Trials@uspto.gov
571-272-7822

Paper 51
Entered: March 26, 2018

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

UNIFIED PATENTS INC.,
Petitioner,

v.

INTELLECTUAL VENTURES I LLC,
Patent Owner.

Case IPR2016-01643
Patent 6,775,745 B1

Before JEFFREY S. SMITH, GEORGIANNA BRADEN, and DANIEL J. GALLIGAN, *Administrative Patent Judges*.

SMITH, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318(a) and 37 C.F.R. § 42.73

I. INTRODUCTION

We have jurisdiction under 35 U.S.C. § 6. This Final Written Decision issues pursuant to 35 U.S.C. § 318(a). For the reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 1, 2, 4, 12, and 14 of U.S. Patent No. 6,775,745 B1 (Ex. 1001, "the '745 patent") are unpatentable. Additionally, we dismiss as moot the Parties' Motions to Exclude and deny Patent Owner's Motion to Strike Petitioner's Reply.

A. Procedural History

Unified Patents Inc. ("Petitioner") filed a Petition for *inter partes* review of claims 1–4, 6–9, and 11–17 of the '745 patent. Paper 2 ("Pet."). Intellectual Ventures I LLC ("Patent Owner") filed a Preliminary Response. Paper 7 ("Prelim. Resp."). Pursuant to 37 C.F.R. §§ 42.4(a) and 42.108 and 35 U.S.C. § 314(a), the Board instituted an *inter partes* review of: (1) claims 1, 4, 12, and 14 as anticipated by Karedla[1] under 35 U.S.C. § 102(b)[2]; (2) claims 1, 4, and 12 as unpatentable under 35 U.S.C. § 103(a) in view of Burton[3] and Karedla; (3) claim 2 as unpatentable under 35 U.S.C. § 103(a) in view of Karedla and Otterness[4]; and (4) claim 2 as unpatentable under 35

[1] Ramakrishna Karedla et al., *Caching Strategies to Improve Disk System Performance*, Computer, Vol. 27, No. 3, 38–46 (March 1994) (Ex. 1004, "Karedla").
[2] The Leahy-Smith America Invents Act ("AIA"), Pub. L. No. 112-29, 125 Stat. 284, 287–88 (2011), revised 35 U.S.C. §§ 102 and 103, effective March 16, 2013. Because the challenged patent was filed before March 16, 2013, we refer to the pre-AIA versions of §§ 102 and 103 in this decision.
[3] U.S. Patent No. 6,738,865, issued May 18, 2004 (Ex. 1006, "Burton").
[4] U.S. Patent No. 6,460,122, issued Oct. 1, 2002 (Ex. 1008, "Otterness").

U.S.C. § 103(a) in view of Burton, Karedla, and Otterness. *See* Paper 9, 56 ("Dec. on Inst.").

After institution of trial, Patent Owner filed a Patent Owner Response (Paper 20, "PO Resp."), to which Petitioner filed a Reply (Paper 25, "Reply"). In addition, Patent Owner filed objections to Evidence in Petitioner's Reply (Papers 22, 27), a Motion to Exclude Evidence (Paper 36), and a Motion to Strike Petitioner's Reply (Paper 32). Petitioner opposed the Motion to Exclude (Paper 43) and the Motion to Strike (Paper 40), and filed a Contingent Motion to Exclude Evidence (Paper 37). Patent Owner submitted a Reply in support of its Motion to Exclude (Paper 46) and an opposition to Petitioner's Contingent Motion to Exclude Evidence (Paper 42). Petitioner submitted a Reply in support of its Motion to Exclude (Paper 47).

An oral argument was held on January 4, 2018. A transcript of the oral argument is included in the record. Paper 50 ("Tr.").

B. Related Matters

The parties identify the following district court cases in which the '745 patent has been asserted: *Intellectual Ventures I, LLC v. NetApp, Inc.*, No. 1:16-cv-10868-IT (D. Mass); and *Intellectual Ventures I, LLC v. Lenovo Grp. Ltd.*, No. 1:16-cv-10860-IT (D. Mass). Pet. 2; Paper 5, 1. The '745 patent is also at issue in IPR2017-00429.

C. The '745 Patent

The '745 patent is directed to a hybrid data caching mechanism. Ex. 1001, 1:6–10. The '745 patent explains as background that disk input/output (I/O) can take significant amounts of time while searching for and loading programs and data. *Id.* at 1:17–20. Accessing a file requires first accessing the directory, which keeps track of file locations, and then accessing the file itself. *Id.* at 1:28–31. "Multiple seeks are typically required to locate a file in the directory, especially if the file is not in the first part of the root directory." *Id.* at 1:31–33. Also,

> When more than one program is requesting data, the file system driver can end up reading a small amount of data for a first program, then seek to a different area on the disk to read another small amount for a second program, then seek back to the original area to read the next small amount of data for the first program, and so forth.

Id. at 1:42–48. To circumvent these problems, file system drivers do a certain amount of caching. *Id.* at 2:1–3. For example, "caching programs strictly use the MRU-LRU (most recently used-least recently used) mechanism as their sole means of deciding what data to keep and what data to discard." *Id.* at 2:14–17. In using the MRU-LRU approach, "an important file that has not been used recently will be discarded when the cache reaches a target capacity and is forced to free up additional capacity." *Id.* at 2:17–20. "Strictly adhering to the MRU-LRU mechanism fails to ensure that files which have been used often, but may not have been used recently, are maintained in the cache." *Id.* at 2:20–23.

The '745 patent discloses a caching mechanism that "minimizes seeks and reads to a hard drive and which keeps data in the cache based upon currency of use and the number of times the data is used." *Id.* at 2:31–35. The patent states "system performance is enhanced through the hybrid caching mechanism where a file that has been used often but not recently is maintained in the cache." *Id.* at 3:48–50.

Specifically, the ’745 patent discloses identifying “a least frequently and least recently used file” and eliminating that file to liberate capacity of the cache. *Id.* at 2:50–53. Figure 2A of the ’745 patent is shown below:

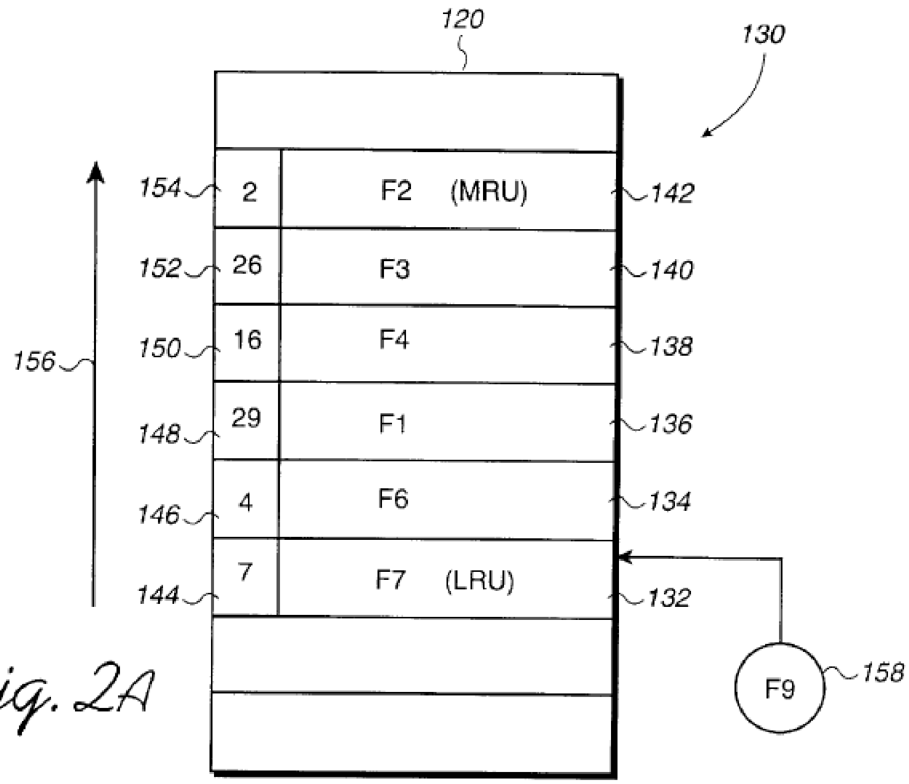

Fig. 2A

Figure 2A above depicts a structure of a cache according to an embodiment of the invention. *Id.* at 5:53–55. Cache 120 includes files F1 through F7. *Id.* at 5:55–56. File F7 132 is the least recently used (LRU) file. *Id.* at 5:56–57. In other words, file F7 132 has gone the longest time as compared to the other files in the cache without being used by the operating system. *Id.* at 5:57–60. File F2 142 is the most recently used (MRU) file. *Id.* at 5:63. A frequency factor is assigned to each file in the cache. *Id.* at 5:65–66. For example, as illustrated above, file F7 has a frequency factor of seven, file F6 has a frequency factor of four, file F1 has a frequency factor of twenty-nine, and so on. *Id.* at 6:4–8. In one embodiment, the frequency factor is increased with each use of the file; in another embodiment, the frequency factor is decreased with the lack of use of the corresponding file. *Id.* at 5:66–6:4.

The ’745 patent discloses that once the target capacity of the cache is reached, the frequency factors are scanned, beginning (in one embodiment) with the frequency factor for the LRU file and ending with the MRU file. *Id.* at 6:8–12. The scanning of the frequency factors results in “a least frequently and least recently used file” being identified. *Id.* at 2:50–51, 12:41–44, 13:1. In one embodiment, the scanning of the frequency factors “determines which file . . . is the least frequently used (LFU) file which has been least recently used.” *Id.* at 6:18–21. The ’745 patent discloses an embodiment in which “the frequency factor for the MRU file is not considered when determining the LFU file since the MRU file has been so recently placed in the cache and has not had an opportunity to be reused.” *Id.* at 6:39–43. In the example illustrated in Figure 2A above, the scan will find file F6 134 as the LRU file with the lowest frequency factor—i.e., the LFU file that has been least recently used. *Id.* at 6:46–49. A new file to be placed in the cache, such as file F9 158, will replace file F6 134. *Id.* at 6:52–55.

