

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.,
Petitioner,

v.

SWARM TECHNOLOGY LLC,
Patent Owner.

IPR2022-00141
Patent 10,592,275 B2

Before MICHAEL R. ZECHER, GREGG I. ANDERSON, and
SCOTT B. HOWARD, *Administrative Patent Judges*.

ZECHER, *Administrative Patent Judge*.

DECISION
Granting Institution of *Inter Partes* Review
35 U.S.C. § 314

I. INTRODUCTION

A. *Background and Summary*

Petitioner, Juniper Networks, Inc. (“Juniper”), filed a Petition requesting an *inter partes* review (“IPR”) of claims 1–7 and 9–17 of U.S. Patent No. 10,592,275 B2 (Ex. 1001, “the ’275 patent”). Paper 1 (“Pet.”). Patent Owner, Swarm Technology LLC (“Swarm”), filed a Preliminary Response. Paper 10 (“Prelim. Resp.”). With our authorization, Juniper filed a Preliminary Reply (Paper 11, “Prelim. Reply”) and Swarm filed a Preliminary Sur-reply (Paper 12, “Prelim. Sur-reply”), each of which were narrowly tailored to address the real party in interest (“RPI”) and 35 U.S.C. § 325(d) arguments raised in Swarm’s Preliminary Response.

We have authority, acting on the designation of the Director, to determine whether to institute an IPR under 35 U.S.C. § 314 and 37 C.F.R. § 42.4(a). An IPR may not be instituted unless “the information presented in the petition filed under section 311 and any response filed under section 313 shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” 35 U.S.C. § 314(a) (2018). “When instituting [an IPR], the Board will authorize the review to proceed on all of the challenged claims and on all grounds of unpatentability asserted for each claim.” 37 C.F.R. § 42.108(a) (2021).

For the reasons set forth below, upon considering the briefing and the preliminary record, we determine that the information presented in the Petition establishes a reasonable likelihood that Juniper will prevail with respect to at least one of the challenged claims. Accordingly, we institute an IPR on all of the challenged claims based on all of the grounds raised in the Petition.

B. Real Parties in Interest

Juniper identifies itself as the only RPI. Paper 7, 1 (Juniper’s Updated Mandatory Notices). Juniper also states that it “will update its mandatory notices to identify Apstra [Inc.] as an RPI, if the Board determines that Apstra should be named as an RPI.” Prelim. Reply 4. Swarm identifies itself as the only RPI. Paper 9, 1 (Swarm’s Updated Mandatory Notices).

C. Related Matters

The parties identify the following two district court cases involving the ’275 patent: (1) *Juniper Networks, Inc. v. Swarm Technology LLC*, No. 3:20-cv-03137-JD (N.D. Cal.) and (2) *Swarm Technology, LLC v. Amazon.com, Inc.*, No. 2:21-cv-00438-DJH (D. Az.). Paper 7, 1; Paper 9, 1–2. Swarm also identifies U.S. Patent No. 9,852,004 B2 (“the ’004 patent”), U.S. Patent No. 9,146,777 B2 (“the ’777 patent”), and U.S. Patent Application No. 16/788,122 that claim or may claim the benefit of the filing date—January 25, 2013—of the ’275 patent. Paper 9, 2. The ’004 patent is the subject of IPR2021-01445 and the ’777 patent is the subject of IPR2021-01317. Prelim. Resp. 3.

D. The ’275 Patent

The ’275 patent is titled “System and Method for Swarm Collaborative Intelligence Using Dynamically Configurable Proactive Autonomous Agents.” Ex. 1001, code (54). The ’275 patent generally is directed to “parallel-process computing, and collaborative intelligence, and particularly relates to a processing architecture which involves autonomous co-processors (such as robotic vehicles, Internet of Things (IoT) components, and networked devices) configured to proactively retrieve tasks from a task pool populated by a central processing unit.” *Id.* at 1:16–24.

According to the '275 patent, “[c]omputer processors traditionally execute machine coded instructions serially. To run a plurality of applications concurrently, a single processor interleaves instructions from various programs and executes them serially, although from the user’s perspective the applications appear to be processed in parallel.” Ex. 1001, 1:49–54. The '275 patent further states that “[t]rue parallel or multi-core processing, on the other hand, is a computational approach that breaks large computational tasks into individual blocks of computations and distributes them among two or more processors.” *Id.* at 1:54–57. “A typical multiprocessor system includes a central processing unit (‘CPU’) and one or more co-processors. The CPU partitions the computational requirements into tasks and distributes the tasks to co-processors. Completed threads are reported to the CPU, which continues to distribute additional threads to the co-processors as needed.” *Id.* at 1:63–2:1.

The '275 patent identifies a problem with using the CPU to control the distribution of tasks:

Presently known multiprocessing approaches are disadvantageous in that a significant amount of CPU bandwidth is consumed by task distribution; waiting for tasks to be completed before distributing new tasks (often with dependencies on previous tasks); responding to interrupts from co-processors when a task is completed; and responding to other messages from co-processors. In addition, co-processors often remain idle while waiting for a new task from the CPU.

Id. at 2:1–10. The '275 patent purportedly addresses this problem using a system that reduces CPU management overhead and which “more effectively harnesses and exploits available co-processing resources.” *Id.* at 2:11–15.

Figure 1 of the '275 patent is reproduced below.




FIG. 1

Ex. 1001, Fig. 1. Figure 1, reproduced above, “is a schematic block diagram of a parallel processing architecture including a CPU, memory, task pool, and a plurality of co-processors configured to communicate through a fabric.” Ex. 1001, 3:64–67. More specifically, Figure 1 shows:

a single or multi-core CPU 11 and one or more solidarity or co-processing cells 12A-12[n] configured to communicate with a task pool 13 through a cross-bar switching fabric 14. The solidarity cells 12 may also communicate with each other through the switching fabric 14 or through a separate cell bus (not shown).

Id. at 4:43–49. “CPU 11 may communicate with the task pool 13 directly or through the switching fabric 14. One or more memory units 15 each contain data and/or instructions.” *Id.* at 4:49–52.

E. Illustrative Claim

Claims 1, 6, and 11 are independent. Independent claims 1 and 6 are each directed to “[a] collaborative intelligence system,” whereas independent claim 11 is directed to “[a] method of using collaborative intelligence to achieve an objective.” Ex. 1001, 14:24, 15:8, 16:25–26. Claims 2–5 directly depend from independent claim 1; claims 7, 9, and 10 directly or indirectly depend from independent claim 6; and claims 12–17 directly depend from independent claim 11. *Id.* at 14:50–15:7, 16:8–11, 16:17–24, 16:42–64. Independent claim 1, reproduced below, is illustrative of the challenged claims.

1. A collaborative intelligence system, comprising:
 - a task pool;
 - a controller configured to populate the task pool with a plurality of first tasks and a plurality of second tasks;
 - a first co-processor configured to successively: proactively retrieve a first task from the task pool; process the first task; generate first resulting data; and update the task pool to reflect completion of the first task, all without any communication between the first co-processor and the controller; and
 - a second co-processor configured to successively: proactively retrieve a second task from the task pool; process the second task; generate second resulting data; and update the task pool to reflect completion of the second task, all without any communication between the second co-processor and the controller;
- wherein the collaborative intelligence system is configured to dynamically accept the first co-processor, the second co-processor, and an additional co-processor into the processing system on a plug-and-play basis without any communication with the controller;

the plurality of first tasks and the plurality of second tasks are associated with a common objective;

the first and second co-processors autonomously work together in solidarity with the task pool to complete the common objective.

Id. at 14:24–49.

F. Asserted Prior Art References

Juniper relies on the prior art references set forth in the tables below.

