assign multiple tasks to multiple processors in a way that maximizes efficiency across the processors. *Id.*, col. 2, lines 4–17 (explaining that it is “infeasible to predict” at “a given instance of processing,” the “optimal type of core for any given processor instance,” which may lead to processing of jobs by “suboptimal types of processing cores”). The patents propose a solution. *Id.*, col. 4, lines 1–9 (explaining that the claimed invention “provides a . . . processor system for executing . . . tasks of differing types on their assigned cores of matching types”). Figure 3 shows an embodiment:



In the embodiment shown, at each time period, “referred to as a core allocation period (CAP),” *id.*, col. 10, lines 61–63, the system first allocates its processor cores among applications (*e.g.*, “software programs,” *id.*, col. 10, line 50) based on the applications’ respective core demands (*e.g.*, which core types and how many it can use, *id.*, col. 9, lines 4–7, 24–28) and entitlements (*e.g.*, how many it is entitled to use, *id.*, col. 13, lines 4–7). *Id.*, col. 10, lines 53–59. The system next selects the highest priority instances of the applications to be executed in the next CAP. *Id.*, col. 10, line 66, through col. 11, line 2. Finally, the system maps selected instances to best-suited processor cores. *Id.*, col. 11, lines 2–4.

For the ’833 patent, claim 15 is representative:

15. A method for assigning instances of software programs to an array of processor cores comprising:

for each of a series of successive core allocation periods (CAPs), selecting, from a group of executable instances of a set of software programs, a subset of the executable instances, referred to as selected instances, for execution on the cores of the array for an upcoming CAP, wherein the selection of the selected instances is based, at least in part, on a respective capacity demand indication of each of the set of software programs, with said indication of a given program (a) being based at least in part on a number of its executable instances that presently have input data available for processing and (b) indicating processor core types demanded by its executable instances;

assigning each of the selected instances for execution on a processor core in the array of processor cores based, at least in part, on matching the respective demanded processor core types associated

with the selected instances with types of processor cores available for assignment; and

executing the selected instances on their assigned cores over the next CAP, at least in part, to process the input data.

Id., col. 21, line 60, through col. 22, line 16.

For the '090 patent, claim 1 is representative:

1. A method for assigning a set of processing tasks to an array of processor cores of configurable types, the method comprising:

executing time variable subsets of the processing tasks of differing types on their assigned processor cores of matching types, wherein the matching type for the assigned processor core for a given processing task of the set corresponds to a type of a processor core demanded by the given processing task;

for each of a series of core allocation periods (CAPs), selecting, from the set of processing tasks, specific tasks, referred to as selected tasks, for execution on the processor cores for a next CAP at least in part based on core capacity demand expressions associated with the processing tasks;

assigning the selected tasks for execution on the processor cores for the next CAP in a manner to maximize, within the array, a number of processor cores whose assigned processing tasks for a present and the next CAP demands the same type of processor core; and

configuring the array such that the type of any given processor core in the array matches the type of processing task assigned for execution on the given processor core for the next CAP.

'090 patent, col. 19, line 59, through col. 20, line 15.

B

In March 2021, then-assignee ThroughPuter sued Microsoft in district court, alleging infringement of the '833 and '090 patents, along with other patents not at issue here. In February 2022, Microsoft petitioned for institution of inter partes reviews (IPRs) of claims 10, 13–15, 17–22, 24, 25, 27, 29, 32–34, and 37 of the '833 patent and claims 1–7 and 9–15 of the '090 patent. The Board instituted the IPRs in September 2022—IPR2022-00527 for the '833 patent and IPR2022-00528 for the '090 patent. The Board issued final written decisions for both patents on September 18, 2023.

1

In its '833 *Decision*, the Board determined that all the challenged claims were unpatentable for obviousness under 35 U.S.C. § 103.¹ The Board first determined that a relevant artisan would have been motivated to combine, and had a reasonable expectation of success in combining, the teachings of two pieces of prior art—Agrawal² and

¹ The '833 and '090 patents claim priority to before March 16, 2013, so the applicable version of § 103 here is the one that pre-dates the Leahy-Smith America Invents Act (AIA), Pub. L. No. 112-29, 125 Stat. 284 (2011). *See* '833 *Decision*, at *3 n.1; '090 *Decision*, at *2 n.1. Because the AIA-made changes do not alter the obviousness analysis in these cases, we cite § 103 without including a date.

² Kunal Agrawal et al., “Adaptive Scheduling with Parallelism Feedback,” Proceedings of the 2006 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2006. J.A. 856–65.