In one embodiment, the caching mechanism described above is self-tuning. *Id*. at 6:56–57. For example, after a user finishes watching a movie on the system and then performs word processing operations, the cache will have frequently used movie files that will not be used again. *Id*. at 6:57–63. The caching mechanism will decrement the frequency factor associated with the files by a factor of one or greater over an arbitrary time period so that the files are removed from the cache. *Id*. at 6:53–7:3. Other tuning features include how often to decrement the frequency factors, how far back to search from the LRU file, how far forward to search towards the LRU file, how much to decrement the factors by, and how to assign weight certain files with frequency factors upon their initial read into the cache to guarantee the certain files remain above others. *Id*. at 7:3–15.

The self-tuning aspects also include embodiments to adjust the frequency factors if "the least frequently and least recently used file" is in a defined proximity to the MRU file so that a file other than the MRU is identified as the least frequently and least recently used file. *Id*. at 7:16–24. In one embodiment, only the frequency factors associated with the files proximate to the LRU are decremented. *Id*. at 7:24–35. In one embodiment, if the least frequently and least recently used file is in the top 1/3 i.e., the 1/3 of the cache unit logically proximate to the MRU file, then "the least frequently and least recently used file" is not chosen from the top 1/3 of the cache unit. *Id*. at 7:36–43. In one embodiment, the frequency factors associated with files in the bottom 1/3 of the unit are decremented until a "least frequently and least recently used file" is identified in the bottom 1/3 of the cache unit. *Id*. at 7:43–47. "In one embodiment, the user is provided an application to customize the tuning to their particular situation." *Id*. at 7:11–13. The self-tuning features can be optimized by the user through a graphical user interface. *Id*. at 7:49–51.

The '745 patent describes implementing the improved caching mechanism "in conjunction with reading extending segments of data." *Id*. at 4:32–36. The '745 patent states that "large reads (greater than 64 Kbyte) eliminate the overhead, i.e., seek times, rotational latencies, transfer times, etc., associated with performing two, four or more reads at 32 Kbytes or less as performed in the prior art." *Id.* at 5:38–42. For example, the '745 patent describes that "when the first block of the file is read[,] an additional 64, 128, or 256 Kbytes of data are read with it." *Id.* at 11:22–24. The '745 patent states that "the reads of the extended segments allows for the minimization of seeks and reads from the storage medium, since the files are transferred to cache upon the initial read." *Id.* at 11:45–48.

D. Illustrative Claim

Among the challenged claims (claims 1, 2, 4, 12, and 14), claims 1, 4, and 12 are independent. Claim 4 is illustrative of the subject matter of the challenged claims and reads as follows:

> 4. A caching method for enhancing system performance of a computer, comprising:
>
> reading an extended segment of data in response to a request from an operating system;
>
> storing copies of files associated with the extended segment in a cache;
>
> assigning frequency factors to each of the files stored in the cache, the frequency factors indicating how often each of the corresponding files are requested by the operating system;
>
> scanning the frequency factors, the scanning being performed in response to a target capacity of the cache being attained;
>
> identifying a least frequently and least recently used file; and
>
> eliminating the least frequently and least recently used file to liberate capacity of the cache.

Id. at 12:54–13:4.

II. DISCUSSION

A. 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 v. John Deere Co. of Kansas City*, 383 U.S. 1, 17 (1966). "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 of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the time of the invention. *In re GPAC, Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995). The level of ordinary skill in the art may be reflected by the prior art of record. *Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001). Factors that may be considered in determining the level of ordinary skill in the art include, but are not limited to, the types of problems encountered in the art, the sophistication of the technology, and educational level of active workers in the field. *GPAC*, 57 F.3d at 1579. In a given case, one or more factors may predominate. *Id.* Generally, it is easier to establish obviousness under a higher level of ordinary skill in the art. *Innovention Toys, LLC v. MGA Entm't, Inc.*, 637 F.3d 1314, 1323 (Fed. Cir. 2011) ("A less sophisticated level of skill generally favors a determination of nonobviousness . . . while a higher level of skill favors the reverse.").

Citing the testimony of Dr. Paul Franzon ("Dr. Franzon"), Petitioner contends that a person of ordinary skill in the art at the relevant time "would have had (i) a B.S. degree in electrical engineering, computer engineering, computer science, or equivalent training, or (ii) approximately two years of experience or research related to computer systems." Pet. 11 (citing Ex. 1002 ¶ 44). Although Patent Owner does not propose its own skill level, Patent Owner's declarant, Dr. Frederic T. Chong ("Dr. Chong"), testifies "Dr. Franzon's formulation of the field of the invention and the level of ordinary skill in the art is sufficiently accurate to resolve all of the issues I analyzed and opined on in this declaration." Ex. 2006 ¶ 23. Based on the evidence of record, including the testimony of the parties' declarants, the subject matter at issue, and the prior art of record, we determine that Petitioner's proposed skill level is appropriate, and we adopt Petitioner's articulation of the level of ordinary skill in the art for our analysis below. Ex. 1002 ¶ 44; Ex. 2006 ¶ 23.

B. Principles of Law

A claim is unpatentable under 35 U.S.C. § 102 if a prior art reference discloses every limitation of the claimed invention, either explicitly or inherently. *Glaxo Inc. v. Novopharm Ltd.,* 52 F.3d 1043, 1047 (Fed. Cir. 1995). If the prior art reference does not expressly set forth a particular element of the claim, that reference still may anticipate if that element is inherent in its disclosure. *In re Robertson*, 169 F.3d 743,745 (Fed. Cir. 1999).

A claim is unpatentable under 35 U.S.C. § 103(a) if "the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains." *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including (1) the scope and content of the prior art;

(2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) objective evidence of nonobviousness, i.e., secondary considerations. *See Graham*, 383 U.S. at 17–18.

"A determination of whether a patent claim is invalid as obvious under § 103 requires consideration of all four *Graham* factors, and it is error to reach a conclusion of obviousness until all those factors are considered." *Apple v. Samsung Elecs. Co., Ltd.*, 839 F.3d 1034, 1048 (Fed. Cir. 2016) (en banc) (citations omitted). "This requirement is in recognition of the fact that each of the *Graham* factors helps inform the ultimate obviousness determination." *Id.*

"In an [*inter partes* review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable." *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016) (citing 35 U.S.C. § 312(a)(3) (requiring *inter partes* review petitions to identify "with particularity . . . the evidence that supports the grounds for the challenge to each claim")). This burden never shifts to Patent Owner. *See Dynamic Drinkware, LLC v. Nat'l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015) (citing *Tech. Licensing Corp. v. Videotek, Inc.*, 545 F.3d 1316, 1326–27 (Fed. Cir. 2008)) (discussing the burden of proof in *inter partes* review). Furthermore, Petitioner cannot satisfy its burden of proving obviousness by employing "mere conclusory statements." *In re Magnum Oil Tools Int'l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016).

Thus, to prevail in an *inter partes* review, Petitioner must explain how the asserted prior art or the proposed combinations of prior art would have rendered the challenged claims unpatentable. At this final stage, we determine whether a preponderance of the evidence of record shows that the challenged claims would have been obvious over the proposed combinations of prior art. We analyze the instituted grounds of unpatentability in accordance with the above-stated principles.

C. Claim Construction

In an *inter partes* review, we construe claim terms in an unexpired patent according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b); *Cuozzo Speed Techs., LLC v. Lee*, 136 S. Ct. 2131, 2144–46 (2016) (upholding the use of the broadest reasonable interpretation standard). Consistent with the broadest reasonable construction, claim terms are presumed to have their ordinary and customary meaning as understood by a person of ordinary skill in the art in the context of the entire patent disclosure. *In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). In addition, the Board may not "construe claims during [an *inter partes* review] so broadly that its constructions are unreasonable under general claim construction principles." *Microsoft Corp. v. Proxyconn, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015) (emphasis omitted), *overruled on other grounds by Aqua Prods., Inc.*, 872 F.3d at 1290. An inventor may provide a meaning for a term that is different from its ordinary meaning by defining the term in the specification with reasonable clarity, deliberateness, and precision. *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994).

Petitioner proposes constructions for the following terms/phrases: "reading extending segment(s) of data"; "files"; "frequency factor(s)"; and "selftuning." Pet. 15–19. Patent Owner only proposes construction of the term "frequency factor." PO Resp. 8–29. For our analysis in this Decision, we determine that only "frequency factor" requires express construction.

The term "frequency factor" (or "frequency factors") is recited in each challenged claim. In the Decision to Institute, the Board determined the broadest reasonable interpretation of "frequency factor" is "an indicator based on frequency of use or access." Dec. on Inst. 14.

Petitioner contends "frequency factor" means "an indicator for distinguishing data based *in part* on its frequency of use." Pet. 18 (emphasis added). Petitioner acknowledges that the '745 patent discloses that the frequency factor reflects how often the file is accessed. *Id.* at 17. Petitioner, however, argues that, according to the '745 patent, "the frequency factors could be 'weighted' based on other considerations, for example, by 'weighting certain files with frequency factors on their initial read into the cache to guarantee they remain above others.'" *Id.* at 17–18 (quoting Ex. 1001, 7:8–10). Petitioner also notes that the specification discloses that the frequency factors for reading directory and file allocation table (FAT) data "would be weighted heavier, i.e., the frequency factor would be incremented by a factor of more than one for each time the directory or FAT data is accessed." *Id.* at 18 (quoting Ex. 1001, 8:67–9:3).

Patent Owner contests Petitioner's position. First, Patent Owner notes claims 4 and 12 recite "frequency factors indicating how often each of the corresponding files are requested by the operating system" and contends the entirety of these "frequency factor" limitations be construed as "frequency factors indicating how often each of the corresponding files are requested by the operating system." PO Resp. 8. Patent Owner argues that the claim language included in the "frequency factor" limitations expressly recites what a "frequency factor" is. *Id.* Patent Owner, relying on testimony of its declarant Dr. Chong, contends a dictionary defines "indicating" as "showing with a fair degree of certainty." PO Resp. 9 (citing Ex. 2006 ¶ 30). According to Patent Owner, the broadest reasonable interpretation of the "frequency factors" limitation of claims 4 and 12 is "frequency factors showing, with a fair degree of certainty, how often each of the corresponding files are requested by the operating system." *Id.* at 9–11 (citing Ex. 2006 ¶¶ 32–35). Patent Owner, relying on testimony of Dr. Chong, also contends that a skilled artisan would understand from the specification of the '745 patent that an important purpose of indicating "how often each of the corresponding files are requested" is to identify "a least frequently used file." *Id.* at 11–13 (citing Ex. 2006 ¶¶ 36–39; Ex. 1001, Abstract, Fig. 2A, 2:45–51, 3:7–16, 5:65–6:28, 10:44–11:2).