Name¹	Reference	Dates	Exhibit No.
Leong	U.S. Patent No. 6,006,249	issued Dec. 21, 1999; filed Aug. 19, 1997	1005

Printed Publication	Exhibit No.
Mike Rogers & Virginia Bare, <i>Hands-on AppleTalk</i> (1989) (“AppleTalkBook”). All citations are to the page numbers added by Juniper.	1006
The Ethernet: A Local Area Network, Data Link Layer and Physical Layer Specifications (Sept. 30, 1980) (“Ethernet Specification”). All citations are to the native pagination.	1007
Dinesh Datla et al., <i>Wireless Distributed Computing: A Survey of Research Challenges</i> , IEEE Communications Magazine, Jan. 2012, at 144–52 (“Datla”). All citations are to the page numbers added by Juniper.	1019

Juniper also relies on the testimony of Dr. Jon B. Weissman (Ex. 1003 (Declaration of Dr. Weissman)) and Sylvia D. Hall-Ellis, Ph.D. (Ex. 1018 (Declaration of Dr. Hall-Ellis)).

¹ For clarity and ease of reference, we only list the first named inventor.

G. Asserted Grounds of Unpatentability

Juniper asserts that claims 1–7 and 9–17 of the '275 patent would have been unpatentable based on the grounds set forth in the table below. Pet. 3, 12–86.

Claim(s) Challenged	35 U.S.C. §²	Reference(s)/Basis
1–7, 9–17	103(a)	Leong, AppleTalkBook
1–7, 9–17	103(a)	Leong, AppleTalkBook, Ethernet Specification
16	103(a)	Leong, AppleTalkBook, Datla
16	103(a)	Leong, AppleTalkBook, Datla, Ethernet Specification

II. ANALYSIS

A. Legal Standards

In *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966), the Supreme Court set out a framework for assessing obviousness under 35 U.S.C. § 103 that requires consideration of four factors: (1) the “level of ordinary skill in the pertinent art,” (2) the “scope and content of the prior art,” (3) the “differences between the prior art and the claims at issue,” and (4) when in evidence, “secondary considerations” of non-obviousness such as “commercial success, long[-]felt but unsolved needs, failure of others, etc.” *Id.* at 17–18. “While the sequence of these questions might be reordered in any particular case,” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S.

² The Leahy-Smith America Invents Act (“AIA”) included revisions to 35 U.S.C. §§ 102, 103 that became effective on March 16, 2013. Because the '275 patent claims the benefit of the priority date of U.S. Patent Application No. 13/750,696, which was filed before March 16, 2013, and neither party has argued that the provisions of the AIA apply, we apply the pre-AIA versions of the statutory bases for unpatentability. *See* Ex. 1001, code (63).

398, 407 (2007), the U.S. Court of Appeals for the Federal Circuit has repeatedly emphasized that “it is error to reach a conclusion of obviousness until all those factors are considered,” *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1328 (Fed. Cir. 2016).³

B. Level of Ordinary Skill in the Art

In determining whether an invention would have been obvious at the time it was made, we consider the level of ordinary skill in the pertinent art at the time of the invention. *Graham*, 383 U.S. at 17. “The importance of resolving the level of ordinary skill in the art lies in the necessity of maintaining objectivity in the obviousness inquiry.” *Ryko Mfg. Co. v. Nu-Star, Inc.*, 950 F.2d 714, 718 (Fed. Cir. 1991). The “person having ordinary skill in the art” is a hypothetical construct from whose vantage point obviousness is assessed. *In re Rouffet*, 149 F.3d 1350, 1357 (Fed. Cir. 1998).

Factors pertinent to a determination of the level of ordinary skill in the art include “(1) the educational level of the inventor; (2) type of problems encountered in the art; (3) prior art solutions to those problems; (4) rapidity with which innovations are made; (5) sophistication of the technology; and (6) educational level of active workers in the field.” *Env’t. Designs, Ltd. v. Union Oil Co. of Cal.*, 713 F.2d 693, 696 (Fed. Cir. 1983) (citing *Orthopedic Equip. Co. v. All Orthopedic Appliances, Inc.*, 707 F.2d 1376, 1381–82 (Fed. Cir. 1983)). “Not all such factors may be present in every case, and one or

³ At this preliminary stage of the proceeding, because neither party addresses objective evidence of non-obviousness, we focus solely on the first three *Graham* factors.

more of these or other factors may predominate in a particular case.” *Id.* at 696–97.

In its Petition, Juniper argues that a person having ordinary skill in the art “would’ve had at least a bachelor’s degree in computer science, electrical engineering, computer engineering, or a closely related field, and one or more years of experience in the design and development of parallel or distributed processing systems.” Pet. 11. Alternatively, Juniper argues that a person having ordinary skill in the art “would have a master’s degree or similar post-graduate work in computer science, electrical engineering, computer engineering, or a closely related field, and less design and development experience.” *Id.* (citing Ex. 1003 ¶¶ 60, 61). In its Preliminary Response, Swarm does not address the level of skill in the art. *See* Prelim. Resp.

Accordingly, for purposes of institution, we adopt Juniper’s proposed formulation of the level of ordinary skill in the art, except that we delete the qualifier “or more” to eliminate vagueness as to the amount of practical experience. The qualifier expands the range indefinitely without an upper bound and, therefore, precludes a meaningful indication of the level of ordinary skill in the art.

C. Claim Construction

We apply the same claim construction standard used in the federal courts, in other words, the claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. § 282(b), which is articulated in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). *See* 37 C.F.R. § 42.100(b). Under the *Phillips* standard, “the words of a claim ‘are generally given their ordinary and customary meaning,’” which is “the meaning that the term would have to a person of ordinary skill

in the art in question at the time of the invention, i.e., as of the effective filing date of the patent application.” *Phillips*, 415 F.3d at 1312–13.

In its Petition, Juniper states that, “[f]or the purposes of this proceeding, [Juniper] construes all terms as having their ordinary and customary meanings.” Pet. 11. In its Preliminary Response, Swarm does not address claim construction. *See* Prelim. Resp.

Based on the preliminary record, because there is no dispute as to the construction of any claim term for purposes of institution, we need not construe any claim term at this time. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (noting that “we need only construe terms ‘that are in controversy, and only to the extent necessary to resolve the controversy’” (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999))).

D. Alleged Failure to Name All RPIs

Swarm argues that Juniper failed to name various entities as RPIs. Prelim. Resp. 2–17; Prelim. Sur-reply 2–3. These entities include Amazon.com, Inc. and Amazon Web Services, Inc. (collectively, “Amazon”) and Apstra, Inc. (“Apstra”). Prelim. Sur-reply 2–3. Swarm also argues that Cisco, Arista, RPX, Intel, Oracle, Qualcomm, HP, AMD, Atmel, Google, Nest, Apple, and NVIDIA have a stake in the outcome of this proceeding. Prelim. Resp. 2–16; Prelim. Sur-reply 2–3. More specifically, Swarm argues that Amazon was originally named as an RPI in IPR2021-01317, and Swarm is asserting the ’275 patent against both Amazon and Juniper. Prelim. Resp. 9–10. Swarm also contends that Apstra is a subsidiary of Juniper and was added as a necessary party to Juniper’s declaratory judgment action seeking a declaration of no infringement. *Id.* at 11–14. Swarm also argues that Juniper had “secret communications and

behind-the scenes coordination with several parties having a vested interest in the validity of the '275 [p]atent.” *Id.* at 14; *see also id.* at 14–17 (discussing communications and the parties).

With respect to Amazon, Swarm further argues that this proceeding is factually distinguishable from the one in our precedential decision, *SharkNinja Operating LLC v. iRobot Corp.*, IPR2020-00734, Paper 11 (PTAB Oct. 6, 2020) (precedential). First, Swarm argues that, unlike *SharkNinja*, which involved an omission of an RPI, this case involves Juniper affirmatively naming Amazon in a related case and then unnamings Amazon. Prelim. Resp. 9. Second, Swarm argues Juniper has not offered to correct its Mandatory Notices to include Amazon. *Id.* at 9–10. Third, Swarm argues that, “[u]nlike in *SharkNinja*, where the omitted [RPI] was simply the corporate parent of the petitioner, Amazon is a major customer of Juniper and is presently a defendant in a separate patent infringement litigation involving the '777, '004, and '275 Patents in Arizona.” *Id.* at 10. Fourth, Swarm argues that, unlike in *SharkNinja*, “the evidence in this case suggests that Amazon was intentionally removed as [an RPI] in the Juniper IPRs to gain an improper legal advantage” by avoiding the statutory estoppel provision of 35 U.S.C. § 315(e)(2).⁴ *Id.*

We are not persuaded by Swarm’s RPI arguments because none of the entities identified by Swarm are time-barred or estopped from challenging the '275 patent, as none were accused of infringing this patent in a litigation more than a year before the Petition’s filing date of November 4, 2021.

