

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF OREGON**

MEMORY INTEGRITY, LLC,

Plaintiff,

v.

INTEL CORPORATION,

Defendant.

Case No. 3:15-cv-00262-SI

**OPINION AND ORDER
ON CLAIM CONSTRUCTION**

John Mansfield, MANSFIELDLAW, 121 SW Morrison Avenue, Suite 400, Portland, OR 97204; Bryan Atkinson and George W. Webb, III, FARNEY DANIELS, PC, 800 S. Austin Avenue, Suite 200, Georgetown, TX 78626; Jonathan Baker, Michael Saunders, and Gurtej Singh, FARNEY DANIELS, PC, 411 Borel Avenue, Suite 350, San Mateo, CA 94402. Of Attorneys for Plaintiff.

Renée E. Rothauge, MARKOWITZ HERBOLD, PC, 1211 SW Fifth Avenue, Suite 3000, Portland, OR 97204; Michael J. Summersgill, Alexandra Amrhein, Jordan L. Hirsch, and Sean K. Thompson, WILMERHALE, LLP, 60 State Street, Boston, MA 02109; Grant K. Rowan, WILMERHALE, LLP, 1875 Pennsylvania Avenue NW, Washington, DC 20006; Arthur W. Coviello, WILMERHALE, LLP, 950 Page Mill Road, Palo Alto, CA 94304. Of Attorneys for Defendant.

Michael H. Simon, District Judge.

Plaintiff Memory Integrity, LLC (“Memory Integrity” or “Plaintiff”) has brought infringement claims against Defendant Intel Corporation (“Intel” or “Defendant”) under five patents: U.S. Patent Nos. 7,103,636 (the “636 patent”), 7,107,409 (the “409 patent”), 7,296,121

(the “121 patent”), 8,572, 206 (the “206 patent”), and 8,898,254 (the “254 patent”). The patents are all directed toward maintaining cache coherence in multiprocessor computer systems. On March 11, 2016, the Court held a claim construction hearing. Based on the documents filed and the expert testimony admitted in the claim construction proceeding, the Court construes the disputed terms below.

STANDARDS

Patent infringement analysis involves two steps. First, the court construes the asserted patent claims. *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 976 (Fed. Cir. 1995) (en banc), *aff’d*, 517 U.S. 370 (1996). Second, the factfinder determines whether the accused product or method infringes the asserted claim as construed. *Id.* The first step, claim construction, is a matter of law. *See Markman*, 517 U.S. at 372; *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996). “It is a ‘bedrock principle’ of patent law that ‘the claims of a patent define the invention to which the patentee is entitled the right to exclude.’” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005) (en banc) (quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1115 (Fed. Cir. 2004)). Patent claims must precisely define the relevant invention and thereby serve to put both the public and competitors on notice of the claimed invention. *See Phillips*, 415 F.3d at 1312.

In interpreting a patent claim during the claim construction process, there are “numerous sources that [a district court] may properly utilize for guidance.” *Vitronics*, 90 F.3d at 1582. Some types of evidence are “more valuable than others.” *Phillips*, 415 F.3d at 1324. The relevant sources of evidence include both intrinsic evidence and extrinsic evidence. *Vitronics*, 90 F.3d at 1582. A court begins claim construction with consideration of the intrinsic evidence in the record, consisting of the claim terms, the patent specification, and the patent prosecution history. *Phillips*, 415 F.3d at 1313. Where the intrinsic evidence alone fails to resolve a claim

construction dispute, a court may consider extrinsic evidence. *Vitronics*, 90 F.3d at 1583.

Extrinsic evidence “consists of all evidence external to the patent and prosecution history, including expert and inventor testimony, dictionaries, and learned treatises.” *Markman*, 52 F.3d at 980.

Intrinsic evidence is the court’s starting point because it is “the most significant source of the legally operative meaning of disputed claim language.” *Vitronics*, 90 F.3d at 1582. The first type of intrinsic evidence is “the words of the claims themselves, both asserted and nonasserted.” *Id.* The words of a claim are “generally given their ordinary and customary meaning.” *Phillips*, 415 F.3d at 1312 (citation and quotation marks omitted). A claim term’s “ordinary and customary meaning” 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.” *Id.* at 1313.

There is a “heavy presumption” that a claim term carries its “ordinary and customary meaning.” *Elbex Video, Ltd. v. Sensormatic Elecs. Corp.*, 508 F.3d 1366, 1371 (Fed. Cir. 2007). There are two exceptions to this general rule: (1) “when a patentee sets out a definition and acts as his own lexicographer;” or (2) “when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012); *see also Hormone Research Found., Inc. v. Genentech, Inc.*, 904 F.2d 1558, 1563 (Fed. Cir. 1990) (“It is a well-established axiom in patent law that a patentee is free to be his or her own lexicographer and thus may use terms in a manner contrary to or inconsistent with one or more of their ordinary meanings.” (citation omitted)).