We do not agree with Patent Owner. One of the sections of the specification cited by Patent Owner discloses scanning frequency factors to determine "the least frequently used (LFU) file which has been least recently used," and the other cited sections disclose scanning the frequency factors to identify "a least frequently and least recently used file," not to identify a least frequently used file. Ex. 1001, Abstract, 2:50–51, 3:15–16, 6:19–21, 11:1–2). Furthermore, the specification discloses "the least frequently used file *takes into consideration* how often a file has been used through the frequency factor." *See id.* at 6:23–24. But, the disclosed "least frequently used file," which "takes into consideration how often a file has been used," may not be the actual least frequently used file, i.e., may not be the file used less frequently than every other file in the cache.

For example, the self-tuning embodiments of the specification disclose adjusting the frequency factors based on criteria other than frequency of use, such as decrementing frequency factors over time, weighting certain frequency factors in their initial read, or adjusting frequency factors of files in only a portion of the cache, while leaving other frequency factors unchanged. Ex. 1001, 6:56–7:51. Thus, in at least the self-tuning embodiments, the frequency factors may not identify the file in the cache that has been used less frequently than every other file in the cache. This disclosure undermines Patent Owner's first argument.

Patent Owner further contends the disclosure of weighting the frequency factors based on criteria other than frequency of use or access does not broaden the claim language so much that it has no meaningful connection to "how often each of the corresponding files are requested by the operating system." PO Resp. 14–16 (citing Ex. 2006 ¶¶ 41–44). According to Patent Owner, one of ordinary skill in the art would understand the disclosed weighting contemplates modest treatment so that the frequency factor can be used to identify the least frequently used file. *Id.* at17–21 (citing Ex. 2006 ¶¶ 45–49). Patent Owner contends that a reasonable construction of the frequency factors recited in claims 4 and 12 must identify a least frequently used file and must allow for a modest degree of variability from a precise count of how often each one of the corresponding files are requested by the operating system. *Id.* at (citing Ex. 2006 ¶ 50). According to Patent Owner, the proposed construction of "showing, with a fair degree of certainty, how often each of the corresponding files are requested by the operating system" has these characteristics. PO Resp. 21.

The specification of the '745 patent, however, does not limit the effects of the tuning criteria to require the frequency factors to show, with a fair degree of certainty, how often each of the corresponding files were accessed. Rather, the specification describes frequency factors in self-tuning embodiments that are based in part on frequency of use, and based in part on other criteria. *See* Ex. 1001, 6:56–7:51. Therefore, the frequency factors in the self-tuning embodiments may not show with a fair degree of certainty how often corresponding files were requested. For example, the specification discloses decrementing a frequency factor of an LRU file over an arbitrary time period. Ex. 1001, 6:63–7:2; *see also* Ex. 1001, 12:51–54 (claim 3 reciting "providing a driver for decrementing a frequency factor after a time period of non-use of a corresponding file").

As another example (and as argued by Petitioner) the specification allows non-frequency considerations such as file type to weight certain files on their initial read into the cache to guarantee they remain above others. *See* Reply 2–3. The '745 patent describes "weighting certain files with frequency factors on their initial read into the cache to guarantee they remain above others. . . ." Ex. 1001, 7:3–11. Petitioner argues, if the lowest frequency factor in the cache is ten, and a frequency factor of a new file is weighted to guarantee that the new file remains above others, then the frequency factor of the new file is higher than ten, even though the new file was accessed only once. Reply 3 (citing Ex. 1001, 7:8–10; 6:28–31; Ex. 1017, 29:10–20). Petitioner persuasively shows that even though a frequency factor greater than ten does not "show, with a fair degree of certainty," that the corresponding file was accessed once, the result remains the same, namely, that the new file, accessed once, would have a frequency factor greater than ten to remain above the others. *See* Reply 2–3.

We agree with Petitioner that Patent Owner's proposed requirement of "showing with fair degree of certainty how often each of the corresponding files are requested" excludes the self-tuning embodiments disclosed in the specification from the scope of the claims. Reply 2–3 (citing Ex. 1001, 7:8–10; Ex. 1017, 29:10–20). Such a result is inconsistent with the claims themselves. For example, dependent claims 7 and 13, which depend, respectively, from claims 4 and 12, expressly recite that the caching methods are "self-tuning," and, therefore, the scope of the independent claims necessarily encompasses self-tuning. Furthermore, the Federal Circuit has stated "the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which [it] appears, but in the context of the entire patent, including the specification." *Phillips v. AWH Corp.*, 415 F.3d 1303, 1315 (Fed. Cir. 2005). The specification "is always highly relevant to the claim construction analysis. Usually, it is dispositive; it is the single best guide to the meaning of a disputed term." *Id.* at 1315. In particular, a "claim construction that excludes [an] . . . embodiment [described in the specification] is rarely, if ever, correct and would require highly persuasive evidentiary support." *Adams Respiratory Therapeutics, Inc. v. Perrigo Co.*, 616 F.3d 1283, 1290 (Fed. Cir. 2010) (citation omitted). Here, Patent Owner's construction would exclude not only embodiments disclosed in the specification but also those expressly claimed, as noted above. We agree with Petitioner that the intrinsic evidence in the specification is more instructive than extrinsic evidence such as a dictionary definition of "indicate." *See Phillips* at 415 F.3d 1317.

Furthermore, claims 4 and 12 do not limit the frequency factor to indicating a count of the number of times each of the corresponding files are requested. Rather, claims 4 and 12 recite frequency factors "indicating *how often* each of the corresponding files are requested." Thus, we determine that a frequency factor is not limited to indicating a count but encompasses indicating a relative metric such as "how often." *See* Ex. 1017, 39:12–17.

Given the use of the word "frequency" in the claim term, the frequency factors must be based to some extent on frequency of access or use. We determine that the broadest reasonable interpretation of "frequency factor" recited in claims 4 and 12, read in light of the specification, including the self-tuning embodiments, encompasses at least "an indicator based on frequency of use or access."

Patent Owner contends that the broadest reasonable interpretation of the "frequency factor of each of the files reflecting how frequent each of the files is accessed" recited in claim 1 is broader than the "frequency factors" limitation recited in claims 4 and 12. PO Resp. 26–29. Having considered the parties' arguments and the evidence of record, we determine that the broadest reasonable interpretation of "frequency factor" recited in claim 1 encompasses at least "an indicator based on frequency of use or access."

D. Asserted Anticipation by Karedla

Petitioner contends that claims 1, 4, 12, and 14 are unpatentable under 35 U.S.C. § 102(b) as anticipated by Karedla. Pet. 3, 4, 19–39. Relying in part on the testimony of Dr. Franzon, Petitioner explains how Karedla discloses each claim limitation. *Id.* at 19–39 (citing Ex. 1002). Relying on the testimony of Dr. Chong, Patent Owner contests Petitioner's position. PO Resp. 30–47 (citing Ex. 2006).

1. Summary of Karedla

Karedla is an article from the March 1994 issue of *Computer* titled "Caching Strategies to Improve Disk System Performance." Ex. 1004, 1, 6.[5] The cover page and table of contents of Karedla indicated a date of March 1994, and the second page of Karedla has a stamp of the Copyright Office of the Library of Congress dated March 22, 1994. *Id.* at 1, 2, 4.

Karedla characterizes its disclosure as follows:

> In this article, we examine the use of caching as a means to increase system response time and improve the data throughput of the disk subsystem. After explaining some basic caching issues . . . , we examine some popular caching strategies and cache replacement algorithms, as well as the advantages and disadvantages of caching at different levels of the computer system hierarchy. Finally, we investigate the performance of three cache replacement algorithms: random replacement (RR), least recently used (LRU), and a frequency-based variation of LRU known as segmented LRU (SLRU).

Id. at 6.

Karedla discloses that one of the main parameters in I/O cache design is the line replacement algorithm. *Id.* at 7. Karedla explains that vendors attempt to design efficient replacement algorithms that can offer higher cache performance for relatively small cache sizes. *Id.* Karedla describes "some of the more popular cache replacement algorithms," including least recently used (LRU), least frequently used (LFU), and variations on LRU. *Id.* at 8. Karedla describes the LRU algorithm as one that "evicts the cache line used least in the recent past on the assumption that it won't be used in

---

[5] Even though Petitioner cites to the original page numbers of Karedla rather than the exhibit page numbers added by Petitioner, for ease of reference to the entire exhibit, we cite to the exhibit page numbers.

the near future." *Id.* Karedla also states "LRU uses only the time since last access and does not take into account the access frequency of lines when making replacement decisions." *Id.* Karedla describes the LFU algorithm as one that "maintains a frequency-of-access count for all its lines and replaces the line that has been used least frequently." *Id.* Karedla discloses that a disadvantage of the LFU algorithm is that "lines with large frequency counts that will not be accessed again (such as a recently active but currently cold file) tend to remain entrenched in the cache, preventing newer additions to the cache from gathering sufficient reference counts to stay in the cache." *Id.*

Karedla describes variations on the LRU algorithm, stating that a growing trend is toward algorithms that "generally are combinations of LRU, LFU, and read-ahead strategies, with varying thresholds on the amount of data that is prefetched." *Id.* Karedla discloses that "[a] read-ahead strategy known as prefetching exploits the principle of spatial locality in an attempt to minimize latency in data access by anticipating future requests for data and bringing it into the cache." *Id.*

One particular algorithm described in Karedla is "[a] frequency-based LRU replacement algorithm," which Karedla states "improves on LRU by partitioning the LRU stack into three sections whose sizes are tunable." Ex. 1004, 8. Karedla further discloses: "On being referenced, a line is placed in the topmost section of the stack. Unlike LFU, the reference count of lines repeatedly referenced in the top section is not incremented." Ex. 1004, 8.