Brent³—to arrive at claims 15, 17, 18, 34, and 37. *’833 Decision*, at *5–18. The Board found that Agrawal discloses a “two-level” scheduler in which a “job scheduler [] allots processors to jobs” and a “task scheduler [] schedules the tasks belonging to a given job onto the allotted processors.” *Id.* at *5 (quoting J.A. 856). The Board also found that Agrawal teaches the technique of “instantaneous parallelism,” in which the job scheduler considers “the number of processors the job can effectively use at the current moment” to decide how to allocate processors to jobs for the upcoming period. *Id.* at *6 (quoting J.A. 857). It found that Brent teaches a system with several different types of processing cores in which “each task may be assigned a preferred core type for optimal execution.” *Id.* (citing J.A. 870 ¶¶ 3, 8; J.A. 871 ¶ 18). The Board determined that a relevant artisan would have had the motivation and ability to combine Agrawal’s teachings with Brent’s disclosure of matching processor core types to task types, *id.* at *11 n.12 (citing J.A. 759–60 ¶ 68 (Microsoft’s expert declaration)), leading to the conclusion that the subject matter of claims 15, 17, 18, 34, and 37 would have been obvious to a relevant artisan, *id.* at *18.

The Board then considered the combination of Agrawal, Brent, and a third piece of prior art, Feitelson⁴—a combination covering all the remaining challenged claims. *Id.* The Board found that Feitelson teaches different approaches to “scheduling on parallel computer systems shared by a number of users.” *Id.* (citing J.A. 874). The Board, citing suggestions in Feitel-

³ U.S. Patent Application Publication No. 2010/0131955 A1. J.A. 866–73.

⁴ Dror G. Feitelson, “Job Scheduling in Multiprogrammed Parallel Systems,” IBM Research Report RC 19790 (87657), 2nd Rev., August 1997. J.A. 874–1041.

son, agreed with Microsoft that a relevant artisan would have found it obvious to combine Feitelson's job-scheduling techniques with the teachings of Agrawal and Brent. *See id.* at *21 (explaining that a relevant artisan "would have found it obvious to apply Feitelson's suggestion" to an Agrawal-Brent combination); *id.* at *26 (agreeing with Microsoft that a relevant artisan would use a suggestion in Feitelson to "avoid[] wasted or unutilized resources" (citing J.A. 2477 ¶ 26)). The Board ultimately held that the subject matter of claims 10, 13, 14, 19–22, 24, 25, 27, 29, 32, and 33 would have been obvious to a relevant artisan based on the combination of Agrawal, Brent, and Feitelson. *Id.* at *30.

2

In its '090 *Decision*, the Board similarly determined that all the challenged claims were unpatentable for obviousness. *Id.* at *19. The Board first found that a relevant artisan would have had the motivation and ability to combine Agrawal with another piece of prior art, Chen,⁵ to arrive at the methods of claims 1–5 and 7. *Id.* at *4–13. The Board found that Chen discloses an operating system, which it models as a rectangular grid of configurable logic blocks (CLBs) (*e.g.*, processors), where processes are assigned to CLBs within the grid. *Id.* at *4–5 (citing J.A. 3321–22). In particular, the Board found that Chen teaches a scheduling algorithm to maximize processing efficiency by scheduling processes to CLBs in the grid to achieve "parallel execution of multiple processes." *Id.* at *5 (citing J.A. 3322–23). The Board made the

⁵ G. Chen et al., "Configuration-Sensitive Process Scheduling for FPGA-Based Computing Platforms," Proceedings Design, Automation and Test in Europe Conference and Exhibition, Vol. 1, pp. 486–93 (2004). J.A. 3321–23.

same findings regarding Agrawal that it did in its '833 *Decision*. *Id.* at *5–6; *see supra* Section I.B.1. The Board found that a relevant artisan “would have been motivated to combine Chen and Agrawal, with a reasonable expectation of success,” *id.* at *9, and that the subject matter of claims 1–5 and 7 would have been obvious to a relevant artisan based on that combination, *id.* at *13.

The Board then held that claims 3 and 4 would have been obvious under the combination of Chen, Agrawal, and a further piece of prior art, Compton.⁶ *Id.* at *13–14. The Board found that Compton “presents a survey of current research in hardware and software systems for reconfigurable computing,” including systems that “customiz[e] how the hardware is used” so that “the same hardware” overall can “execute different applications.” *Id.* at *13 (quoting IPR2022-00528, Ex. 1006, at 171, 173–74). The Board agreed with Microsoft that a relevant artisan would have combined Compton’s teachings on “reconfigurable devices,” which cover Field-Programmable Gate Arrays (FGPAs) like those taught by Chen, with the systems disclosed by Chen and Agrawal, *id.* at *13, and that “the limitations of each of claims 3 and 4 are rendered obvious by the combination of Chen, Agrawal, and Compton,” *id.* at *14.