In some cases, the “ordinary meaning of claim language” may be “readily apparent to lay judges,” in which case, “claim construction . . . involves little more than the application of the

widely accepted meaning of commonly understood words.” *Phillips*, 415 F.3d at 1314. In such cases, commonplace terms or those that a juror can easily use without further direction from the court “do not need to be construed because they are neither unfamiliar to the jury, confusing to the jury, nor affected by the specification or prosecution history.” *Bd. of Trs. of Leland Stanford Junior Univ. v. Roche Molecular Sys., Inc.*, 528 F. Supp. 2d 967, 976 (N.D. Cal. 2007). In those circumstances, it is enough to hold that the ordinary and customary meaning controls. *Finjan, Inc. v. Secure Computing Corp.*, 626 F.3d 1197, 1206-07 (Fed. Cir. 2010) (holding that the district court “was not obligated to provide additional guidance to the jury” beyond directing the jury to apply the “ordinary meaning” of a claim term).

“[T]he context in which a term is used in the asserted claim can be highly instructive.” *Phillips*, 415 F.3d at 1314. For example, it is improper for a court to interpret a claim term in a manner that renders subsequent claim terms superfluous. *See Stubmo v. Eastman Outdoors, Inc.*, 508 F.3d 1358, 1362 (Fed. Cir. 2007); *Merck & Co. v. Teva Pharms. USA, Inc.*, 395 F.3d 1364, 1372 (Fed. Cir. 2005) (“A claim construction that gives meaning to all the terms of the claim is preferred over one that does not do so.”). The doctrine of claim differentiation further provides that the limitations in each claim are presumed to be distinct from one another, although simply “describing claim elements or limitations in different words does not invariably change the scope of the claim.” *Inpro II Licensing, S.A.R.L. v. T-Mobile USA, Inc.*, 450 F.3d 1350, 1354 (Fed. Cir. 2006).

The second type of intrinsic evidence is the patent specification. A court does not determine the ordinary meaning of a disputed term “in a vacuum.” *Medrad, Inc. v. MRI Devices Corp.*, 401 F.3d 1313, 1319 (Fed. Cir. 2005). The claims “must be read in view of the specification, of which they are a part.” *Markman*, 52 F.3d at 979. “[T]he 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.” *Vitronics*, 90 F.3d at 1582. The purpose of a patent’s specification is to “teach and enable those of skill in the art to make and use the invention and to provide a best mode for doing so.” *Phillips*, 415 F.3d at 1323. A specification may limit or disavow the scope of a claim, however, when a patentee makes a “clear and unmistakable disclaimer” deviating from the ordinary meaning. *Thorner*, 669 F.3d at 1366-67. For example, where a specification makes clear that an invention does not include a particular feature, “that feature is deemed to be outside the reach of the claims of the patent, even though the language of the claims, read without reference to the specification, might be considered broad enough to encompass the feature in question.” *Thorner*, 669 F.3d at 1366 (quoting *SciMed Life Sys., Inc. v. Advanced Cardiovascular Sys., Inc.*, 242 F.3d 1337, 1341 (Fed. Cir. 2001)).

The Federal Circuit has repeatedly warned, however, that “although the specification often describes very specific embodiments of the invention,” the claims should not necessarily be confined to those embodiments. *Phillips*, 415 F.3d at 1323; *Intel Corp. v. U.S. Int’l Trade Comm’n*, 946 F.2d 821, 836 (Fed. Cir. 1991) (holding that “[w]here a specification does not require a limitation, that limitation should not be read from the specification into the claims” (citation, quotation, and emphasis omitted); see also *Amhil Enters., Ltd. v. Wawa, Inc.*, 81 F.3d 1554, 1559 (Fed. Cir. 1996) (holding that “[a] preferred embodiment . . . is just that, and the scope of a patentee’s claims is not necessarily or automatically limited to the preferred embodiment”). Ultimately, there is “no magic formula or catechism for conducting claim construction.” *Phillips*, 415 F.3d at 1324. The court must “read the specification in light of its purposes in order to determine whether the patentee is setting out specific examples of the invention to accomplish those goals, or whether the patentee instead intends for the claims and

the embodiments in the specification to be strictly coextensive.” *Decisioning.com, Inc. v. Federated Dept. Stores, Inc.*, 527 F.3d 1300, 1308 (Fed. Cir. 2008) (quoting *Phillips*, 415 F.3d at 1323).