⁴ Amazon filed a separate petition requesting an IPR of claims 1–17 of the '275 patent. *Amazon.com, Inc. v. Swarm Tech. LLC*, IPR2022-00633, Paper 2 (PTAB Feb. 25, 2022).

Prelim. Reply 1–2 (citing Ex. 1022 (listing cases involving the ’275 patent)). *SharkNinja* holds that the Board need not consider whether an unnamed party is an RPI when the addition of that party “would not create a time bar or estoppel under 35 U.S.C. § 315.” *SharkNinja*, Paper 11 at 18. In reaching that conclusion, the Board recognized that, even if the omitted party was an unnamed RPI, “our jurisdiction to consider a petition does not require a ‘correct’ identification of all RPIs in a petition,” and the Director can allow the petitioner to add the omitted party. *Id.* (citing *Lumentum Holdings, Inc. v. Capella Photonics, Inc.*, IPR2015-00739, Paper 38 at 6 (PTAB Mar. 4, 2016) (precedential) (“*Lumentum*”); *Blue Coat Sys., Inc. v. Finjan, Inc.*, IPR2016-01444, Paper 11 at 10 (PTAB July 18, 2017); *Mayne Pharma Int’l Pty. Ltd. v. Merck Sharp & Dohme Corp.*, 927 F.3d 1232, 1240 (Fed. Cir. 2019)). Because the failure to name an RPI in those circumstances does not bar the proceeding, not considering the RPI issue “better serves the interest of cost and efficiency.” *Id.* at 20.

In this proceeding, Swarm has not argued that the addition of any of the alleged omitted RPIs to the Petition would create a time bar under 35 U.S.C. § 315(b) or estoppel under 35 U.S.C. § 315(e)(1).⁵ Nor do we find any of the reasons given by Swarm to distinguish this case from *SharkNinja* persuasive. Although the facts of this case are not identical to those in *SharkNinja*, none of those differences go to the heart of *SharkNinja*’s holding—namely, “cost and efficiency.” Accordingly, based on *SharkNinja* and in the interests of efficiency and controlling costs, we do

⁵ Any potential estoppel in a civil action or other proceeding, under 35 U.S.C. § 315(e)(2), would not arise until after “an [IPR] . . . that results in a final written decision.” Any future estoppel in another proceeding has no bearing on the outcome of this proceeding.

not consider whether any of the fifteen identified entities should have been named as RPIs.

E. Alleged Failure to Identify All Related Matters

Swarm also argues that we should deny institution because Juniper failed to accurately identify a co-pending litigation and failed to identify other patents that share the same effective filing date as the '275 patent. Prelim. Resp. 17–19 (citing 37 C.F.R. § 42.8); *see* Prelim. Sur-reply 3–4.

We are not persuaded by Swarm's argument in this regard for four reasons. First, Juniper has offered to “stipulate to [Swarm's] February 9, 2022[,] Section 42.8(b)(2) Related Matters and/or update its Mandatory Notices to identify those matters.” Prelim. Reply 5.

Second, Swarm has not identified any case where we have declined to institute for this reason alone. On the contrary, our precedent holds that other formalities under 35 U.S.C. § 312(a),⁶ such as the identification of the RPI, are not jurisdictional and can be cured. *See Lumentum*, Paper 38 at 3–6.

Third, to the extent that 37 C.F.R. § 42.8 requires the parties to file notice of the co-pending litigation and other patents as related judicial and administrative matters, Swarm satisfied the mandatory notice requirement. Swarm identified the related district court litigation—including Juniper and Apstra in the California case caption—and the '777 and '004 patents in its mandatory notices. *See* Paper 9, 1–2. In *SharkNinja*, the decision explained,

⁶ Section 312(a) of the United States Code requires, *inter alia*, that the Petition must include “such other information as the Director may require by regulation.” Our regulations require that the Petition must “[i]dentify any other judicial or administrative matter that would affect, or be affected by, a decision in the proceeding.” 37 C.F.R. § 42.8.

and we agree, that § 312 “serves important notice functions . . . to the Board, to identify conflicts of interests that are not readily apparent from the identity of the petitioner.” *SharkNinja*, Paper 11 at 17 (citing *NOF Corp. v. Nektar Therapeutics*, IPR2019-01397, Paper 24 at 6 (PTAB Feb. 10, 2020); Consolidated Trial Practice Guide 12 (Nov. 2019)⁷). Similarly, we see the mandatory notice requirements of 37 C.F.R. § 42.8 serving the same function of identifying conflicts of interest for the Board.

Fourth, we see no error in the identification of the co-pending litigation. Swarm does not direct us to, nor are we aware of, a rule that requires the listing of every party in a case caption. Indeed, the inclusion of the names of multiple entities—as Swarm requires—is inconsistent with proper Bluebook citation.

F. Alleged Failure to Satisfy the Requirements under 35 U.S.C. § 312(a)(3)

Swarm contends that Juniper violates the statutory requirements under 35 U.S.C. § 312(a)(3). Prelim. Resp. 19–22. Under § 312(a)(3), a petition may only be considered if it “identifies, in writing and with particularity, each claim challenged, the grounds on which the challenge to each claim is based, and the evidence that supports the grounds for the challenge to each claim.” 35 U.S.C. § 312(a)(3); *see also* 37 C.F.R. § 42.104(b) (specifying necessary elements of a petition).

As the Federal Circuit has explained, “[i]n an IPR, the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.” *Adaptics Ltd. v. Perfect Co.*, IPR2018-01596, Paper 20 at 15–16 (March 6, 2019) (quoting *Harmonic Inc. v. Avid Tech.*,

⁷ Available at <https://www.uspto.gov/TrialPracticeGuideConsolidated>.

Inc., 815 F.3d 1356, 1363 (Fed. Cir. 2016)). Juniper relies on Leong as the primary reference for each of its asserted grounds of unpatentability. Pet. 3, 12–86. Under each of these asserted obviousness grounds, Juniper also relies on AppleTalkBook to address a plug-and-play processing system, without any communication with a controller (identified as claim elements 1.5, 6.5, 11.2), and a plurality of autonomous co-processors operating in solidarity to complete a common objective (identified as claim element 11.3). *See* Pet. 32–38, 59, 73; Prelim. Resp. 20. Swarm, however, contends that despite Juniper’s designation of the asserted obviousness ground solely based on Leong and AppleTalkBook, the arguments under this particular ground rely on, and require support from, Ethernet Specification. Prelim. Resp. 20 (citing Pet. 36–37, 44–46, 49–53, 58–69, 75–78). Swarm also argues that Juniper’s identification and analysis of Ethernet Specification as a reference under the remaining asserted obviousness grounds are deficient. *Id.* at 21–22.