Finally, the Board turned to Microsoft’s contention that claims 1, 2, 5–7, and 9–15 are unpatentable for obviousness in light of Chen, Agrawal, and Brent. *Id.* at *14–19. The Board made the same findings regarding Brent that it did in the '833 *Decision*. *Id.* at *14, *see supra* Section I.B.1. Then, the Board found that that a relevant

⁶ Katherine Compton & Scott Hauck, “Reconfigurable Computing: A Survey of Systems and Software,” *ACM Computing Surveys*, Vol. 34, No. 2, June 2002, pp. 171–210. IPR2022-00528, Ex. 1006.

artisan implementing a Chen-Agrawal system “would have naturally looked to Brent for an example of a system where tasks include explicit indications regarding the best-fitting or more suitable region or core type for the respective task.” *Id.* at *15 (citing J.A. 3280–81 ¶ 155). The Board ultimately held that the combination of Chen, Agrawal, and Brent would have rendered claims 1, 2, 5–7, and 9–15 unpatentable for obviousness. *Id.* at *19.

II

We have jurisdiction to review the two Board decisions under 28 U.S.C. § 1295(a)(4)(A). The obviousness determination by the Board rests on a familiar basis—that “a skilled artisan not only *could have made* but *would have been motivated to make* the combinations or modifications of prior art to arrive at the claimed invention.” *Belden Inc. v. Berk-Tek LLC*, 805 F.3d 1064, 1073 (Fed. Cir. 2015) (citing *InTouch Technologies, Inc. v. VGO Communications, Inc.*, 751 F.3d 1327, 1352 (Fed. Cir. 2014)). We review the Board’s ultimate determination of obviousness de novo but the underlying factual determinations for substantial-evidence support. *Personal Web Technologies, LLC v. Apple, Inc.*, 848 F.3d 987, 991 (Fed. Cir. 2017) (citing *In re Gartside*, 203 F.3d 1305, 1312 (Fed. Cir. 2000)). On the factual components of the inquiry, we “ask[] ‘whether a reasonable fact finder could have arrived at the agency’s decision,’” which “requires examination of the ‘record as a whole, taking into account evidence that both justifies and detracts from an agency’s decision.”” *Intelligent Bio-Systems, Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1366 (Fed. Cir. 2016) (quoting *Gartside*, 203 F.3d at 1312). When the Board’s claim constructions are challenged, the overall decision about the proper construction is a legal conclusion, reviewed de novo on appeal, but if there are underlying factual findings, they are reviewed for substantial-evidence support. *See St. Jude Medical, LLC v. Snyders Heart Valve LLC*, 977 F.3d 1232, 1238 (Fed. Cir. 2020).

On appeal, Mr. Sandstrom argues that the Board's decisions contain several prejudicial errors, so they must be set aside. We reject Mr. Sandstrom's challenges.

A

Mr. Sandstrom challenges the Board's finding that a relevant artisan would have been motivated to combine Agrawal and Brent in the '833 *Decision* or to combine Chen and Agrawal in the '090 *Decision*. Sandstrom Informal Br. at 19–23. He argues that there is no evidence of any real-world practical combination of Agrawal and Brent, *id.* at 21–22, and that “the stated goals of Chen and Agrawal would not have seemed compatible” to a relevant artisan, *id.* at 19. We disagree.

The Board explained its findings that a relevant artisan would have been motivated to make the relevant combinations, and in doing so it cited bases that we readily conclude make the findings reasonable in light of any contrary evidence cited to us. In its '833 *Decision*, the Board explained that a relevant artisan would have been motivated to combine Agrawal and Brent because both are “directed to systems for scheduling tasks in a parallel processing environment.” '833 *Decision*, at *11 n.12 (citing J.A. 758 ¶ 66 (Microsoft's expert declaration); J.A. 856–58 (Agrawal); J.A. 870–72 ¶¶ 13, 20, 26 (Brent)). The Board found that the combination would have been motivated by a desire “to achieve predictable benefits, including an improved ability to optimize task execution on a processor whose instruction set matches the instruction types provided in a task.” *Id.* (citing J.A. 758–59 ¶ 67 (Microsoft's expert declaration); J.A. 866, Abstract; J.A. 871 ¶¶ 16, 18, 19 (Brent)). To the extent Mr. Sandstrom suggests that the obviousness showing requires identification of an existing real-world combination, he cites, and we know of, no persuasive authority for such a requirement, which would run counter to the “would have been obvious” language of § 103.

In its *'090 Decision*, the Board made comparable findings, supported by substantial evidence, that a relevant artisan “would have been motivated to combine Chen and Agrawal, with a reasonable expectation of success.” *'090 Decision*, at *9. The Board found, citing evidence, that a relevant artisan “would have recognized that, like Chen, Agrawal is directed to a space-sharing system for scheduling tasks to hardware by means of its greedy task scheduler.” *Id.* (citing J.A. 3263–64 ¶ 103 (Microsoft’s expert declaration); J.A. 858 (Agrawal)). The Board also determined, citing evidence, that such an artisan “would have been motivated to modify Chen’s method with Agrawal’s use of scheduling quanta [*i.e.*, making scheduling assignments at certain time intervals] for making favorable choices for space allocation across the array.” *Id.* (citing J.A. 3263–64 ¶ 103 (Microsoft’s expert declaration)). Mr. Sandstrom has not shown why those findings are not reasonable considering the evidence as a whole.