The third type of intrinsic evidence is the prosecution history of a patent (or the “file wrapper”), if it is in evidence. *Markman*, 52 F.3d at 980; *see also Graham v. John Deere Co.*, 383 U.S. 1, 33 (1966) (“[A]n invention is construed not only in light of the claims, but also with reference to the file wrapper or prosecution history in the Patent Office.”). The patent history with the U.S. Patent and Trademark Office (“PTO”) “contains the complete proceedings . . . , including any express representations made by the applicant regarding the scope of the claims.” *Vitronics*, 90 F.3d at 1582. Because the prosecution represents an “ongoing negotiation between the PTO and the applicant, rather than the final product of that negotiation, it often lacks the clarity of the specification and thus is less useful for claim construction purposes.” *Phillips*, 415 F.3d at 1317. The prosecution history, however, can inform the meaning of a claim term where a patentee makes “a clear and unmistakable disavowal of scope during prosecution.” *Creative Integrated Sys., Inc. v. Nintendo of Am., Inc.*, 526 F. App’x 927, 934 (Fed. Cir. 2013). “A patentee could do so, for example, by clearly characterizing the invention in a way to try to overcome rejections based on prior art.” *Computer Docking Station Corp. v. Dell, Inc.*, 519 F.3d 1366, 1374 (Fed. Cir. 2008); *see also Chimie v. PPG Indus., Inc.*, 402 F.3d 1371, 1384 (Fed. Cir. 2005) (“The purpose of consulting the prosecution history in construing a claim is to exclude any interpretation that was disclaimed during prosecution.” (citation and quotation marks omitted)). If the alleged disavowal of the claim scope is ambiguous, it will not limit the scope of the claim. *Creative Integrated Sys., Inc.*, 526 F. App’x at 934.

Generally an “analysis of the intrinsic evidence alone will resolve any ambiguity in a disputed claim term.” *Vitronics*, 90 F.3d at 1583. Where that is the case, “it is improper to rely on extrinsic evidence.” *Id.* Where the ordinary and customary meaning of a claim term is not apparent based on the intrinsic evidence, courts are authorized to consult extrinsic evidence. *Markman*, 52 F.3d at 980. Relevant extrinsic evidence includes “expert and inventor testimony, dictionaries, and learned treatises.” *Id.* at 980. Judges may “rely on dictionary definitions when construing claims, ‘so long as the dictionary definition does not contradict any definition found in or ascertained by a reading of the patent documents.’” *Meyer Intellectual Props. Ltd. v. Bodum, Inc.*, 690 F.3d 1354, 1368 (Fed. Cir. 2012) (quoting *Phillips*, 415 F.3d at 1322). Indeed, “heavy reliance on the dictionary divorced from the intrinsic evidence risks transforming the meaning of the claim term to the artisan into a meaning of the term in the abstract, out of its particular context, which is the specification.” *Phillips*, 415 F.3d at 1321. Thus, a patentee is not “entitled to a claim construction divorced from the context of the written description and prosecution history.” *Nystrom v. Trex Co., Inc.*, 424 F.3d 1136, 1144-45 (Fed. Cir. 2005). Extrinsic evidence, therefore, is “less reliable than the patent and its prosecution history in determining how to read claim terms,” and its consideration is within the court’s sound discretion. *Phillips*, 415 F.3d at 1318-19.

BACKGROUND

Memory Integrity is the assignee of the five asserted patents. All five patents purport to solve the cache coherency problem in multiprocessor computer systems. This problem arises in multiprocessor computer systems where the processors share a main memory. The main memory stores data needed or generated by the system, but each individual processor also uses its own smaller, faster “cache” memory to store copies of data upon which the processor regularly operates. When the data is stored in the processor’s cache, the processor may change the data.

The data is then saved back to the main memory after operations conclude. In the interim, changes to data in an individual cache may cause the master copy in the main memory to become “stale” or out-of-date, and multiple processors may start using different versions of the same data, leading to cache incoherence. Methods of maintaining cache coherence, such as those described by the asserted patents, ensure that processors have access to the most up-to-date copies of data and that the system does not generate inconsistent versions of data.

The five asserted patents share common inventors and have overlapping specifications. The patents all describe methods of maintaining cache coherence in the same type of multiprocessor system: a system of multiple processor clusters interconnected in a “point-to-point architecture.” Dkt. 136-1 at 27 (’636 patent at 20:55-63); Dkt. 136-2 at 25 (’409 patent at 18:5-14); Dkt. 136-3 (’121 patent at 31:23-27); Dkt. 136-4 at 24 (the ’206 patent at 20:57-59); Dkt. 136-5 at 23 (the ’254 patent at 18:37-39).¹ The system described by each of the patents also uses a “cache coherence controller” or “interconnection controller” to maintain cache coherence across multiple clusters of processors. *See* ’636 patent at 9:45-49; ’409 patent at 8:48-52; ’121 patent at 8:50-54; ’206 patent 3:67-4:3; ’254 patent at 4:3-6.

The ’636 and ’409 patents detail a method known as “speculative probing.” A processor’s request for data in the multiprocessor system goes through a “serialization point,” defined in the patents as “[a]ny mechanism for ordering data access requests.” ’636 patent at 5:43-44; ’409 patent at 4:65-66. Before the requests arrive at the serialization point, the system’s cache coherence controller sends out “probes” to determine whether any cache memory locations have

¹ The docket numbers of the asserted patents are included only in the first citation to the patents. The patents are cited in the order in which they issued to either Memory Integrity or its predecessor in interest. The ’636, ’409, ’121, ’206, and ’254 patents issued on September 5, 2006, September 12, 2006, November 13, 2007, October 29, 2013, and November 25, 2014, respectively.

modified copies of the requested data. '636 patent at 6:16-18; '409 patent at 5:38-40. According to the patents, sending speculative probes before a request reaches the serialization point increases the efficiency of the system. Speculative probing helps ensure cache coherence, but may also generate excessive probe traffic if all remote or all local caches are probed regardless of whether they contain the requested data. The '636 patent is directed at speculative probing of "remote" processor clusters, *i.e.* those that do not contain the processors requesting the data. '636 patent at 3:3-7. The '409 patent is directed at speculative probing of "local" processor clusters, those that contain the processors requesting the data. '409 patent at 2:67-3:5.