We are not persuaded that Juniper violated the statutory requirements of § 312(a)(3). First, Juniper uses Ethernet Specification in its analysis of its asserted obviousness ground solely based on Leong and AppleTalkBook merely as evidence of the state of the art. The Federal Circuit has explained that “[a]rt can legitimately serve to document the knowledge that skilled artisans would bring to bear in reading the prior art identified as producing obviousness.” *Ariosa Diagnostics v. Verinata Health, Inc.*, 805 F.3d 1359, 1365 (Fed. Cir. 2015); *see also Randall Mfg. v. Rea*, 733 F.3d 1355, 1362–63 (Fed. Cir. 2013) (emphasizing that additional prior art references or evidence are not for the purposes of changing the prior art combination that forms the basis of the asserted ground, but rather “constitute important

evidence of the state of the art,” including the general background knowledge of a person having ordinary skill in the art).⁸

Here, Juniper does not map any limitation of the challenged claims to Ethernet Specification. *See* Pet. 12–78. Rather, Juniper cites Ethernet Specification as evidence of the general background knowledge of a person having ordinary skill in the art. More specifically, Juniper cited Ethernet Specification to explain that “[t]he Ethernet standard has a goal that all Ethernet implementations ‘should be capable of exchanging data at the data link level.’” Pet. 44 (quoting Ex. 1007, 12). Juniper also cited Ethernet Specification to show that a person having ordinary skill in the art would have understood that Leong’s processing units—connected and communicating through an Ethernet network—use the underlying mechanisms of an Ethernet network, including standard communication parameters in Ethernet data frames (i.e., source address, destination address, and data fields or payloads). Pet. 36–37, 44–46, 49–53, 58–69, 75–78. Because Juniper’s citation to Ethernet Specification shows that Leong’s operations are consistent with the standards for participation and interoperation in a standard Ethernet network, we are persuaded that Ethernet Specification was used as evidence of the state of the art, not to form the basis of the asserted obviousness ground.

⁸ *Cf.* Applicant Admitted Prior Art Memorandum at 9, *available at* https://www.uspto.gov/sites/default/files/documents/signed_aapa_guidance_memo.pdf (“Statements made in the specification of the patent that is being challenged in an IPR can be used as evidence of [the] general knowledge [of a person of ordinary skill in the art], and thus, . . . can be used to . . . supply missing claim limitations that were generally known in the art prior to the invention.”).

Second, to the extent that Juniper’s citation to Ethernet Specification under its asserted obviousness ground solely based on Leong and AppleTalkBook goes beyond evidence of the state of the art, Juniper includes Ethernet Specification as part of the prior art combination forming the basis of its remaining asserted obviousness grounds. Therefore, we are not persuaded that Juniper violated the statutory requirements under § 312(a)(3).

G. Discretionary Denial under 35 U.S.C. § 325(d)

1. Legal Standard

Swarm urges us to exercise our discretion to deny institution under 35 U.S.C. § 325(d). Prelim. Resp. 33–48. Under § 325(d), we have discretion to deny a petition that presents the same or substantially the same prior art or arguments as previously presented to the Office. In evaluating whether the factual predicate under § 325(d) is met, we consider a number of non-exclusive factors, as set forth in *Becton, Dickinson & Co. v. B. Braun Melsungen AG*, IPR2017-01586, Paper 8 at 17–18 (PTAB Dec. 15, 2017) (Institution Decision) (precedential as to Section III.C.5, first paragraph) (“the *Becton Dickinson* factors”):

(a) the similarities and material differences between the asserted art and the prior art involved during examination;

(b) the cumulative nature of the asserted art and the prior art evaluated during examination;

(c) the extent to which the asserted art was evaluated during examination, including whether the prior art was the basis for rejection;

(d) the extent of the overlap between the arguments made during examination and the manner in which Petitioner relies on the prior art or Patent Owner distinguishes the prior art;

(e) whether Petitioner has pointed out sufficiently how the Examiner erred in its evaluation of the asserted prior art; and

(f) the extent to which additional evidence and facts presented in the Petition warrant reconsideration of the prior art or arguments.

Becton, Dickinson, IPR2017-01586, Paper 8 at 17–18.

In performing an analysis under § 325(d), “the Board uses the following two-part framework:

(1) whether the same or substantially the same art previously was presented to the Office or whether the same or substantially the same arguments previously were presented to the Office; and (2) if either condition of first part of the framework is satisfied, whether the petitioner has demonstrated that the Office erred in a manner material to the patentability of challenged claims. . . . If, after review of [*Becton, Dickinson*] factors (a), (b), and (d), it is determined that the same or substantially the same art or arguments previously were presented to the Office, then factors (c), (e), and (f) relate to whether the petitioner has demonstrated a material error by the Office.

Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte GmbH, IPR2019-01469, Paper 6 at 8, 10 (PTAB Feb. 13, 2020) (Institution Decision) (precedential) (footnote omitted). Consequently, we first turn to an analysis of *Becton Dickinson* factors (a), (b), and (d) to determine whether the same or substantially the same art previously was presented to the Office or whether the same or substantially the same arguments previously were presented to the Office.

2. Part One of the Advanced Bionics Analysis

Becton Dickinson factors (a), (b), and (d) require us to determine, respectively: (a) “the similarities and material differences between the asserted art and the prior art involved during examination;” (b) “the cumulative nature of the asserted art and the prior art evaluated during

examination,” and (d) “the extent of the overlap between the arguments made during examination and the manner in which Petitioner relies on the prior art or Patent Owner distinguishes the prior art.” *Becton, Dickinson*, IPR2017-01586, Paper 8 at 17.

a) Prosecution History of the '275 Patent

During prosecution of the '275 patent, the Examiner considered Bates (U.S. Patent Publication No. 2007/0074207 A1 (Ex. 1008)), Inoue (U.S. Patent Publication No. 2005/0188372 A1 (Ex. 2011)), Sikdar (U.S. Patent Publication No. 2006/0026377 A1 (Ex. 2013)), DeMent (U.S. Patent Publication No. 2008/0168318 A1 (Ex. 2012)), and Tozlu (Serbulent Tozlu et al., *Wi-Fi Enabled Sensors for Internet of Things: A Practical Approach*, IEEE Commc'ns Mag., June 2012, at 134 (Ex. 2014)), and rejected original claims 1–20 of the '275 patent application as obvious over the combined teachings of Bates and Tozlu. Prelim. Resp. 40 (citing Ex. 1002, 106–11, 123–24, 359–60, 559). The Examiner, however, did not consider Leong, AppleTalkBook, Ethernet Specification, or Datla either alone or in combination during prosecution of the '275 patent. Therefore, Swarm argues that these references are substantially the same as those considered by the Examiner during prosecution of the '275 patent. We address each of Swarm's arguments in turn.

b) Leong

Swarm argues that Leong is substantially the same as Bates, which was considered during prosecution of the '275 patent. Prelim. Resp. 40, 43–44; Ex. 1002, 123–24, 359–60, 559. Swarm's argument consists of an annotated figure showing the purported similarity of features and the following paragraph:

Like Leong, Bates discloses a plurality of processing units (104) connected via a memory bus to a memory (106) containing a task queue (114). Bates at 2–3, ¶¶ 0032–0033, 0036. Like Leong, Bates discloses that the plurality of processing units 104 select tasks from the task queue. Bates at 6, ¶ 0057. [Juniper’s] use of Leong throughout the Petition mirrors Examiner Rashid’s use of Bates throughout the 2015 [Non-Final Office Action] of the ’777 Patent (Ex. 2007 at 10–17), the prosecution of the ’004 Patent[] (Ex. 2008 at 5–10, Ex. 2009 at 3–8, and Ex. 2010 at 3–8), and the prosecution of the ’275 Patent (Ex. 1002 at 106–111).

Id. at 43–44.


We are not persuaded by Swarm’s argument that Leong and Bates are substantially the same. Although these two references share many similar features, there are material differences between how the systems of Leong and Bates operate. Unlike the ’275 patent, which operates without communication between the co-processors and the controller, Bates’s power processor units (PPU) “act[] as a controller” for the synergistic processing units (SPUs). Ex. 1008 ¶ 33. During prosecution of the ’275 patent, Swarm acknowledged that “a signal, command, or instruction must be transmitted between the PPU and the SPU. . . . Bates thus required direct communication between the PPU and an SPU.” Ex. 1002, 337–38 (citing Ex. 1008 ¶ 0059). In contrast, Leong teaches “processing units 12 each of which is self[-]controlled using a software application running on its own operating system. Thus, there is no central management and/or control unit.” Ex. 1005, 2:45–48; Prelim. Reply 6 (citing Pet. 23–38) (“Unlike Inoue/Bates, Leong’s other processing units don’t contact the surveying unit.”). Therefore, Leong’s system does not operate in substantially the same way as Bates’s system.