2. Analysis

a. Karedla as prior art

Petitioner contends that Karedla is prior art to the ’745 patent under 35 U.S.C. § 102(b). Pet. 4, 19 (citing, e.g., Ex. 1004, 2–3; Ex. 1013; Ex. 1012 ¶ 3). Patent Owner does not challenge in its Response that Karedla is prior art under 35 U.S.C. § 102(b). *See generally* PO. Resp.

We determine that Karedla was published in March 1994. *See* Pet. 19; Ex. 1004, 1 (“March 1994”), 2 (Library of Congress stamp dated “MAR 22 1994”), 4 (“March 1994”); Ex. 1013 (printout from Library of Congress website identifying publication history and frequency). Accordingly, we find that Karedla is prior art to the ’745 patent, which issued from an application filed September 7, 2001, under 35 U.S.C. § 102(b).

b. Claim 4

Petitioner contends Karedla discloses each limitation of independent claim 4. Pet. 20–28. As explained more fully below, we agree with Petitioner and find Karedla discloses the subject matter of claim 4, and, in particular, we focus on Karedla’s disclosure of a “frequency-based LRU replacement algorithm.” *See* Ex. 1004, 8.

“reading an extended segment of data”

Claim 4 is directed to “[a] caching method for enhancing system performance” and recites “reading an extended segment of data in response to a request from an operating system.” Petitioner argues that Karedla describes “[a] caching method for enhancing system performance.” Pet. 20–21 (citing Ex. 1004, 6; Ex. 1002 ¶¶ 75–76). We agree because, as its title indicates, Karedla is directed to caching strategies for improving performance. Ex. 1004, 6 (“Caching Strategies to Improve Disk System

Performance"); *see also id.* ("This article describes a caching strategy that offers the performance of caches twice its size.").

Petitioner argues that Karedla discloses the claimed "reading an extended segment of data in response to a request from an operating system" by describing a "read-ahead strategy known as prefetching," which "anticipat[es] future requests for data and bring[s] it into the cache." Pet. 21 (quoting Ex. 1004, 8). In particular, Petitioner contends that Karedla discloses caching strategies, including reading ahead, that exploit the principle that if a block of data is accessed, then nearby blocks of data will also soon be accessed. *Id.* at 21 (citing Ex. 1004, 7–8; Ex. 1002 ¶ 79). Petitioner also contends Karedla discloses that prefetching exploits the principle of spatial locality by anticipating future requests and bringing them into the cache. *Id.* (citing Ex. 1004, 8; Ex. 1002 ¶ 80).

Dr. Franzon testifies that "one of ordinary skill in the art would find *Karedla*'s disclosure of reading ahead to prefetch additional blocks of data and storing the data in larger cache lines satisfies the element 'reading extended segments of data.' Further, it is common to use these strategies with a cache." Ex. 1002 ¶ 82. We credit this testimony, and we find that Karedla discloses "reading an extended segment of data in response to a request from an operating system." For example, Karedla discloses: "A read-ahead strategy known as prefetching exploits the principle of spatial locality in an attempt to minimize latency in data access by anticipating future requests for data and bringing it into the cache. Most disk drives implement prefetch in an onboard cache." Ex. 1004, 8.

"storing copies of files associated with the extended segment"

Claim 4 further recites "storing copies of the files associated with the extended segment in a cache." Petitioner argues that Karedla discloses this subject matter by describing that caches keep "copies of backing store data" so that they "can service some requests at faster memory speeds." *Id*. at 22 (quoting Ex. 1004, 7). In describing "[b]asic caching terminology," Karedla discloses that "[a] *cache buffer* is faster memory used to enhance the performance of a slower memory (a disk drive, for example), known as the backing store. By keeping copies of backing store data, caches can service some requests at faster memory speeds." Ex. 1004, 7. Furthermore, as discussed above, Karedla also discloses "anticipating future requests for data and bringing it into the cache." Ex. 1004, 8. Therefore, we determine Karedla stores files associated with the extended segment in the cache. We, therefore, agree with Petitioner, and we find Karedla discloses "storing copies of the files associated with the extended segment in a cache."

"assigning frequency factors to each of the files in the cache"

Claim 4 further recites "assigning frequency factors to each of the files stored in the cache, the frequency factors indicating how often each of the corresponding files are requested by the operating system." Petitioner notes that, in describing an LFU algorithm, Karedla discloses that "[t]he cache maintains a frequency-of-access count for all its lines". Pet. 23 (quoting Ex. 1004, 8); Ex. 1002 ¶ 84. In describing the "frequency-based LRU replacement algorithm," Karedla discloses that, "[u]nlike LFU, the reference count of lines repeatedly referenced in the top section is not incremented." Ex. 1004, 8.

Patent Owner contends that Karedla's frequency-based LRU replacement algorithm does not disclose "assigning frequency factors to each of the files stored in the cache" because this algorithm is a three-section cache that does not assign frequency factors to files in the top section. PO Resp. 43–44 (citing Ex. 1004, 8; Ex. 2006[6] ¶ 73). The cited section of Karedla discloses, however, "the reference count of lines repeatedly referenced in the top section is not incremented." Ex. 1004, 8. Contrary to Patent Owner's contention, each of the lines, or "files," in the top section has a reference count, or "frequency factor." *Id*. Although Karedla discloses not incrementing the reference count of lines in the top section, the lines in the top section still have a reference count. *Id*. Karedla further discloses: "Eventually, the line ages into the bottom part of the list by LRU replacement and is finally evicted if not referenced." Ex. 1004, 8. The reference counts of lines in the middle and bottom sections are incremented when the files are referenced, upon which those files are "placed in the topmost section of the stack." Ex. 1004, 8.

Patent Owner argues: "Karedla does not keep track of how many times cache lines in the top section of the cache are accessed. Therefore, Karedla does not assign 'frequency factors to *each* of the files stored in the cache.'" PO Resp. 44. Claim 4, however, does not require that the

---

[6] The Patent Owner's Response cites "Ex. 2003 ¶¶ 73." Exhibit 2003 is the transcript of the deposition of Petitioner's declarant, Dr. Franzon, and the transcript does not contain paragraph numbers. Page 73 of the Franzon transcript contains testimony pertaining to the Burton reference. We, therefore, understand Patent Owner's citation to be referring to paragraph 73 of Exhibit 2006, which is the declaration of Dr. Chong, and we have considered this testimony in making our determination.

frequency factor "keep track" of the number of accesses even under Patent Owner's construction of "showing, with a fair degree of certainty." *See* PO Resp. 9 (emphasis omitted). As discussed in the claim construction section above, claim 4 encompasses "self-tuning" embodiments, in which the frequency factor is adjusted so that it does *not* track how many times cache lines are accessed. Ex. 1001, 6:56–7:51. One self-tuning implementation is described in the passage below.

> The caching mechanism described above is self tuning in one embodiment. For example, should a user finish watching a movie on the system and proceed to perform word processing operations, the cache will have files from the movie which were frequently used. Because of the frequent use, these files will have a high frequency factor associated with them even though the files will not be used upon completion of the movie viewing. In one embodiment, the caching mechanism will decrement the frequency factor associated with these files to bring down the frequency factor so that the files are removed from the cache.

Ex. 1001, 6:56–66. Incrementing a frequency factor and then decrementing it results in the frequency factor *not* tracking how many times the cache lines were accessed, as does simply not incrementing the frequency factor, as disclosed in Karedla. Ex. 1004, 8. Claim 4 does not recite a step of updating the frequency factor in any particular manner.

Based on the foregoing, we find that each file in Karedla's frequency-based LRU has a reference count, and we find that the reference count of Karedla is "an indicator based on frequency of use or access," which discloses the claimed frequency factor.

Having considered the parties' arguments and the evidence of record, we agree with Petitioner, and we find Karedla discloses "assigning frequency factors to each of the files stored in the cache, the frequency factors indicating how often each of the corresponding files are requested by the operating system."

"scanning the frequency factors"

Claim 4 further recites "scanning the frequency factors, the scanning being performed in response to a target capacity of the cache being attained." In particular, Petitioner contends this limitation is disclosed by Karedla's description that a "cache replacement algorithm decides which cache line to discard when replacement is required." Pet. 24 (citing Ex. 1004, 7). Petitioner further contends this limitation is disclosed by Karedla's description of the frequency-based LRU algorithm, which partitions the LRU stack into three sections. Pet. 24 (citing Ex. 1004, 8). Petitioner contends that lines not referenced frequently will age into the bottom part of the list by LRU replacement and will be evicted finally if not referenced. *Id*.

Also, Dr. Franzon testifies:

> As I explained above, *Karedla* discloses assigning frequency factors to each of the files stored in the cache. One of ordinary skill in the art would understand that the decision as to which cache line to replace is based on the specific replacement algorithm used, and *Karedla* discloses at least two algorithms similar to the concept disclosed in the '745 patent.
>
> In particular, *Karedla* explains that "*cache replacement algorithm* decides which cache line to discard when replacement is required." [Ex. 1004, 7]. Replacement is needed due to the finite capacity of the cache. "Cache replacement, also known as eviction, is the operation of discarding data from the cache to make room for new data." [Ex.1004, 7]. Thus, in my opinion, the element "the scanning being performed in response to a target capacity of the cache being attained" is satisfied.

Ex. 1002 ¶¶ 87–88. Dr. Franzon further testifies that "one of ordinary skill in the art would understand that in order to identify the LFU and LRU file, the cache mechanism must scan the frequency factors in the cache lines." Ex. 1002 ¶ 91. We credit this testimony because it is consistent with Karedla's disclosure of the very purpose of cache replacement (Ex. 1004, 7), and we find that Karedla discloses "scanning the frequency factors, the scanning being performed in response to a target capacity of the cache being attained."

"identifying a least frequently and least recently used file"

Petitioner argues that Karedla's discloses "identifying a least frequently and least recently used file," as further recited in claim 4. In particular, Petitioner relies on Karedla's disclosure of the frequency-based LRU algorithm in which lines eventually age into the bottom part of the list and are evicted if not referenced again. Pet. 25 (citing Ex. 1004, 8). Karedla discloses that "[e]ventually, the line ages into the bottom part of the list by LRU replacement and is finally evicted if not referenced." Ex. 1004, 8. Dr. Franzon testifies that, in his opinion, "the evicted lines [of Karedla] are those least frequently referenced as well as least frequently used." Ex. 1002 ¶ 93. We agree with Petitioner, and we find that Karedla discloses "identifying a least frequently and least recently used file."