The '121 patent describes a technique for further reducing probe traffic. The patent purports to reduce probe traffic by using a probe filtering unit ("PFU"). In the system claimed by the patent, the cache coherence controller receives a request for data and then sends a probe to the PFU. The PFU contains "probe filtering information," which allows the PFU to determine if the requested data is located in cache memory within the system. '121 patent at 2:67-3:1. If no cache memory contains the requested data, the PFU does not send any probes. If the PFU determines that a processor's cache may contain a copy of requested data, the PFU sends a probe to that cache alone.

The '254 patent is a continuation of the '206 patent, and both describe a technique for purportedly increasing the speed of memory transactions. The patents describe a multiprocessor system that has both "local memory" and "remote memory." '206 patent at 7:32-38; '254 patent at 7:35-40. Local memory is stored in the same cluster as a given processor, and remote memory is all memory outside a given processor's cluster. "Protocol engines" process requests for local and remote memory. '206 patent at 1:48-67; '254 patent at 1:51-2:4. The patents describe a way

to improve the efficiency of memory transaction processing by dividing the processing workload among multiple protocol engines, according to whether the targeted memory is local or remote.

CLAIM CONSTRUCTION

A. Term One: “point-to-point architecture”

Memory Integrity’s Proposed Construction	Intel’s Proposed Construction
“an architecture in which multiple processors or processing nodes are directly connected to each other through point-to-point links and share the same memory address space”	“an architecture in which multiple processors are directly connected to each other through point-to-point links”

The term “point-to-point architecture” appears in the asserted claims of each of the five patents-in-suit. The parties disagree on two issues: (1) whether “point-to-point architecture” must connect multiple “processors” or “processors *or processing nodes*”; and (2) whether “point-to-point architecture” must “share the same memory address space.” The parties propose one construction of the term across all five patents, although, in the alternative, Memory Integrity proposes that the Court construe “point-to-point architecture” one way for the ’636 and ’409 patents, which describe interconnected “processors,” and another way for the ’121, ’206, and ’254 patents, which describe interconnected “processing nodes.”

1. Processors or Processing Nodes

The Court looks first to the intrinsic evidence—the text of the claims, the specifications, and the prosecution history—to obtain the term’s meaning. *Goldenberg v. Cytogen, Inc.*, 373 F.3d 1158, 1164 (Fed. Cir. 2004). The ’636 and ’409 patents claim:

A computer system, comprising: a first cluster including a first plurality of processors and a first cache coherence controller, the first plurality of processors, and the first cache coherence controller interconnected in a point-to-point architecture; a second cluster including a second plurality of processors and the second cache coherence controller interconnected in a point-to-point

architecture, the first cache coherence controller coupled to the second cache coherence controller;

'636 patent at 20:53-64, claim 15; '409 patent 17:40-50, claim 1.

The '121 patent claims: “A computer system comprising a plurality of processing nodes interconnected by a first point-to-point architecture, each processing node having a cache memory associated therewith” '121 patent at 30:65-31:1, claim 1. The '206 and '254 patents claim a cluster consisting of processing nodes and “an interconnection controller coupled to the plurality of processing nodes . . . wherein the processing nodes and interconnection controller are interconnected in a point-to-point architecture.” '206 patent at 20:40-59, claims 30-31; '254 patent at 14-39, claims 1-2. Patent '206 also claims “[a]n interconnection controller for use in a computer system having one or more processor clusters, each cluster including a plurality of local nodes and an instance of the interconnection controller interconnected by a local point-to-point architecture.” '206 patent at 20:3-7, claim 21.

The claims do not specifically define “point-to-point architecture,” leaving open the possibility that the term could refer to interconnected processors, interconnected processors and cache coherence controllers, or interconnected processing nodes. The text of the '636, '409, and '121 patent specifications clarifies the meaning of “point-to-point architecture” by stating: “In a point to point architecture, a cluster of processors includes *multiple processors directly connected to each other* through point-to-point links.” '636 patent at 5:8-11; '409 patent at 4:32-35; '121 patent at 4:38-40 (emphasis added). The '254 and '206 patents do not contain this definition; the latter two patents instead state that the multiprocessor system “includes a plurality of local nodes and an interconnection controller interconnected by a local point-to-point architecture.” '206 and '254 patents, Abstracts. The '206 and '254 patents also note, however, that “[m]ulti-processor architectures having *point-to-point communication among their*

processors are suitable for implementing specific embodiments of the present invention.”