Swarm also argues that Leong is substantially the same as Inoue, which was considered during prosecution of the ’275 patent. Prelim. Resp.

40–43; Ex. 1002, 123–24, 359–60, 559. Swarm’s argument consists of an annotated figure and a discussion of the purported similarity of features.

Prelim. Resp. 40–43. Although these references share many similar features, the issue is not just whether the same features can be found in each of the references, but whether these features function the same way. That is, even though Leong and Inoue teach similar features, that does not mean that those features interact with each other in the same way or perform the same functions as recited in the claims.

In fact, there are material differences between how the systems of Leong and Inoue operate. Unlike the ’275 patent, which operates without communication between the co-processors and the controller, Inoue teaches that “the sub-processing units select processor tasks from the shared memory for execution” and “[t]he sub-processing units 102 also preferably conduct an initialization routine in order to create the task table 280 and the task queue 282,” as illustrated in Figure 11, reproduced below. Ex. 2011 ¶¶ 10, 90; Prelim. Reply 6 (citing Pet. 23–38) (“Unlike Inoue/Bates, Leong’s other processing units don’t contact the surveying unit.”).



Ex. 2011, Fig. 11. Figure 11 of Inoue, reproduced above, illustrates the initialization processes. By action 354, “the service processor [(e.g., the main processing unit 102)] is preferably no longer involved in the maintenance and/or allocation of the processor tasks among the sub-processing units 102” and, at action 356, “the sub-processing units 102 initialize with one another to determine which SPU is to prepare the task table 280 and the task queue 282.” *Id.* ¶¶ 90–92. Accordingly, the sub-processing units, which act as both “a controller configured to populate the task pool with a plurality of . . . tasks” and as “co-processors” configured to retrieve tasks, process tasks, generate resulting data, and update the task pool, are communicating with one another during the initialization processes. *Compare* Ex. 2011 ¶¶ 92–93 (stating that “the initialization of the task table 280 and task queue 282 is preferably performed by having each SPU kernel execute an initial task” before “looking for [a] next task”), *with*

Ex. 1001, 14:26–39. In contrast, Leong teaches that its “surveying agents may be programmed to only survey the memory 14” and, therefore, avoid communications with the other processing units. Ex. 1005, 4:17–25.

Therefore, Swarm has not sufficiently shown that Leong is substantially the same as Inoue.

c) Ethernet Specification

Swarm argues that Ethernet Specification is substantially the same as Sikdar, which was considered during prosecution of the ’275 patent. Prelim. Resp. 40, 44–45; Ex. 1002, 123–24, 359–60, 559. Once again, Swarm’s argument consists of an annotated figure and a discussion of the purported similarity of features. Prelim. Resp. 44–45. Although these references teach many of the same features, Juniper cites Ethernet Specification as evidence of the general background knowledge of a person having ordinary skill in the art, as discussed above. *See supra* Section II.F. In addition, Juniper argues that, “[u]nlike Sikdar, Ethernet [Specification] provides details on the specification’s data frame fields.” Prelim. Reply 6 (citing Pet. 79–81). The Examiner also never used Sikdar, alone or combined with the teachings of Leong and AppleTalkBook, in a rejection of the original claims that eventually issued in the ’275 patent. *See* Ex. 1002. Thus, Swarm has not sufficiently shown that the asserted obviousness combinations including Ethernet Specification are substantially the same as Sikdar.

d) Plug-and-Play Limitation

Swarm argues that Juniper’s arguments with respect to the plug-and-play limitation are the same or substantially the same as those previously made by the Office. Prelim. Resp. 45–46. Swarm contends that Juniper’s use of AppleTalkBook is substantially the same as the Examiner’s use of

DeMent during prosecution of the '004 patent. *Id.* (quoting Ex. 2010, 5⁹ (citing Ex. 2012 ¶¶ 33, 44, 45, 48, 50)).

The Examiner relied on DeMent during prosecution of the '004 patent to argue the following:

It would have been obvious to a person having ordinary skill in the art at the time of invention to incorporate the teachings of DeMent to plug into a printed circuit board the package that is housing the PPU and SPU to at least make the PPU and SPU die operational. The *plugging* is done without any “communication” with the controller.

Ex. 2010, 6 (emphasis added). The Examiner, however, did not argue that DeMent teaches “playing” without any communication with the controller. In fact, DeMent teaches “multiple synergistic processor element (SPE) cores 320–334 that each has its own synergistic processor unit (SPU) 340–354, memory flow control 355–362, local memory or store (LS) 363–370, and bus interface unit (BIU) 380–394.” Ex. 2012 ¶ 44, Fig. 3. Particularly, “[t]he memory flow control units (MFCs) 355–362 serve as an interface for an [SPE] to the rest of the system,” including “provid[ing] the primary mechanism for data transfer, protection, and synchronization between main storage and the local storages.” *Id.* ¶ 50. Further, “SPE cores 320–334 are coupled to [memory interface controller (MIC)] 398 and [bus interface controller (BIC)] 397 via the [element interconnect bus (EIB)] 396,” where “MIC 398 provides a communication interface to shared memory 399” and “BIC 397 provides a communication interface between the CBE 300 and other external buses and devices.” *Id.* ¶ 46, Fig. 3. Therefore, unlike

⁹ All references to the prosecution history of the '004 patent refer to the page numbers inserted by Juniper in the bottom right-hand corner of each page in Exhibit 2010.

AppleTalkBook, DeMent teaches that new co-processors communicate with a controller when accessing a task pool in the shared memory. *Id.* ¶¶ 44–46, 50, Fig. 3.

In addition, Juniper argues that, “[u]nlike DeMent, AppleTalkBook details plug-and-play implementations using broadcast packets.” Prelim. Reply 6 (citing Pet. 17–19, 32–38). Thus, even though DeMent was cited in two Information Disclosure Statements that were submitted during prosecution of the ’275 patent, Swarm has not sufficiently shown that AppleTalkBook is substantially the same as DeMent. *See* Ex. 1002, 125, 360.

e) Datla

Swarm argues that Datla is substantially the same as Tozlu, which was considered and used by the Examiner to reject original claims 11, 13–18, and 19 during prosecution of the ’275 patent. Prelim. Resp. 40, 46–47; Ex. 1002, 109–11, 123–24, 359–60, 559. Swarm contends that Juniper’s use of Datla in the asserted obviousness grounds challenging claim 16—based on the combined teachings of Leong, AppleTalkBook, and Datla, with and without the teachings of the Ethernet Specification—is substantially the same as the Examiner’s use of Tozlu during prosecution of the ’004 patent. Prelim. Resp. 46–47 (citing Ex. 2008, 10 (arguing that Tozlu teaches “integrating the first cell into a first device prior to adding the programmed cell to the network,” devices comprising IoT devices, “establishing a communication link between the first cell and the task pool,” and co-processors not residing on the monolithic integrated circuit)).

Although these references teach many of the same general features related to utilizing processing power available in other IoT devices, Tozlu “investigate[s] the feasibility of WiFi for battery-powered, low-power

sensors.” Ex. 2014, 1. Specifically, Tozlu focuses on the power consumption of sensors when performing periodic and event-triggered data transmissions. *Id.* at 1–4. In contrast, Juniper relies on Datla to teach how “[w]ireless distributed computing (WDC) exploits wireless connectivity to share processing-intensive tasks among multiple devices.” Pet. 82–86; Prelim. Reply 6 (citing Pet. 82–86) (arguing that, “unlike Tozlu, Datla’s wireless nodes share processing-intensive tasks”). Specifically, Juniper relies on Datla to teach the enabling technology and possible implementations. *See* Pet. 82–86 (citing Ex. 1019, 2–4, 6–7) (describing the benefits of WDC, including in-network processing, decentralized control, and collaboration on executing tasks). Therefore, even though Datla and Tozlu are in the same field of endeavor, Datla is not substantially the same as Tozlu.

In addition, the Examiner did not combine Tozlu with Leong, AppleTalkBook, or the Ethernet Specification during prosecution of the ’275 patent. Accordingly, Swarm has not sufficiently shown that the asserted obviousness combinations including Datla are substantially the same as Tozlu.