"eliminating the least frequently and least recently used file"

Petitioner, relying on testimony of Dr. Franzon, contends that Karedla discloses "eliminating the least frequently and least recently used file to liberate capacity of the cache." Pet. 28 (citing Ex. 1002 ¶ 99). Dr. Franzon testifies that cache replacement, also known as eviction, is the operation of discarding data from the cache to make room for new data and that Karedla evicts the least frequently and least recently used file to make room for new data. Ex. 1002 ¶ 99 (citing Ex. 1004, 7). We credit this testimony because it is consistent with Karedla's disclosure that cache replacement "is the operation of discarding data from the cache to make room for new data." Ex. 1004, 7. Based on the evidence of record, we find that Karedla discloses "eliminating the least frequently and least recently used file to liberate capacity of the cache."

Alleged mixing of embodiments

Patent Owner contends that one cannot combine elements from different embodiments to support a finding of anticipation. PO Resp. 30 (citing *NetMoneyIN, Inc. v. Verisign, Inc.*, 545 F.3d 1359, 1371 (Fed. Cir. 2008)). Patent Owner asserts that the Board erred as a matter of law by improperly relying on a combination of embodiments disclosed in Karedla to institute on a ground of anticipation. PO Resp. 32–33. According to Patent Owner, the Board's reasoning that "the general statements in Karedla regarding read-ahead strategies are applicable to the teachings regarding variations on LRU and segmented LRU" is too vague to support a finding of anticipation. *Id.* at 33 (citing Dec. on Inst., 19). Patent Owner contends that the Board's reasoning that the read-ahead strategies of Karedla are applicable to the LRU algorithms of Karedla is directed to obviousness. *Id.* at 34. In particular, Patent Owner, relying on testimony of Dr. Chong, contends that the variations on LRU algorithms disclosed by Karedla do not describe using prefetching because the algorithms described in the first two paragraphs of the section in Karedla labeled "Variations on LRU" do not disclose using prefetching. *Id.* at 37–40 (citing Ex. 1004, 8; Ex. 2006 ¶¶ 69, 72, 75, 80, 84, 88).

We do not agree with Patent Owner. "Anticipation occurs when a prior art reference discloses each element of the claimed invention, not only where a particular embodiment within a reference discloses each element of the claimed invention." *In re AT&T Intellectual Property II, L.P.*, 856 F.3d 991, 996 (Fed. Cir. 2017) (citing *Blue Calypso, LLC v. Groupon, Inc.*, 815 F.3d 1331, 1341(Fed. Cir. 2016)). "Instead, a reference may still anticipate if that reference teaches that the disclosed components or functionalities may be combined and one of skill in the art would be able to implement the combination." *Blue Calypso*, 815 F.3d at 1344. Because all the limitations of claim 4 are disclosed in Karedla, the question for the purposes of anticipation is whether using read-ahead strategies with a variation on an LRU algorithm is expressly disclosed by Karedla or, if not, would "be immediately apparent to one of ordinary skill in the art." *Kennametal, Inc. v. Ingersoll Cutting Tool Co.*, 780 F.3d 1376, 1382 (Fed. Cir. 2015).

In the same section in which Karedla discloses the "frequency-based LRU replacement algorithm," Karedla discloses observing "a growing trend among vendors toward algorithms that adapt to changing access patterns. These algorithms generally are combinations of LRU, LFU, and read-ahead strategies, with varying thresholds on the amount of data that is prefetched." Ex. 1004, 8. In one section, therefore, Karedla discloses an algorithm that combines aspects of LFU and LRU algorithms (the "frequency-based LRU replacement algorithm") and also discloses that such algorithms generally are implemented with "read-ahead strategies." Ex. 1004, 8. At the very least, Karedla's observations of vendors trending toward "algorithms [that] are combinations of LRU, LFU, and read-ahead strategies" shows that a person of skill in the art, reading the variations on LRU algorithms of Karedla, would immediately envisage applying the read-ahead strategy to the LRU algorithms. *See Kennametal*, 780 F.3d at 1384.

Patent Owner, relying on testimony of Dr. Chong, contends that Karedla does not disclose "reading an extended segment of data" because Karedla does not disclose the implementation details of cache replacement algorithms that combine read-ahead strategies with LRU and LFU functions. PO Resp. 40 (citing Ex. 2006 ¶¶ 126–127). In particular, Dr. Chong testifies that the "authors are merely disclosing that they are aware of ***other, undisclosed*** algorithms that are 'combinations of LRU, LFU, and read-ahead strategies,' but they do not disclose any additional detail about these algorithms. Ex. 2006 ¶ 127. According to Patent Owner, without implementation details, a skilled artisan would not understand how the read-ahead strategies would be used in such algorithms. PO Resp. 41 (citing Ex. 2006 ¶¶ 123–134).

We disagree with Patent Owner. Here, Karedla discloses a limited number of caching strategies to improve disk system performance. Dr. Franzon testifies that

> one of ordinary skill in the art would find Karedla's disclosure of reading ahead to prefetch additional blocks of data and storing the data in larger cache lines satisfies the element "reading extended segments of data." Further, it is common to use these strategies with a cache. Indeed, Karedla explains that "[replacement] algorithms generally are combinations of LRU, LFU, and read-ahead strategies, with varying thresholds on the amount of data that is prefetched."

Ex. 1002 ¶ 82 (citing Ex. 1004, 8). As discussed above, we determine that, given Karedla's disclosure, using a read-ahead strategy with a variation on an LRU algorithm would be immediately apparent to one of ordinary skill in the art. *See Kennametal*, 780 F.3d at 1384. Also, as discussed above, Dr. Franzon testifies that "one of ordinary skill in the art would find Karedla's disclosure of reading ahead to prefetch additional blocks of data and storing the data in larger cache lines satisfies the element 'reading extended segments of data.' Further, *it is common to use these strategies with a cache*." Ex. 1002 ¶ 82 (emphasis added). Contrary to Dr. Chong's testimony that "a person of ordinary skill in the art would conclude that Karedla's generic disclosure of 'combinations of LRU, LFU, and read-ahead strategies' is wholly separate from Karedla's more specific disclosure of specific cache replacement algorithms" (Ex. 2006 ¶ 128), Dr. Franzon's testimony that it is common to use read-ahead strategies with a cache is consistent with Karedla's discussion of replacement algorithms that "generally are combinations of LRU, LFU, and read-ahead strategies." Ex. 1004, 8. We find that a skilled artisan would at once envisage using a read-ahead strategy with a variation on LRU algorithm, such as the frequency-based LRU algorithm of Karedla.

Having reviewed the entirety of the record, including the testimony of Dr. Franzon and Dr. Chong, we determine that Petitioner has shown by a preponderance of the evidence that claim 4 is unpatentable under 35 U.S.C. § 102(b) as anticipated by Karedla. *See* Pet. 20–28. We also determine, for this ground, that Petitioner's citations to Karedla alone are sufficient and we do not rely on any of Petitioner's citations to the Robinson article or Robinson '885. *See, e.g.*, Pet. 23, 24–25 (citing Exs. 1005 and 1007).

c. Claims 1 and 12

Petitioner asserts that claims 1 and 12 are "substantively identical" to claim 4, noting that claims 1 and 12 are directed to a similar method and media for identifying a least frequently and least recently used file, respectively. Pet. 28. Petitioner asserts that Karedla discloses the feature recited in claims 1 and 12, but not in claim 4, of the cache being located in a random access memory. *Id.* at 28–31 (citing, e.g., Ex. 1004, 7, 9; Ex. 1002 ¶¶ 104, 107). Patent Owner contests Petitioner's position, raising the arguments discussed above in connection with claim 4 also for claims 1 and 12. *See* PO. Resp. 30–46.

Claim 1 is a method claim that recites limitations that are very similar to those of method claim 4, with the additional requirement of "the cache being located in a random access memory (RAM) of the computer." Claim 12 recites "[a] computer readable media having program instructions for a caching method" comprising program instructions for performing functions nearly identical to those recited in claim 4. Like claim 1, claim 12 additionally requires that the cache be located in a random access memory of the computer.

Karedla discloses that "[a] cache buffer is faster memory used to enhance the performance of a slower memory (a disk drive, for example), known as the backing store." Ex 1004, 7. Karedla further discloses that "[a]lthough RAM costs are decreasing, cache memory is still expensive." *Id*. Karedla further discloses a disadvantage of host caching is data volatility (unless battery backed). *Id*. at 9. Dr. Franzon testifies that "[o]ne of ordinary skill in the art would understand that Karedla's description of a 'faster memory' or 'a volatile cache' is referring to a volatile memory, which is a random access memory (RAM) and not read-only memory (ROM)." Ex. 1002 ¶ 107. We rely on this testimony, and we find that Karedla discloses "the cache being located in a random access memory (RAM) of the computer" as recited in claims 1 and 12. For the reasons explained above in connection with claim 4, and because Petitioner shows by a preponderance of the evidence that Karedla discloses that the cache is located in a random access memory, we determine that Petitioner shows that claims 1 and 12 are unpatentable under 35 U.S.C. § 102(b) as anticipated by Karedla. We also determine, for this ground, that Petitioner's citations to Karedla are sufficient and we do not rely on any of Petitioner's citations to the Robinson article. *See, e.g.*, Pet. 29 (citing Ex. 1005).

d. Dependent claim 14

Claim 14 recites that the computer readable media as recited in claim 12 further comprises "program instructions for allowing a user to adjust a size of the extended segment that is read." Petitioner relies on its analysis of claim 4 and additionally relies, for example, on a passage in Karedla's description of read-ahead strategies stating that "vendors also offer a tunable read-ahead threshold to minimize both cache pollution and the latency in transferring large amounts of data from the disk to the I/O bus." Pet. 38–39 (citing Ex. 1002 ¶ 130). Patent Owner does not raise, in its Response, any arguments specific to the limitation added by dependent claim 14. *See* PO. Resp. 30–46. The burden, however, remains on Petitioner to demonstrate unpatentability. *See Dynamic Drinkware*, 800 F.3d at 1378.