’206 patent 2:62-65; ’254 patent 2:65-3:1 (emphasis added). Further, the ’206 and ’254 patents explain that “[a]lthough generally a node may correspond to one or a plurality of resources (including, for example, a processor), it should be noted that the terms node and processor are often used interchangeably herein.” ’206 patent at 6:36-39; ’254 patent at 6:40-43. The ’206 and ’254 specifications are consistent with a definition of “point-to-point architecture” in which the interconnected processing nodes are processors.

Additionally, Figure 2 of the specifications (below), which appears in all five of the patents, shows an architecture in which multiple processors are directly connected to each other through point-to-point links. Each of the five patents includes the same discussion of Figure 2, showing “processors 202a-202d” and “point-to-point communication links 208a-208e.” ’636 patent at 7:31-37; ’409 patent at 6:34-39; ’121 patent at 6:36-41; ’206 patent at 3:34-39; ’254 patent at 3:37-42. Although some processors are also connected to the controller and I/O switch through point-to-point links, all processors are connected to at least one other processor.




Fig. 2

The specifications do not necessarily limit the scope of the claims. *Phillips*, 415 F.3d at 1323. The specifications are, however, “highly relevant to the claim construction analysis” and often provide “the single best guide to the meaning of a disputed term.” *Vitronics*, 90 F.3d at 1582. As the Federal Circuit has emphasized, “[t]here is a fine line between construing the claims in light of the specification and improperly importing a limitation from the specification into the claims.” *Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 653 F.3d 1296, 1305 (Fed. Cir. 2011). A court must strive to “tether the claims to what the specifications indicate the inventor actually invented.” *Id.* Although the specifications do not confine “point-to-point architecture” to an embodiment in which *only* processors are directly connected through point-to-point links, the specifications indicate that, at a minimum, “point-to-point architecture” must include multiple processors that have point-to-point connections.

The prosecution history of ’409 confirms this understanding of “point-to-point architecture,” at least with respect to the term’s use in the ’409 patent.² During the prosecution of this patent, the patent examiner initially rejected the claims based on prior art U.S. Patent No. 6,760,819 to Sang Hoo Dhong (“Dhong”). The Dhong patent disclosed direct point-to-point connections between one main memory, containing a cache coherence controller, and multiple processing units with cache memories. Based on that disclosure, the patent examiner concluded

² The Court recognizes that statements made during the prosecution history of an earlier patent, even one filed by the same inventor, do not limit the claims of a later patent when the later patent is “not filed as a continuation, continuation-in-part, or divisional application.” *Abbott Labs. v. Dey, L.P.*, 287 F.3d 1097, 1105 (Fed. Cir. 2002). Other than the ’206 and ’254 patents, the patents-in-suit were not continuations, continuations-in-part, or divisional applications of each other. Accordingly, the Court does not construe the prosecution history of the ’409 patent—the first of the five asserted patents to be filed—as binding upon the other patents. *Cf. Elkay Mfg. Co. v. EbcO Mfg. Co.*, 192 F.3d 973, 980 (Fed. Cir. 1999) (“When multiple patents derive from the same initial application, the prosecution history regarding a claim limitation in any patent that has issued applies with equal force to subsequently issued patents that contain the same claim limitation.”).

that, to a person of ordinary skill in the art at the time the invention was made, it would have been obvious to have a plurality of processors and a coherence controller interconnected in a “point-to-point architecture.”

The patentees responded that other than the connections between the several individual cache memories and the one main memory with the cache coherence controller, “Dhong does not teach any other point-to-point connections.” Dkt. 136-10 at 11. The patentees repeatedly emphasized that their invention required “a cache coherence controller and processors interconnected in a point-to-point architecture” and that “Dhong actually teaches away from the practice of connecting processing nodes to a cache coherence controller in a point-to-point architecture.” *Id.* at 11-12. Although the description of the invention allowed for “processing nodes” (left undefined) to be interconnected, the patentees’ explanation indicated that there must be point-to-point connections between processors as well.

Having considered the intrinsic evidence, the Court finds some remaining ambiguities at least with the respect to the meaning of “point-to-point architecture” in the ’206 and ’254 patents, which allow for the term “processing nodes” to include processors but do not precisely define the boundaries of “processing nodes.” The Court therefore turns to extrinsic evidence from expert witnesses.

Eric Morton, one of the named inventors of the ’121 patent, testified in his deposition that “[p]oint-to-point architecture is . . . where the processors are connected using single point-to-point links. There is only one processor on each side of the link” Dkt. 137-1 at 5.³

³ The testimony of an inventor “cannot be relied on to change the meaning of the claims.” *Markman*, 52 F.3d at 983. In particular, “[t]he subjective intent of the inventor when he used a particular term is of little or no probative weight in determining the scope of a claim (except as documented in the prosecution history).” *Id.* at 985. This rule remains true “[w]hether an inventor’s testimony is consistent with a broader or narrower claim scope” because “that

Professor Arvind Mithal similarly opined in his declaration that the claims, specifications, and prosecution history indicate that in order to have a “‘point-to-point architecture’ . . . it is not sufficient if the system has direct connections between other components besides processors.” Dkt. 147 ¶ 37.