3. Conclusion

Because we find that the evidence of record demonstrates that *Becton Dickinson* factors (a), (b), and (d) have not been satisfied, we find that the same or substantially the same art or arguments were not presented previously to the Office. Thus, our analysis ends at this point, and we need not reach step two of the *Advanced Bionics* framework. *See Advanced Bionics*, IPR2019-01469, Paper 6 at 8. Consequently, we decline to exercise our discretion to deny institution based on § 325(d).

H. Asserted Obviousness in View of Leong and AppleTalkBook

Juniper argues that claims 1–7 and 9–17 would have been obvious over the combined teachings of Leong and AppleTalkBook. *See* Pet. 12–78. Based on the preliminary record, we are sufficiently persuaded that Juniper has established a reasonable likelihood of prevailing on this asserted obviousness ground with respect to at least one of claims 1–7 and 9–17.

1. Overview of Leong

Leong is titled “Method and Apparatus for Concurrent Data Processing.” Ex. 1005, code (54). Leong is directed “to multi-tasking systems employing networked stand alone and independently operating micro-processing units each of which perform one or more tasks for manipulating electronic data.” *Id.* at 1:7–11. Leong states it is an improvement on prior art multi-tasking computing systems that used a CPU coupled to a plurality of processing units. *Id.* at 1:19–24, Fig. 1.

Leong Figure 2 is reproduced below.




Figure 2, reproduced above, “is a block diagram of a multi-tasking computing system.” Ex. 1005, 2:14–16. Each processing unit 12 “is self controlled using a software application running on its own operating system”

and “there is no central management and/or control unit.” *Id.* at 2:43–48. One or more processing units 12 determine which tasks need to be performed and post status information about the tasks on bulletin board 14A. *Id.* at 4:9–16. Additionally, each processing unit 12 is capable of reading the tasks on the bulletin board and executing tasks that it is capable of performing. *Id.* at 3:6–18.

2. Overview of AppleTalkBook

AppleTalkBook describes itself as “a practical book” that “is about the process of choosing and using an AppleTalk network.” Ex. 1006, 11. Relevant to this proceeding, AppleTalkBook discusses “dynamic configuration.” *Id.* at 104. Dynamic configuration refers to the “self-configuring, ‘automatic’ nature of AppleTalk” and is characterized by each device having “the ability to pick its own unique network address” and the ability of devices “to look up across the network and identify, at any given time, what resources are available.” *Id.* at 105.

3. Analysis of Independent Claim 1

a) Undisputed Limitations

The preamble of independent claim 1 recites “[a] collaborative intelligence system.” Ex. 1001, 14:24–25. To the extent the preamble should be treated as limiting, Juniper argues Leong teaches the feature recited in the preamble. Pet. 20–22. Specifically, Juniper argues Leong teaches a processing system including several processing units associated with a bulletin board. *Id.* at 20 (citing Ex. 1005, Fig. 2; Ex. 1003 ¶ 74).

Independent claim 1 further recites “a task pool.” Ex. 1001, 14:25. Juniper argues Leong teaches this limitation. Pet. 22–23. Specifically, Juniper argues Leong teaches an electronic bulletin board that resides in a region of physical memory accessible by the CPU. *Id.* at 22 (citing

Ex. 1005, 3:9–11, Fig. 2; Ex. 1003 ¶ 80). Juniper also argues that “Leong’s processing units read the bulletin board when seeking additional tasks.”¹⁰ *Id.* (citing Ex. 1005, 3:9–18; Ex. 1003 ¶ 80); *see also id.* at 22–23.

Independent claim 1 further recites “a first co-processor configured to successively: proactively retrieve a first task from the task pool; process the first task; generate first resulting data; and update the task pool to reflect completion of the first task, all without any communication between the first co-processor and the controller.” Ex. 1001, 14:28–33. Juniper argues Leong teaches that limitation. Pet. 27–31. Specifically, Juniper argues that “Leong’s processing units proactively seek out matching tasks,” “retrieve the task’s data for processing,” “execut[e] (i.e., process[]) the task,” “generat[e] first resulting data,” and “change the task status information to ‘complete.’” *Id.* at 29–30 (citing Ex. 1005, 3:55–61, 5:36–45, 6:21–24, 7:34–52, 7:53–61, 8:25–30, Fig. 7C; Ex. 1003 ¶¶ 96–98). Juniper further argues that Leong’s processing units are “independent and autonomous co-processors” because there is no “central management and/or control unit” and the co-processors “perform all these functions without communicating with the specialized surveying agents (i.e., controller).” *Id.* at 27–28, 31 (citing Ex. 1005, 2:45–48, Figs. 7A–7C, 9; Ex. 1003 ¶ 93, 100).

Independent claim 1 further recites “a second co-processor configured to successively: proactively retrieve a second task from the task pool; process the second task; generate second resulting data; and update the task pool to reflect completion of the second task, all without any communication between the second co-processor and the controller.” Ex. 1001, 14:34–39.

¹⁰ Juniper uses color to denote various elements. All coloring has been removed from the quotes.

Juniper argues Leong teaches that limitation. Pet. 31. Specifically, Juniper relies on essentially the same arguments and evidence described above regarding the “first co-processor” limitation to explain how Leong teaches the “second co-processor” limitation. *Id.*

Independent claim 1 further recites “the plurality of first tasks and the plurality of second tasks are associated with a common objective.” Ex. 1001, 14:45–46. Juniper argues Leong teaches this limitation. Pet. 38–39. Specifically, Juniper contends that each of the tasks posted in Leong’s task pool “is a ‘unit[] of work’ associated with a ‘common objective’ because they’re derived from the ‘overall work to be performed.’” *Id.* at 38 (quoting Ex. 1005, 2:50–64) (citing Ex. 1003 ¶ 117).

Independent claim 1 further recites “the first and second co-processors autonomously work together in solidarity with the task pool to complete the common objective.” Ex. 1001, 14:47–49. Juniper argues Leong teaches this limitation. Pet. 40–41. Specifically, Juniper contends that “Leong’s processing units cooperate to complete the larger task from which the bulletin board tasks are created by autonomously retrieving and completing the tasks.” *Id.* at 40 (citing Ex. 1005, 2:54–64, 3:6–18, 3:26–4:3, 4:18–35, 6:47–61, 7:34–50, Figs. 7A, 7C; Ex. 1003 ¶ 124). Juniper relies on essentially the same arguments and evidence described above regarding the “collaborative intelligence system,” “first co-processor,” “second co-processor,” and “common objective” limitations to explain how Leong teaches the “in solidarity” limitation. *Id.*


After reviewing Juniper’s arguments and supporting evidence, including the Weissman Declaration (Ex. 1003), which Swarm does not address at this preliminary stage (*see* Prelim. Resp.), we are persuaded that Juniper sufficiently demonstrates that the combined teachings of Leong and

AppleTalkBook account for all of the undisputed limitations recited in independent claim 1, and that a person having ordinary skill in the art would have had sufficient reasons to combine the teachings of Leong and AppleTalkBook with a reasonable expectation of success.

b) “A Controller Configured to Populate the Task Pool with a Plurality of First Tasks and a Plurality of Second Tasks”

(1) The Parties’ Arguments

Independent claim 1 further recites “a controller configured to populate the task pool with a plurality of first tasks and a plurality of second tasks.” Ex. 1001, 14:26–27. Juniper argues Leong teaches this limitation. Pet. 23–27. Specifically, Juniper argues that “Leong’s processing agents can specialize and ‘survey’ the queue in memory to determine whether new data has arrived” and, “when new data is added to the memory queue 102, the surveying agent examines it, determines which tasks need to be performed to accomplish the overall project 108, defines the status information for each task 108 . . . , and posts status information to the task bulletin board 110.” *Id.* at 24–25 (quoting Ex. 1005, 4:9–25) (citing Ex. 1005, 4:38–48, 5:52–55, 6:36–46, 9:1–4, 9:42–45, Figs. 5, 6; Ex. 1003 ¶¶ 86, 87). Specifically, as shown in an annotated version of Leong’s Figure 2, reproduced below, Juniper identifies “processing unit 1” as the surveying unit and maps that processor to the recited “controller.” *Id.* at 24.