B

Mr. Sandstrom argues that the Board did not show that the prior art would have enabled a relevant artisan to make and use the claimed inventions. *See generally* Sandstrom Informal Br. at 16–25 (citing *Raytheon Technologies Corp. v. General Electric Co.*, 993 F.3d 1374, 1380 (Fed. Cir. 2021) (“To render a claim obvious, the prior art, taken as a whole, must enable a skilled artisan to make and use the claimed invention.”)). Specifically, he argues that the Board in its *'833 Decision* did not explain how the prior art would have enabled claim 15’s requirement of “selecting, from a group of executable instances of a set of software programs . . . wherein the selection of the selected instances is based, at least in part, on a respective capacity demand indication of each of the set of software programs.” *'833 patent*, col. 21, line 63, through col. 22, line 2; *see also* Sandstrom Informal Br. at 9–10, 12–15, 22. He also argues that the Board in its *'090 Decision* did not explain that the prior art enabled a relevant artisan to

achieve claim 1's limitation, "for each of a series of core allocation periods (CAPs), selecting . . . specific tasks . . . for execution on the processor cores for a next CAP at least in part based on core capacity demand expressions associated with the processing tasks." '090 patent, col. 20, lines 1–6; Sandstrom Informal Br. at 21 (emphasis omitted).

Mr. Sandstrom did not make either of these enablement arguments to the Board. *See generally* J.A. 2240–77; J.A. 2645–76; J.A. 3693–3704; J.A. 3720–37. The arguments are therefore forfeited, given the important principle that strongly precludes faulting the Board for inadequacies not timely and fairly called to its attention. *See In re Google Technology Holdings LLC*, 980 F.3d 858, 863 (Fed. Cir. 2020) (“[A] position not presented in the tribunal under review will not be considered on appeal in the absence of exceptional circumstances.”).

On the merits, too, the argument fails on the record actually developed. The Board in its '833 *Decision* found that Agrawal, which itself discloses using “instantaneous parallelism,” *i.e.*, “the number of processors the job can effectively use at the current moment,” would have met claim 15's limitation. '833 *Decision*, at *10 (citing J.A. 857 (Agrawal)). And in its '090 *Decision*, the Board, with substantial support from Microsoft's expert declaration and the prior art references themselves, determined that Agrawal's disclosure of “scheduling quanta,” *i.e.*, scheduling tasks at regular time periods, and Chen's disclosure of considering the processing needs of different applications in scheduling “to maximize FPGA utilization,” would have enabled the claim 1 limitation Mr. Sandstrom highlights. '090 *Decision*, at *9–10 (citing J.A. 782 ¶ 102 (Microsoft's expert declaration), J.A. 857 (Agrawal); J.A. 3323 (Chen)). Mr. Sandstrom has shown no reversible error regarding enablement for obviousness purposes.

C

Mr. Sandstrom makes what amounts to a claim-construction argument about claim 1 of the '090 patent. Sandstrom Informal Br. at 8–12. He argues that the Board ignored claim 1's feature of “concurrently seeking collective optimality” and requirement of first selecting tasks to run and then assigning processors to tasks. *Id.* at 8. We disagree.

First, while Mr. Sandstrom perceives collective optimality to be an intended benefit of his invention, there is no such limitation within the claim. *See Howmedica Osteonics Corp. v. Wright Medical Technology, Inc.*, 540 F.3d 1337, 1347 (Fed. Cir. 2008) (“[T]he inventor’s subjective intent is irrelevant to claim construction.”). Second, Mr. Sandstrom did not present his order-of-steps claim-construction argument to Board, J.A. 3693–3704; J.A. 3720–37, so it is forfeited. *See In re Google*, 980 F.3d at 863. We further note that although claim 1 does recite selecting tasks before it recites assigning tasks to processors, method claims like claim 1 of the '090 patent “are not ordinarily construed to require” that recited steps (one of which necessarily is *recited* before the other) must be *performed* in the recited order. *See, e.g., Interactive Gift Express, Inc. v. Compuserve, Inc.*, 256 F.3d 1323, 1342 (Fed. Cir. 2001). Mr. Sandstrom has not shown that the matter he now raises is an exception.

III

We have considered Mr. Sandstrom’s other arguments and find them unpersuasive. For the foregoing reasons, we affirm the Board’s decisions.

The parties shall bear their own costs.

AFFIRMED