In contrast, Professor Mark Jones concludes in his declaration that a “point-to-point architecture” need only include “processing nodes” because the claims of the ’121, ’206, and ’254 patents only mention “nodes” or “processing nodes.” Dkt. 139 ¶¶ 11-15. According to Professor Jones, a “node” can be a single processor but could also include multiple processors or other subunits such as “I/O bridges.” *Id.* ¶ 14. Professor Jones offers no other definition of “processing nodes.”

The Court agrees with Professor Jones that the claims, specifications, and prosecution history indicate that a “point-to-point architecture” could include “processing nodes” that are not limited to single processors. The Court, however, finds nothing in the claims, specifications, or prosecution history to indicate that a “point-to-point architecture” could exist without any processor directly connected to another processor. The patents disclose only embodiments in which at least two processors are directly connected to each other. “Point-to-point architecture” must at least include multiple processors.

2. Shared Memory Address Space

Memory Integrity argues that use of a shared memory address space distinguishes “point-to-point architecture” from other methods of interconnection such as external networks. Intel

testimony is still limited by the fact that an inventor understands the invention but may not understand the claims, which are typically drafted by the attorney prosecuting the patent application.” *Howmedica Osteonics Corp. v. Wright Med. Tech., Inc.*, 540 F.3d 1337, 1346-47 (Fed. Cir. 2008). Nonetheless, an inventor’s testimony “may be pertinent as a form of expert testimony, for example, as to understanding the established meaning of particular terms in the relevant art.” *Id.* at 1347 n.5. This is the case here.

responds that the patents omitted “shared memory address space” from their definitions of “point-to-point architecture.” The Court begins with the intrinsic evidence.

Some of the patents include claims that discuss a shared memory address space while other patents do not. The ’409 patent most clearly addresses shared memory space. The patent claims a cache coherence controller “wherein the plurality of local processors in the cluster share a memory address space with a plurality of non-local processors in the non-local cluster.” ’409 patent at 18:37-39, claim 9. Additionally, the ’409 patent claims a cache coherency method including “the non-local cache coherence controller associated with a remote cluster of processors connected through a point-to-point architecture, wherein the remote cluster shares an address space with the local cluster of processors.” *Id.* at 19:38-42, claim 25. Claim 34 of the ’409 patent describes a “non-local cache coherence controller associated with a remote cluster of processors connected through a point-to-point architecture, wherein the remote cluster of processors shares an address space with the local cluster of processors.” *Id.* at 20:13-17, claim 34. Claim 42 of the ’409 patent describes a system consisting of a first cluster with a first processor and a second cluster with a second processor, “the first and second processors sharing a common virtual address space.” *Id.* at 20:50-51, claim 42.

The ’636 patent claims a cache coherency method in which “a plurality of request cluster processors in the request cluster share a memory space with a plurality of home processors in the home cluster.” ’636 patent at 20:46-48, claim 13. The ’121 patent claims do not mention shared address space at all.

The ’206 patent claims a computer system “wherein the interconnection controller in each cluster is operable to uniquely map selected ones of locally generated transactions directed to others of the clusters to a global transaction space, and remotely generated transactions

directed to the associated local nodes to the local transaction space.” ’206 patent at 19: 34-41, claim 16. Other than this mention of “global transaction space,” the ’206 patents’ claims do not describe any shared memory address space.

Claim one of the ’254 patent mentions “global memory space” several times. For example, the patent describes circuitry that selects a first protocol engine based on “destination information associated with a memory transaction of the memory transactions where the destination information corresponds to the first subset of the global memory space” and selects a second protocol engine “where the destination information corresponds to the second subset of the global memory space.” ’254 patent at 18:32-36, claim 1. The ’254 patent also claims a cluster “wherein the interconnection controller is further configured such that the second subset of the global memory space is mutually exclusive of the first subset of global memory space.” *Id.* at 18:57-60, claim 8.

The claims do not indicate whether “point-to-point architecture” requires use of shared memory address space. The specifications, however, provide some guidance. *See Retractable Techs.*, 653 F.3d at 1305 (“It is axiomatic that the claim construction process entails more than viewing the claim language in isolation. Claim language must always be read in view of the written description . . .”). The ’636, ’409, and ’121 patents state: “In a point-to-point architecture, a cluster of processors includes multiple processors directly connected to each other through point-to-point links. By using point-to-point links instead of a conventional shared bus or external network, the multiple processors are used efficiently in a system sharing the same memory space.” ’636 patents at 5:8-14; ’409 patent at 4:32-38; ’121 patent at 4:38-43. Each of these three patents also teaches that “delay in an architecture using a shared memory space is significantly less than the delay in conventional message passing environments using external

networks such as Ethernet or Token Ring.” ’636 patent at 6:6-9; ’409 patent at 5:28-31; ’121 patent at 5:35-38. The fact that the patents distinguish prior art that does not use shared memory space must inform the claim construction. *See Retractable Techs.*, 653 F.3d at 1305 (holding that the district court erred by construing “body” to include multiple pieces where the patent expressly “distinguish[ed] prior art syringes comprised of multiple pieces” and disclosed only embodiments “having a body that is a single piece”); *O.I. Corp. v. Tekmar Co.*, 115 F.3d 1576, 1581 (Fed. Cir. 1997) (limiting “passage” to non-smooth structures where the patent’s “description expressly distinguish[ed] over prior art passages by stating that those passages are generally smooth-walled”).