Leong FIG. 2 (Annotated)

Id. The annotated version of Leong’s Figure 2, reproduced above, “is a block diagram of a multi-tasking computing system.” Ex. 1005, 2:14–16. Juniper annotated Figure 2 to identify the bulletin board in red as part of memory 14 and the various processing units in yellow. Pet. 24. Juniper also identifies “Processing Unit 1” as the “surveying unit.” *Id.* According to Juniper, the surveying unit posts the plurality of tasks to the bulletin board. *Id.* at 23–27.

In response, Swarm argues that “Leong accomplishes [its] goal by completely removing the Central Managing Unit from the computing system.” PO Resp. 28 (citing Ex. 1005, Fig. 2, 2:47–48). According to Swarm, “[w]hen read in its proper light, Leong not only fails to teach a controller of the type found in the ’275 patent, but Leong expressly teaches away from incorporating such a controller into the claimed system.” *Id.* at 28.

(2) Our Analysis

Although Leong teaches away from using a centralized controller, it does not teach away from a processor performing some of the tasks that a centralized controller would perform. On the contrary, as Juniper argues,

Leong explicitly teaches that one or more of Leong's processors—which Juniper refers to as the “surveying unit”—will determine when new data has arrived, determine what tasks need to be performed, and then populate the bulletin board—which is mapped to the task pool—with those tasks:

It is preferred that one or more micro-processing units 12 be predisposed as agents to “survey” the queue in memory 14 to determine whether new data has arrived. When new data has arrived in the memory 14, the surveying agents 12 determine what task(s) need to be performed on the data and define the status information for those tasks. The surveying agents 12 then post the status information for the tasks on the bulletin board 14a.

Ex. 1005, 4:9–16. Thus, although Leong teaches away from a centralized controller that manages all aspects of the tasks, it explicitly teaches using that surveying agent to perform the tasks required by the “controller” in independent claim 1. That is, based on the preliminary record, we are sufficiently persuaded that Leong's first processor—which Juniper calls the “surveying agent”—performs the functions of the claimed “controller.”

Swarm's argument is based on semantics: because Leong does not have a processor called a controller, then Leong does not teach the “controller” recited in independent claim 1. Based on the preliminary record, we are not persuaded by this argument. There is no requirement in an obviousness analysis for the prior art to “contain a description of the subject matter of the appealed claim in *ipsissimis verbis*.” *In re May*, 574 F.2d 1082, 1090 (CCPA 1978).

Unlike in IPR2021-01317, Juniper did not map two structural limitations to the same structure. As shown in the annotated version of Leong's Figure 2, reproduced previously, Juniper identifies a specific processor—processing unit 1—as performing the task of the claimed “controller.” Pet. 24. Juniper further identifies the processing units

generically as the first and second co-processors. *Id.* at 27–31. Thus, processing unit 2 can be considered the first co-processor and processing unit n as the second co-processor. *See* Pet. 24 (reproducing Leong’s Figure 2). Based on the preliminary record, Juniper has identified different structural elements for each of the structural limitations—including the “controller” and “first and second co-processors”—recited in independent claim 1.

Accordingly, based on the preliminary record, Juniper presents sufficient evidence that would support a finding that Leong teaches “a controller configured to populate the task pool with a plurality of first tasks and a plurality of second tasks,” as recited in independent claim 1.

c) *“Wherein the Collaborative Intelligence System Is Configured to Dynamically Accept the First Co-Processor, the Second Co-Processor, and an Additional Co-Processor into the Processing System on a Plug-and-Play Basis Without Any Communication with the Controller”*

(1) The Parties’ Arguments

Independent claim 1 further recites “wherein the processing system is configured to dynamically accept the first co-processor, the second co-processor, and an additional co-processor into the processing system on a plug-and-play basis without any communication with the controller.”

Ex. 1001, 14:40–44. Juniper argues that the combined teachings of Leong and AppleTalkBook account for this limitation. Pet. 32–38. Specifically, Juniper argues Leong teaches the need for a “readily scalable” computer system in which the number of processing units available to execute a given task changes. *Id.* at 32–33 (quoting Ex. 1005, 1:49–52) (citing Ex. 1005, 1:44–48, 4:30–5:11, 5:26–35, 6:62–7:34; Ex. 1003 ¶¶ 105–07). According to Juniper, a person having ordinary skill “would’ve investigated techniques

to introduce new processing units to Leong’s bulletin board with minimal configuration, even if the new units have different hardware or software configurations.” *Id.* at 33 (citing Ex. 1005, 6:4–14; Ex. 1003 ¶ 108).

Juniper further argues that a person having ordinary skill in the art would have considered using plug-and-play technology like that taught in AppleTalkBook “to avoid the need for manual configuration when adding new processing units.” Pet. 34–35 (citing Ex. 1003 ¶ 111); *see also* Ex. 1006, 104–105, 108–109 (discussing how AppleTalkBook works); Pet. 34–38 (same). According to Juniper, implementing AppleTalkBook’s plug-and-play teachings into Leong’s system creates a plug-and-play system in which “a new processing unit in Leong’s system would independently discover the bulletin board and obtain work.” Pet. 35 (citing Ex. 1003 ¶ 111). Juniper further argues that applying AppleTalkBook’s plug-and-play teachings to Leong’s standard Ethernet network “would’ve been an exercise of routine skill by a [person of ordinary skill in the art] with a reasonable expectation of success.” *Id.* at 36–37 (citing Ex. 1005, 5:21–25; Ex. 1006, 126–30; Ex. 1003 ¶ 113).

In response, Swarm argues that “AppleTalk fails to teach, and indeed explicitly teaches away from ‘dynamically accept[ing] the . . . co-processor into the processing system on a plug-and-play basis without any communication with the controller.’” Prelim. Resp. 33. Particularly, Swarm contends that “to implement the AppleTalk protocol and enable devices to set their own unique addresses (node numbers) and identify available resources on the AppleTalk network, the AppleTalk protocol relies on a special packet of information known as the ‘broadcast packet,’” which “doesn’t travel to one particular node, but is sent to every node on the network.” *Id.* at 31 (emphasis omitted) (citing Ex. 1006, 108). Swarm

further argues that combining the teachings of AppleTalkBook with those of Leong would explicitly teach away from the “without any communication with the controller” limitation because the broadcast packets would be sent to each processing unit. *Id.* at 31–32.

(2) *Our Analysis*

AppleTalkBook teaches the limitation of “dynamically accepting . . . co-processor[s] . . . without any communication with the controller.” *See* Ex. 1006, 108–09. The Federal Circuit has explained that “a reference need not state a feature’s absence in order to disclose a negative limitation.” *AC Techs., S.A. v. Amazon.com, Inc.*, 912 F.3d 1358, 1367 (Fed. Cir. 2019) (citing *Sud-Chemie, Inc. v. Multisorb Techs., Inc.*, 554 F.3d 1001, 1004–05 (Fed. Cir. 2009) (affirming finding that reference disclosed “uncoated” film where it did not describe the film as coated and did not suggest necessity of coatings)). Here, AppleTalkBook need not disclose that there is no communication between the new co-processor and the controller to teach the negative limitation because it does not suggest the necessity of those communications. AppleTalkBook supports that assertion by teaching that new co-processors are added to a system “without a system administrator assigning an address for the [resource] and distributing it to the network.” Ex. 1006, 108–09.

Even when the broadcast packets are sent to every node on the network, AppleTalkBook does not require communication with a controller for the system to accept the new co-processor into the processing system. *See* Ex. 1006, 108–09. AppleTalkBook teaches a system requiring only communication between the new co-processor and Leong’s bulletin board (i.e., task pool) to accept the new co-processor into the system. *Id.* (explaining that the system accepts new resources “without a system

administrator” because a user can ask nodes to identify themselves and, “if appropriate, respond[] to” the broadcast packet); *see also* Ex. 1005, 2:42–48, 6:4–14, 7:35–52. Thus, AppleTalkBook explains that a resource would “only ha[ve] to be plugged in, attached to the net, and turned on” for a new co-processor to identify the bulletin board, retrieve a task from the bulletin board, process the task, generate resulting data, and update the task pool. *See* Ex. 1006, 108–09.