Karedla discloses "vendors also offer a tunable read-ahead threshold to minimize both cache pollution and the latency in transferring large amounts of data from the disk to the I/O bus . . . . Also, to avoid cache pollution, vendors offer user-selectable upper bounds on the request size that the cache will process." Ex. 1004, 8. Based on this disclosure, we find Karedla discloses "program instructions for allowing a user to adjust a size of the extended segment that is read." We determine that Petitioner shows, by a preponderance of the evidence, that claim 14 is unpatentable under 35 U.S.C. § 102(b) as anticipated Karedla. *See* Pet. 36–39.

e. Summary

For the above reasons, we determine that Petitioner shows that claims 1, 4, 12, and 14 are unpatentable under 35 U.S.C. § 102(b) as anticipated by Karedla alone.

F. Asserted Obviousness over Burton and Karedla

Petitioner contends that claims 1, 4, and 12 of the ’745 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over Burton and Karedla. Pet. 39–56. Relying in part on the testimony of Dr. Franzon, Petitioner explains how the references allegedly teach or suggest the claim limitations and provides reasoning for combining the teachings of the references. *Id.* (citing Ex. 1002 ¶¶ 133–135, 141, 142, 144–148, 150, 151, 154–158). Relying on the testimony of Dr. Chong, Patent Owner contests Petitioner’s position. PO Resp. 47–60 (citing Ex. 2006).

1. Summary of Burton

Burton is a U.S. patent titled “Method, System, and Program for Demoting Data From Cache Based on Least Recently Accessed and Least Frequently Accessed Data.” Ex. 1006, [54]. Burton issued May 18, 2004 from an application filed June 9, 2000, and is prior art to the ’745 under 35 U.S.C. § 102(e). *Compare id.* [22], [45], *with* Ex. 1001, [22], [45].

Burton describes as background a least recently used (LRU) algorithm for managing data in a cache. Ex. 1006, 1:39–49. Burton explains that when a track is added to the cache, a pointer to the track is placed at the top of the LRU linked list. *Id.* at 1:39–42. Burton explains that when the cache manager determines that data must be demoted or removed from cache to make room for subsequent data accesses, the cache manager will demote tracks having pointers at the bottom of the list, representing those tracks that were accessed the longest time ago relative to other tracks in the cache. *Id.* at 1:44–49. Burton states that "LRU algorithms are not optimal for randomly accessed data" because a demoted track in a cache, even though it is the least recently accessed, may be more likely to be subsequently accessed than those less frequently accessed. *Id.* at 1:50–58.

Burton discloses a caching method that "is an improvement over prior art techniques that do not take into account the frequency of access when demoting cache entries." *Id.* at 2:31–34. According to Burton, "an entry that is more frequently accessed relative to another entry is more likely to be accessed again in the future." *Id.* at 2:26–29. Burton discloses "select[ing] for demotion those entries less likely to be accessed in the future with respect to other entries." *Id.* at 2:29–31.

Figure 1 of Burton illustrates Burton's cache architecture and is shown below:

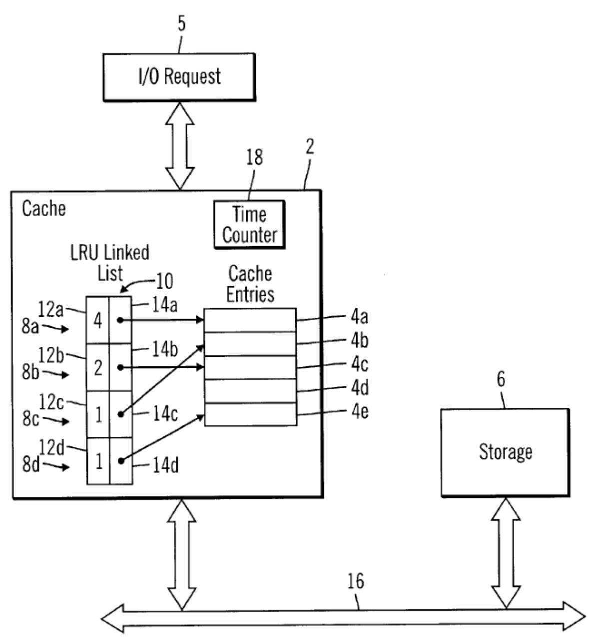

FIG. 1

Id. at 2:55–56. As illustrated in Figure 1 above, cache 2 includes cache entries 4a, 4b, 4c, 4d, into which tracks or pages of data from storage 6 may be placed. *Id.* at 2:55–58. Burton discloses that when a track or page of data is staged into a cache entry, new entry 8a is added to the top of LRU linked list 10. *Id.* at 2:58–60. According to Burton, for each cache entry 4a, 4b, 4c, 4d, LRU linked list 10 includes one entry 8a, 8b, 8c, 8d, respectively. *Id.* at 2:60–61. Burton discloses that each entry in LRU linked list 10 includes LRU rank value 12a, b, c, d and pointer 14 a, b, c, d to one cache entry 4a, b, c, d, respectively. *Id.* at 2:62–66. Burton discloses that when an input/output request is serviced from the cache (i.e., a cache hit), then the entry in the LRU linked list including the pointer "is moved to the top of the LRU linked list." *Id.* at 2:66–3:3.

Burton discloses that the LRU rank "provides a value for each entry in [the cache] that indicates both how frequently the entry is accessed and the time of last access." *Id.* at 3:11–13. Burton states that "[t]he purpose of this weighting is to allow the process [of] demoting entries from cache [] to take into account the frequency with which an entry was accessed and remove those less frequently accessed entries." *Id.* at 3:13–17 (reference numerals omitted). Burton describes its ranking methodology, which uses time counter 18 (shown in Figure 1 above), as follows:

> The . . . LRU rank is weighted for previous accesses. Because the time counter 18 is incremented for every I/O access, the value added to the LRU rank for subsequently accessed cache entries increases substantially in a short period of time. Thus, LRU entries recently accessed or added to cache will have a substantially higher LRU rank than the LRU rank of entries that have not been accessed recently, even those entries accessed numerous times. However, to the extent entries have not been accessed for the same amount of time, the entry that was accessed

> more frequently will have a higher LRU rank because its LRU rank will be weighted with previous accesses. Thus, to the extent entries were last accessed at about the same time, which in terms of cache operations is within fractions of a second of each other, the more frequently accessed entry will have a greater weighting.

Id. at 3:54–4:2. Burton explains that the more frequently accessed entry is given a greater weighting because "a cache entry more frequently accessed in the past is likelier to be accessed in the future over less frequently accessed entries that were last accessed at about the same time." *Id.* at 4:2–5.

Burton also discloses that when the number of cached entries reaches an upper threshold, the cache determines from the last 1024 entries in the LRU linked list thirty-two entries that have the lowest LRU rank. *Id.* at 4:7–11. According to Burton, the cache then demotes those thirty-two entries from the cache. *Id.* at 4:11–13. Burton states that "those entries having the lowest LRU rank marked for demotion are both the least recently accessed and[,] among those entries least recently accessed, are less frequently accessed." *Id.* at 4:19–22. Thus, according to Burton, "the algorithm . . . ensures that more frequently accessed cache entries . . . remain in cache longer." *Id.* at 4:22–24.

2. Analysis

a. Claim 4

Petitioner contends that claim 4 of the '745 patent is unpatentable under 35 U.S.C. § 103(a) as obvious over Burton and Karedla. Pet. 3, 40–50. Relying in part on the testimony of Dr. Franzon, Petitioner explains how the references allegedly teach or suggest the claim limitations and provides reasoning for combining the teachings of the references. *Id.* at 40–50 (citing Ex. 1002 ¶¶ 134, 135, 141, 142, 144–148, 150).

"storing copies of files . . . in a cache"

Petitioner argues that Burton discloses "storing copies of files . . . in a cache" recited in claim 4 in describing a cache architecture having a cache that includes cache entries into which tracks or pages of data from storage may be placed. Pet. 40–43 (citing Ex. 1006, Figs. 1 and 2, 2:1–3, 2:55–59, 3:46–50). Burton discloses "[d]ata from a device, such as a volatile memory device or non-volatile storage device, is maintained in entries in a cache." Ex. 1006, 2:1–3. We agree with Petitioner, and we find that Burton discloses "storing copies of files . . . in a cache."

"assigning frequency factors to each of the files stored in the cache"

Petitioner argues that Burton discloses "assigning frequency factors to each of the files stored in the cache, the frequency factors indicating how often each of the corresponding files are requested by the operating system" as recited in claim 4 in describing a least recently used (LRU) linked list that includes an LRU rank value for each entry in the cache. Pet. 43–44, (citing Ex. 1006, Fig. 1, 2:3–5, 2:60–63, 5:27–30). Petitioner contends Burton discloses that each LRU rank value indicates both how frequently the entry in cache is accessed and the time of last access, for the purpose of allowing the process demoting entries from cache "to take into account the frequency with which an entry was accessed and remove those less frequently accessed entries." Pet. 44 (citing Ex. 1006, Fig. 1, 3:11–17).

Patent Owner asserts that Burton does not disclose "assigning frequency factors to each of the files," as claim 4 requires. PO Resp. 49–54. Patent Owner asserts Burton's LRU rank value reflects or indicates a file's recency of use, not frequency of use. *Id.* at 49. According to Patent Owner, Burton "discloses that recency of use makes a dominant contribution to a block's LRU rank value and that frequency of use makes only a *de minimus* contribution that is overwhelmed by the recency of use." *Id.* at 50–51 (citing Ex. 2006 ¶¶ 161–163). In particular, Patent Owner contends that the LRU rank of Burton does not show, with a fair degree of certainty, how often each of the corresponding files is requested. *Id.* at 51. Patent Owner explains its position using an example application of the algorithm illustrated in Burton's Figure 2. *Id.* at 52–54 (citing Ex. 2006 ¶¶ 164–176). As Patent Owner argues, Burton discloses that the LRU rank for the most recently accessed cache entry is incremented, per block 110 of Figure 2, "by the modulo (*i.e.*, the remainder) of the time counter divided by 512." *Id.* at 52–53. Patent Owner describes an example in which three tracks of the cache have a higher LRU than a track labeled "A," even though they were accessed fewer times than track A. *Id.* at 53–54. According to Patent Owner, "Burton's LRU rank is not sufficiently related to the number of times that a given track has been accessed to show the frequency of use of a track with a fair degree of certainty." (emphasis omitted) *Id.* at 57.