Additionally, the disclosures in the patents repeatedly refer to shared or global memory space. All five patents include Figure 1A, “a diagrammatic representation of one example of a multiple cluster, multiple processor system which may employ the techniques of the present invention.” *See, e.g.*, ’636 patent at 7:5-7. The patents state, “The processing clusters 101, 103, 105, and 107 are connected to each other through point-to-point links 111a-f.” *See, e.g., id.* at 7:9-11. The ’206 and ’254 patents further explain, “The multiple processors in the multiple cluster architecture shown in FIG. 1A share a global memory space.” ’206 patent at 3:15-15; ’254 patent at 3:18-19. The ’636, ’409, and ’121 patents include a similar statement: “In one embodiment, the multiple processors in the multiple cluster architecture shown in FIG. 1A share the same memory space.” ’636 patent at 7: 11-13; ’409 patent at 6:14-16; ’121 patent at 6:16-18. Further, in the ’206 and ’254 patents, Figure 11 “is an exemplary mapping of protocol engines in a processor cluster to a global memory space in a multi-cluster system.” ’206 patent at 2:42-44; ’254 patent at 2:45-47.

Yet, although the patent claims and specifications indicate that “point-to-point architecture” used by the inventions is frequently used in systems with shared or global memory space, the intrinsic evidence does not clearly indicate whether the architecture requires such a memory space. The Court therefore turns to the extrinsic evidence presented. In his declaration, Professor Mithal opined, “The manner in which processors are connected (*e.g.*, point-to-point links or some other type of connection) and the manner in which processors access data (*e.g.*, shared memory address space or non-shared memory address space) are independent choices.” Dkt. 147 ¶ 43. According to Professor Mithal, “Processors directly connected through point-to-point links may utilize a single shared memory address space or each processor may have a private memory that is not shared.” *Id.* Professor Jones disagreed based upon how the patents distinguish prior art computer systems, such as Ethernet networks, in which separate computers “typically” have their “own memory and memory address space.” Dkt. 139 ¶ 18. In his declaration, Professor Jones also emphasized the repeated references to shared or global memory space in the ’206 and ’254 patents.

The patents distinguish over prior art architecture that does not use shared or global memory address space, but nothing in the claims or specifications indicate that “point-to-point architecture” requires such a memory address space. The patents all suggest that Figure 1A, showing point-to-point architecture in a shared memory address space, is only one embodiment of the architecture, and specific embodiments described in the specifications do not limit the claims. *See SciMed Life Sys.*, 242 F.3d at 1340 (noting that “the cardinal sins of patent law” is “reading a limitation from the written description into the claims”). Moreover, even Professor Jones concedes that prior art computer systems that use other architectures only “typically” have private, as opposed to shared, memory address spaces, meaning that a particular type of memory

address space is not inherent in any specific architecture. Therefore, the Court concludes that one skilled in the art reading the claims, description, and prosecution history would understand that the term “point-to-point architecture” in the patents is not limited to a computer system using shared memory address space. As Professor Mithal explained, for the purposes of the patents-in-suit, the type of connection and the manner of accessing data are independent variables.

3. Construction

Based on the above analysis, the Court construes “point-to-point architecture” as follows: “an architecture including multiple processors that are directly connected to each other through point-to-point links.”

B. Term Two: “a cache access request”

Memory Integrity’s Proposed Construction	Intel’s Proposed Construction
“a request for access to data that may be stored in cache”	“a request for access to data stored in cache”

The term “a cache access request” appears in the asserted claims of the ’636 and ’409 patents. The parties primarily disagree on one issue: whether the requested data must necessarily be present in the cache or just *may* be present in the cache. Memory Integrity argues that a person of ordinary skill in the art would understand “a cache access request” to encompass access requests that result in “cache misses”—when the requested portion of memory (called a “line” or a “block”) is not present in the cache. Intel responds that its definition includes the possibility of a cache miss. Intel further argues that Memory Integrity’s definition improperly conflates the concepts of “a cache access request” and a “probe,” which the parties agree to define as “a mechanism for eliciting a response from a node to maintain cache coherency in a system.” Dkt. 125 at 3.

The parties agree that a person of ordinary skill in the art would understand that data sought by “a cache access request” is not guaranteed to be in the cache because the cache contains only a limited subset of the overall data in main memory. According to Intel, the plain meaning of “request” allows for a “request” that does not produce a desired response. Intel analogizes to a request for a book from a library: if the book is checked out, that result does not change the nature of the request; it was still a request for a book in the library.