Accordingly, based on the preliminary record, Juniper presents sufficient evidence that would support a finding that the combined teachings of Leong and AppleTalkBook account for “wherein the processing system is configured to dynamically accept the first co-processor, the second co-processor, and an additional co-processor into the processing system on a plug-and-play basis without any communication with the controller,” as recited in independent claim 1.

d) Conclusion

For the reasons set forth above, Juniper has demonstrated a reasonable likelihood of prevailing on its assertion that the subject matter of independent claim 1 would have been obvious over the combined teachings of Leong and AppleTalkBook.

4. Claims 2–12

Because Juniper has demonstrated a reasonable likelihood of success in proving that at least one claim of the ’275 patent is unpatentable, we institute on all challenged claims and all grounds raised in the Petition. *See* 37 C.F.R. 42.108(a). Therefore, at this preliminary stage, it is not necessary for us to provide an assessment of every challenge raised by Juniper, especially as Swarm has not presented any responsive argument beyond those addressed above.

Nevertheless, we note that Juniper provides detailed explanations supported by the testimony of Dr. Weissman and specific citations to the relevant references indicating where in each reference Juniper argues the limitations of claims 2–7 and 9–17 are taught, and why a person having ordinary skill in the art would have combined the various teachings of Leong and AppleTalkBook with a reasonable expectation of success. *See* Pet. 41–78. Accordingly, we are persuaded the information presented in the Petition establishes there is a reasonable likelihood that Juniper would prevail on its assertion that the subject matter of claims 2–7 and 9–17 would have been obvious over the combined teachings of Leong and AppleTalkBook.

I. Asserted Obviousness in View of Leong, AppleTalkBook, and Ethernet Specification

Juniper argues that claims 1–7 and 9–17 would have been obvious over the combined teachings of Leong, AppleTalkBook, and Ethernet Specification. *See* Pet. 79–81. Based on the preliminary record, we are sufficiently persuaded that Juniper has established a reasonable likelihood of prevailing on this asserted obviousness ground with respect to at least one of claims 1–7 and 9–17.

1. Overview of Ethernet Specification

Ethernet Specification “contains the specification of the Ethernet, a local area network.” Ex. 1007, i. Ethernet Specification “is intended as a design reference document, rather than an introduction or tutorial.” *Id.*

Relevant to Juniper’s argument, chapter 6 provides the “Ethernet Data Link Layer Specification.” Ex. 1007, 19–44. Figure 6-1 is reproduced below.




Figure 6-1: Data Link Layer Frame Format

Ex. 1007, 20. Figure 6-1, reproduced above, depicts a “Data Link Layer Frame Format,” which includes a destination, source, type, data, and frame check sequence. *Id.*

2. Analysis of Claims 1–7 and 9–17

As discussed above, because Juniper has demonstrated a reasonable likelihood of success in demonstrating that at least one claim of the ’275 patent is unpatentable, we institute on all challenged claims and all grounds raised in the Petition. Therefore, at this preliminary stage, it is not necessary for us to provide an assessment of every challenge raised by Juniper, especially as Swarm has not presented any additional argument not previously discussed. *See Prelim. Resp.*

Nevertheless, we note that Juniper provides detailed explanations supported by the testimony of Dr. Weissman and specific citations to the

relevant references indicating where in each reference Juniper argues the limitations of claims 1–7 and 9–17 are taught, and why a person having ordinary skill in the art would have combined the various teachings of Leong, AppleTalkBook, and Ethernet Specification with a reasonable expectation of success. *See* Pet. 19–81. Accordingly, we are persuaded the information presented in the Petition establishes there is a reasonable likelihood that Juniper would prevail on its assertion that the subject matter of claims 1–7 and 9–17 would have been obvious over the combined teachings of Leong, AppleTalkBook, and Ethernet Specification.

J. Asserted Obviousness in View of (1) Leong, AppleTalkBook, and Datla and (2) Leong, AppleTalkBook, Ethernet Specification, and Datla

Juniper argues that claim 16 would have been obvious over the combined teachings of (1) Leong, AppleTalkBook, and Datla and (2) Leong, AppleTalkBook, Ethernet Specification, and Datla. *See* Pet. 82–86. Based on the preliminary record, we are sufficiently persuaded that Juniper has established a reasonable likelihood of prevailing on both asserted obviousness grounds with respect to dependent claim 16.

1. Overview of Datla

Datla is titled “Wireless Distributed Computing: A Survey of Research Challenges.” Ex. 1019, 2. Datla aims “to quantify the benefits of distributed processing over local processing, extend traditional distributed computing (DC) approaches to allow operation in dynamic radio environments, and meet design and implementation challenges unique to [wireless distributed computing].” *Id.*

Relevant to Juniper’s argument, Datla discloses how “[w]ireless distributed computing (WDC) exploits wireless connectivity to share

processing-intensive tasks among multiple devices.” Ex. 1019, 2. For example, “[t]hese technologies can offer capabilities such as collaboration among wirelessly connected computational devices, flexible link design, and adaptation and autonomous operation in dynamic [wireless distributed computing] environments.” *Id.*

2. Analysis of Claim 16

As discussed above, because Juniper has demonstrated a reasonable likelihood of success in demonstrating that at least one claim of the ’275 patent is unpatentable, we institute on all challenged claims and all grounds raised in the Petition. Therefore, at this preliminary stage, it is not necessary for us to provide an assessment of every challenge raised by Juniper, especially as Swarm has not presented any responsive argument directed to the additional limitation set forth in dependent claim 16. *See* Prelim. Resp.

Nevertheless, we note that Juniper provides detailed explanations supported by the testimony of Dr. Weissman and specific citations to the relevant references indicating where in each reference Juniper argues the limitation of dependent claim 16 is taught, and why a person having ordinary skill in the art would have combined the various teachings of Leong, AppleTalkBook, and Datla, both with and without Ethernet Specification, with a reasonable expectation of success. *See* Pet. 82–86. Accordingly, we are persuaded the information presented in the Petition establishes there is a reasonable likelihood that Juniper would prevail on its assertion that the subject matter of dependent claim 16 would have been obvious over both the combined teachings of (1) Leong, AppleTalkBook, and Datla and (2) Leong, AppleTalkBook, Ethernet Specification, and Datla.

III. CONCLUSION

Following § 314, we have determined that the totality of the information presented at this preliminary stage shows there is a reasonable likelihood that Juniper would prevail with respect to at least one of the challenged claims, and we institute on all challenged claims and all grounds raised in the Petition.

Our factual findings, conclusions of law, and determinations at this stage of the proceeding are preliminary and based on the evidentiary record developed thus far. This is not a final decision as to the patentability of claims for which IPR is instituted. Our final decision will be based on the record as fully developed during trial.

IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that an IPR of all challenged claims of the '275 patent is instituted with respect to all grounds raised in the Petition; and

FURTHER ORDERED that pursuant to 35 U.S.C. § 314(a), an IPR of the '275 patent is hereby instituted commencing on the entry date of this Decision, and pursuant to 35 U.S.C. § 314(c) and 37 C.F.R. § 42.4, notice is hereby given of the institution of a trial.

IPR2022-00141
Patent 10,592,275 B2

FOR PETITIONER:

Joseph F. Edell
Kyle K. Tsui
FISCH SIGLER LLP
Joe.Edell.IPR@fischllp.com
kyle.tsui@fischllp.com

FOR PATENT OWNER:

Daniel Pote
JENNINGS STROUSS & SALMON PLC
p@jsslaw.com
dpote@hsslaw.com

Michael K. Kelly
Daniel J. Anderson
BEUS GILBERT MCGRODER PLLC
mkelly@beusgilbert.com
danderson@beusgilbert.com