Burton discloses, however, that the algorithm of Figure 2 is weighted for previous accesses. Ex. 1006, 3:54. The purpose of this weighting is to allow the process to take into account the frequency with which an entry was accessed and to remove those entries accessed less frequently. Ex. 1006, 3:13–17; Ex. 1002 ¶ 144. Burton's algorithm of Figure 2 increments the LRU rank for the most recently accessed cache entry, such that the incremented LRU rank is (a) weighted for previous accesses, and (b) will have a substantially higher LRU rank than the LRU rank of entries that have not been accessed recently, even those accessed numerous times. *See* Ex. 1006, 3:55–4:2; Ex. 1002 ¶¶ 144–45.

Thus, we find the LRU rank of Burton is based, in part on frequency of use, and in part on recency of use. *See* Ex. 1006, 3:11–13 ("The LRU rank . . . provides a value for each entry . . . in cache 2 that indicates both how frequently the entry is accessed and the time of last access"). As we discuss in our claim construction, the specification of the '745 patent does not require the frequency factors to show, with a fair degree of certainty, how often each of the corresponding files were accessed. Rather, the specification describes frequency factors that are based in part on frequency of use, and based in part on other criteria. *See* Ex. 1001, 6:56–7:51. We find that the LRU rank of Burton, which is weighted for previous accesses, describes "an indicator based on frequency of use or access." We agree with Petitioner that the LRU rank of Burton, which provides a value that indicates both how frequently the entry is accessed and the time of last access, describes "frequency factors indicating how often each of the corresponding files are requested by the operating system."

"scanning the frequency factors"

Petitioner argues that Burton discloses "scanning the frequency factors, the scanning being performed in response to a target capacity of the cache being attained," in describing determining when a number of cache entries has reached a threshold, and identifying entries that have the lowest LRU rank. Pet. 46–47 (citing Ex. 1006, Fig. 3, 2:6–11, 4:7–8). Dr. Franzon testifies that in response to reaching the threshold, the cache of Burton determines, from the last 1024 entries in the LRU linked list, thirty-two entries that have the lowest LRU rank, then demotes those thirty-two entries from the cache. Ex. 1002 ¶ 149 (citing Ex. 1006, 4:8–13, 19–22).

We agree with Petitioner, and based on the disclosure of Burton (Ex. 1006, 4:7–13, 19–22) and the testimony of Dr. Franzon, we find that Burton teaches the recited limitation of "scanning the frequency factors, the scanning being performed in response to a target capacity of the cache being attained."

"identifying a least frequently and least recently used file"

Petitioner argues that Burton discloses "identifying a least frequently and least recently used file" in describing those entries having the lowest LRU rank marked for demotion are both the least recently accessed and, among those entries least recently accessed, are less frequently accessed. Pet. 48 (citing Ex. 1006, Fig. 3, 4:19–22).

Patent Owner contends that Burton discloses that a relatively low variable value indicates that the cache entry is one of a least recently accessed entry and/or least frequently accessed entry. PO Resp. 48 (citing Ex. 1006, 2:11–13, claims 2, 6, 9, 13, 16, 20). According to Patent Owner, Burton may choose tracks for elimination that are least recently accessed, or least frequently accessed, but not both "least frequently and least recently used" as claimed. PO Resp. 48.

The claimed "least frequently and least recently used file," read in light of the specification, encompasses a file used less frequently than another file, and used less recently than another file. *See* Ex. 1001, Fig. 2A, 6:46–49 (selecting file F6 as the least frequently and least recently used file, even though file F2 was used less frequently, and file F7 was used less recently). Burton discloses that "those entries having the lowest LRU rank marked for demotion are both the least recently accessed and among those entries least recently accessed recently, are less frequently accessed. Ex. 1006, 4:19–22; Ex. 1002 ¶ 151. Dr. Franzon testifies that the LRU rank of Burton is used as an indicator for distinguishing data based in part on its frequency of use. Ex. 1002 ¶ 152. We agree with Petitioner and based on the disclosure of Burton (Ex. 1006, 4:19–22) and the testimony of Dr. Franzon, we find that Burton teaches "identifying a least frequently and least recently used file."

"eliminating the least frequently and least recently used file"

Petitioner argues that Burton discloses "eliminating the least frequently and least recently used file to liberate capacity of the cache" in describing that demoting data from the cache will make room for subsequently accessed tracks or data more recently accessed from storage. Pet. 48 (citing Ex. 1006, 1:39, 2:23).

Burton discloses "determining which entries to demote from cache to make room for subsequently accessed data." Ex. 1006, 2:23–24. Burton further discloses "those entries having the lowest LRU rank marked for demotion are both the least recently accessed and among those entries least recently accessed recently, are less frequently accessed." Ex. 1006, 4:19–22. Dr. Franzon testifies that the entries demoted from the cache of Burton would include the least frequently and least recently used entry among them. Ex. 1002 ¶ 154.

We agree with Petitioner, and based on the disclosure of Burton and the testimony of Dr. Franzon, we find that Burton teaches "identifying a least frequently and least recently used file."

"reading an extended segment of data"

Petitioner asserts that Burton discloses all of the limitations of claim 4, except "reading an extended segment of data in response to a request from an operating system," which Petitioner asserts Karedla discloses. *Id.* at 49 (citing Ex. 1004, 8). Karedla discloses "[a] read-ahead strategy known as prefetching exploits the principle of spatial locality . . . by anticipating future requests for data and bringing it into the cache." Ex. 1004, 8. We agree with Petitioner and based on the disclosure of Karedla, we find that Karedla teaches "reading an extended segment of data in response to a request from an operating system."

b. Alleged Rationale to Combine Asserted Art

Petitioner also asserts that a person of ordinary skill in the art would have been motivated to combine Burton's cache replacement mechanism with Karedla's method of reading ahead for prefetching data to improve system performance. Pet. 49–50 (citing Ex. 1002 ¶ 158; Ex. 1006, 1:61–62; Ex. 1004, 8). According to Petitioner, the combination provides a cost-effective method to increase performance and data throughput, minimize latency in data access, reduce cache lookup time, and increase the speed of cache directory searches. *Id.*

Based on (1) the teachings of Burton that are directed to determining which entries to demote from cache to make room for subsequently accessed data (Ex. 1006, 2:23–24), (2) the teachings of Karedla directed to caching strategies to improve system performance, including a read-ahead strategy known as prefetching (Ex. 1004, 8), and (3) the testimony of Dr. Franzon demonstrating the applicability of Burton to Karedla and explaining why a person of ordinary skill in the art would have combined the prior art teachings (Ex. 1002 ¶ 158), we find Petitioner provides articulated reasoning with rational underpinning for combining the teachings of Burton and Karedla as claimed—to improve system performance such as to minimize latency in data excess, to reduce cache lookup time, and to increase the speed of cache directory searches. *See* Pet. 49–50.

Accordingly, we determine that Petitioner shows, by a preponderance of the evidence, that claim 4 is unpatentable under 35 U.S.C. § 103(a) as obvious over the combination of Burton and Karedla. *See* Pet. 40–50.

c. Claims 1 and 12

Petitioner asserts that claims 1 and 12 are "substantively identical" to claim 4, noting that claims 1 and 12 are respectively directed to a similar method and media for identifying a least frequently and least recently used file. Pet. 50. Petitioner also asserts that Burton discloses the feature recited in claims 1 and 12 (but not in claim 4) of the cache being located in a random access memory. *Id.* at 51–53 (citing, e.g., Ex. 1006, 1:16–17, 1:26–27; Ex. 1002 ¶ 160). Patent Owner contests Petitioner's position, raising the same arguments for claims 1 and 12 discussed above in connection with claim 4. *See* PO. Resp. 47–60 (citing Ex. 2006 ¶¶ 161–176).

Claim 1 is a method claim that recites limitations that are very similar to those of method claim 4, with the additional requirement of "the cache being located in a random access memory (RAM) of the computer." Claim 12 recites "[a] computer readable media having program instructions for a caching method" comprising program instructions for performing functions nearly identical to those recited in claim 4. Like claim 1, claim 12 additionally requires that the cache be located in a random access memory of the computer.

Burton discloses placing pages or tracks of data copied from a storage device such as a hard disk drive, "into a volatile, electronic memory area referred to as a cache." Ex. 1006, 1:13–17. Dr. Franzon testifies that a "random access memory (RAM)" is a type of volatile memory and that it would have been "obvious that the cache in *Burton* may be located in a random access memory (RAM)." Ex. 1002 ¶ 160. Based on the disclosure of Burton and the testimony of Dr. Franzon, we are persuaded that locating the cache in a random access memory would have been obvious to a person of ordinary skill in the art.

For the reasons explained above in connection with claim 4, and because Petitioner shows it would have been obvious for the cache to be located in a random access memory, Petitioner shows by a preponderance of the evidence that claims 1 and 12 are unpatentable for obviousness over Burton and Karedla.

d. Summary

For the above reasons, we determine that Petitioner shows that claims 1, 4, and 12 are unpatentable under 35 U.S.C. § 103(a) as obvious over the combination of Burton and Karedla.

G. Asserted Obviousness over either (i) Karedla and Otterness or (ii) Burton, Karedla, and Otterness

Petitioner contends that claim 2 of the '745 patent is unpatentable under 35 U.S.C. § 103(a) as obvious over either (i) Karedla and Otterness or (ii) Burton, Karedla, and Otterness. Pet. 3–4, 67–70. Relying, in part, on the testimony of Dr. Franzon, Petitioner explains how the references allegedly teach or suggest the claim limitations and provides reasoning for combining the teachings of the references. *Id.* at 67–70 (citing Ex. 1002

¶¶ 191–195). Patent Owner contests Petitioner's position, arguing Otterness fails to teach the underlying limitations of independent claim 1, and argues that claim 2 is patentable for the same reasons as claim 1. PO Resp. 60.

1. Summary of Otterness

Otterness is a U.S. patent that issued October 1, 2002 from an application filed September 30, 1999, and is prior art to the '745 patent under 35 U.S.C. § 102(e). *Compare* Ex. 1008, [22], [45], *with* Ex. 1001, [22], [45]. Otterness discloses a multiple level caching method that distributes caching operations across multiple processors to provide higher input/output processing performance. Ex. 1008, 1:10–17. Otterness discloses using a cache line descriptor (CLD) "to keep track of all of the data stored in the cache. . . ." *Id.* at 9:23–25.

Figure 5 of Otterness is shown below:

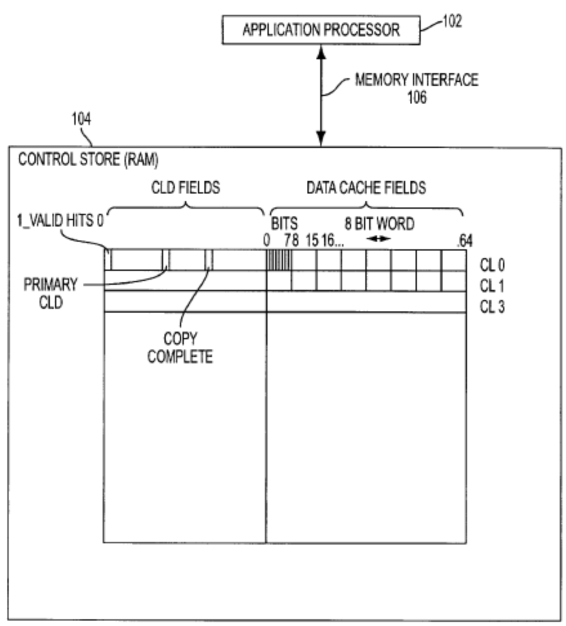

FIG. 5

Figure 5, above, illustrates an exemplary embodiment in which each cache data line has its own associated CLD. *Id.* at 10:36–37. Otterness discloses that "[t]he CLDs provide various pointers to manage data movement." *Id.* at 9:46–47. Otterness also discloses, "the cache line stores data (for example 8 kbyte, 16 kbyte, 32 kbyte, 64 kbyte, or the like of data)." *Id.* at 10:37–40.

2. Analysis

Claim 2 recites "[t]he method as recited in claim 1, wherein the extended segments are one of 64 Kbytes, 128 Kbytes and 256 Kbytes in size." Petitioner asserts that a person of ordinary skill in the art reading Otterness would have understood that, although Otterness "does not expressly list out cache line sizes of 128 kbyte or 256 kbyte," a skilled artisan would have understood that Otterness's disclosure of "'or the like of data' refers to such sizes." Pet. 68 (citing Ex. 1002 ¶ 193). Petitioner also notes that the three references "relate[] to the same field of endeavor (using cache to enhance computer system performance)." *Id.* at 68 (citing Ex. 1002 ¶ 191).

In its Response, Patent Owner does not raise any arguments specific to the limitation added by dependent claim 2 or in response to Petitioner's assertions regarding the reasons to combine the teachings of Otterness, Karedla, and Burton. *See* PO. Resp. 60. The burden, however, remains on Petitioner to demonstrate unpatentability. *See Dynamic Drinkware*, 800 F.3d at 1378.

We determine that Petitioner shows by a preponderance of the evidence that the combination of either Karedla and Otterness or Burton, Karedla, and Otterness renders obvious "wherein the extended segments are one of 64 Kbytes, 128 Kbytes and 256 Kbytes in size." *See, e.g.*, Ex. 1008, 10:37–40. As explained above, we determine that Petitioner shows that claim 1 is unpatentable under 35 U.S.C. § 102(b) as anticipated by Karedla and unpatentable under 35 U.S.C. § 103(a) as obvious over Burton and Karedla. Otterness teaches cache lines of various sizes, including 64 KB. Ex. 1008, 10:37–39. As discussed above, Karedla's disclosure of reading ahead (prefetching data) discloses reading extended segments of data. Dr. Franzon testifies that "one of ordinary skill in the art would have been motivated to combine *Karedla*, *Burton*, and *Otterness* to potentially improve the performance of systems." Ex. 1002 ¶ 194. Karedla discloses using "read-ahead strategies, with varying thresholds on the amount of data that is prefetched." Ex. 1004, 8. Karedla also discloses that, "[f]or large cache lines, cache directory searches can be faster, reducing cache lookup time." *Id*. Karedla, therefore, suggests varying the cache line size to improve performance, consistent with Dr. Franzon's testimony. Based on the evidence of record, therefore, we find that a person of ordinary skill in the art would have had reason to combine Otterness with Karedla and Burton, and we conclude that the subject matter of claim 2 would have been obvious based on the combined teachings of Karedla and Otterness and based on the combined teachings of Karedla, Burton, and Otterness.

Thus, for these reasons and the reasons explained above in connection with Petitioner's contentions that claim 1 is unpatentable as anticipated by Karedla or obvious over Burton and Karedla, Petitioner shows by a preponderance of the evidence that claim 2 is unpatentable under 35 U.S.C. § 103(a) as obvious over Karedla and Otterness or over Burton, Karedla, and Otterness.

H. Remaining Patent Owner Arguments

Patent Owner makes additional arguments stating that post grant review proceedings, such as this trial, are unconstitutional and are an impermissible taking of a private right without Article III oversight. PO Resp. 60–61. We decline to consider the constitutional challenge as, generally, "administrative agencies do not have jurisdiction to decide the constitutionality of congressional enactments" where consideration of the constitutional question would "require the agency to question its own statutory authority or to disregard any instructions Congress has given it." *Riggin v. Office of Senate Fair Employment Practices*, 61 F.3d 1563, 1569–70 (Fed. Cir. 1995).[7]

III. MOTIONS TO EXCLUDE EVIDENCE

Patent Owner filed a motion to exclude Exhibit 1015. Paper 36. Because we do not rely on Exhibit 1015 in our Final Decision we dismiss the motion to exclude Exhibit 1015 as moot. Petitioner filed a contingent motion to exclude Exhibits 2011–2015 in the event that the Board excludes Exhibit 1015. Paper 37. Because we do not exclude Exhibit 1015, we dismiss the contingent motion to exclude Exhibits 2011–2015 as moot.

Patent Owner's motion to exclude is dismissed, and Petitioner's contingent motion to exclude is dismissed.

---

[7] On June 12, 2017, the United States Supreme Court granted a petition for a writ of certiorari in *Oil States Energy Services, LLC v. Green's Energy Group, LLC*, 639 Fed. App'x 639 (Fed. Cir 2016), cert. granted, 198 L. Ed. 2d 677 (U.S. Jun. 12, 2017) (No. 16-712). The Court will answer the question of whether the USPTO's statutorily created IPR process is unconstitutional.

IV. MOTION TO STRIKE

Patent Owner filed a motion to strike Petitioner's Reply, because of alleged improper reply arguments (Paper 32) to which Petitioner filed an opposition (Paper 40). Patent Owner lists several portions of Petitioner's Reply and evidence as allegedly beyond the scope of what can be considered appropriate for a reply. *See* Paper 32.

We have considered Patent Owner's motion, but we disagree that the cited portions of Petitioner's Reply and reply evidence are beyond the scope of what is appropriate for a reply. Replies are a vehicle for responding to arguments raised in a corresponding patent owner response. Petitioner's arguments and evidence to which Patent Owner objects are not beyond the proper scope of a reply because we find that they fairly respond to Patent Owner's arguments raised in Patent Owner's Response. *See Idemitsu Kosan Co., LTD. v. SFC Co. LTD*, 870 F.3d 1376, 1381 ("This back-and-forth shows that what Idemitsu characterizes as an argument raised 'too late' is simply the by-product of one party necessarily getting the last word. If anything, Idemitsu is the party that first raised this issue, by arguing—at least implicitly—that Arkane teaches away from non-energy-gap combinations. SFC simply countered, as it was entitled to do."). We also note for each of the items in Patent Owner's list, Petitioner cites to pages in the Petition for support. *See e.g.*, Paper 40, 2–3 ("the summary of SLRU on page 23 of Petitioner's Reply is fully supported by the Petition. *See* Pet. 20, 23, 26–28.").

Accordingly, Patent Owner's motion to strike the Reply is denied.

V. CONCLUSION

On this record, Petitioner has shown by a preponderance of the evidence that claims 1, 2, 4, 12, and 14 of the ’745 patent are unpatentable.

VI. ORDER

Accordingly, it is:

ORDERED that

1. Claims 1, 4, 12, and 14 of the ’745 patent are unpatentable on the ground of anticipation by Karedla;
2. Claims 1, 4, and 12 of the ’745 patent are unpatentable on the ground of obviousness over Burton and Karedla;
3. Claim 2 of the ’745 patent is unpatentable on the ground of obviousness over Karedla and Otterness;
4. Claim 2 of the ’745 patent is unpatentable on the ground of obviousness over Burton, Karedla, and Otterness;

FURTHER ORDERED that Patent Owner’s motion to exclude Exhibit 1015 is dismissed;

FURTHER ORDERED that Petitioner’s contingent motion to exclude Exhibits 2011–2015 is dismissed;

FURTHER ORDERED that Patent Owner’s motion to strike Petitioner’s Reply is denied; and

FURTHER ORDERED that because this Final Written Decision is final, a party to the proceeding seeking judicial review of the Decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

FOR PETITIONER:

Roshan Suresh Mansinghani
Jonathan Stroud
UNIFIED PATENTS INC.
roshan@unifiedpatents.com
jonathan@unifiedpatents.com

Lionel M. Lavenue
James Stein
FINNEGAN, HENDERSON, FARABOW
GARRETT & DUNNER, LLP
lionel.lavenue@finnegan.com
james.stein@finnegan.com

FOR PATENT OWNER:

John R. King
Ted M. Cannon
Brenton R. Babcock
KNOBBE, MARTENS, OLSON & BEAR, LLP
2jrk@knobbe.com
2tmc@knobbe.com
2brb@knobbe.com

Tim Seeley
James R. Hietala
INTELLECTUAL VENTURES
tims@intven.com
jhietala@intven.com