1. Whether Data *May* Be Stored in Cache

The claims in the '636 and '409 patents do not define “a cache access request.” The claims do, however, recite “a cache access request” and a “probe” as distinct claim elements. *See, e.g.*, '636 patent at 20:31-40, claim 11 (referring to “receiving a cache access request from a request cluster processor” and “sending a probe associated with the cache access request”); '409 patent at 22:10-15, claim 52 (referring to “a cache access request originating from a first cluster of processors” and “sending a probe to nodes associated with the first cluster of processors”). The specifications also refer to “a cache access request” and a “probe” as separate and distinct. *See* '636 patent at 10:14-16 (“Data or cache access requests usually target the home node memory controller. Probes are used to query each cache in the system.”); '409 patent at 9:17-19 (same).

The claims indicate that a request is not the equivalent of a probe. The specifications show that requests “target” while probes “query.” The plain and ordinary meaning of the patents thus establishes that “a cache access request” is, at least typically, distinct from a probe and more than a query for data that may or may not be in cache memory. Although the request may result in a “cache miss,” the request is more definite than the definition proposed by Memory Integrity suggests.

This understanding of “cache access request” comports with extrinsic evidence. *See Omega Eng’g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1332 (Fed. Cir. 2003) (“Although expert testimony and declarations are useful to confirm that the construed meaning is consistent with the denotation ascribed by those in the field of the art, such extrinsic evidence cannot be used to vary the plain language of the patent document.” (citations omitted)). For example, in his deposition, inventor David Glasco states that a “probe” is “different” from “a cache access request” and that the two terms refer to “unique messages.” Dkt. 162-1 at 16.⁴ Additionally, Professor Mithal explains, “[E]ven in the case of a cache miss, the request itself was still a request for access to data stored in cache.” Dkt. 147 ¶ 49. Finally, Professor Jones states that “a cache access request” asks for data that “*may* be stored in the cache” or “*may not* be stored in the cache” because “there is no way to know whether or not a cache access request will miss until the cache is actually checked.” Dkt. 139 ¶ 25. Professor Jones does not, however, indicate that a “request” is the equivalent of a “probe” and did not take issue with any definition of “a cache access request” that could allow for cache misses.

Before oral argument, the Court explained to the parties in a minute order that the Court is concerned that Memory Integrity’s proposal appears to conflate request and probe, while Intel’s proposed construction does not expressly incorporate the possibility of a cache miss. In

⁴ As noted previously, *see supra* note 3, an inventor’s subjective intent is of little weight when it comes to understanding patent claims. An inventor’s testimony, however, “may be pertinent as a form of expert testimony, for example, as to understanding the established meaning of particular terms in the relevant art.” *Howmedica Osteonics*, 540 F.3d at 1346-47. As with Mr. Morton, the Court finds Mr. Glasco’s testimony pertinent as a form of expert testimony as to understanding the meaning of the disputed terms in the field. His testimony concerning “requests” as opposed to “probes” resembles the inventor testimony that the Federal Circuit found pertinent in *Apple Inc. v. Samsung Electronics Company*, 2016 WL 761884 (Fed. Cir. Feb. 26, 2016). There, the inventor testified that he understood “a shared library program and a server” to be “two different ways of implementing the function described in the [patent-in-suit].” *Id.* at *5 n.6. The Federal Circuit found that this testimony “undermin[ed] Apple’s arguments that a shared library program can be a separately running server.” *Id.*

response to the Court’s concern, the parties submitted revised proposals for this term’s construction.

Memory Integrity’s Revised Proposed Construction	Intel’s Revised Proposed Construction
“a request for access to data stored in cache, including a request that may result in a cache miss”	“a request to cache, distinct from a probe, for access to data”

Intel’s revised proposal requires that a “request” to cache always and necessarily be distinct from a probe. Although, as explained above, the specifications do indicate that the two terms are separate and distinct, the Court is unwilling to expressly import this limitation into the claims when the claims do not include the limitation. *See Specialty Composites v. Cabot Corp.*, 845 F.2d 981, 987 (Fed. Cir. 1988) (citing *Lemelson v. United States*, 752 F.2d 1538, 1551-52 (Fed. Cir. 1985)) (“Where a specification does not **require** a limitation, that limitation should not be read from the specification into the claims.” (emphasis in original)). The specifications recite only what a “request” does “usually.” *See* ’636 patent at 10:14-16. Moreover, Intel provides no prosecution history showing that Memory Integrity disavowed a “request” that may, in some limited circumstances, have overlapping features with a “probe.”

If and when a “request” may possibly act like a “probe,” while still satisfying the definition allowed by the evidence, is a factual question appropriate for the jury. *See Eon Corp. IP Holdings LLC v. Silver Spring Networks, Inc.*, 2016 WL 766661, at *3 (Fed. Cir. Feb. 29, 2016) (“[C]ourts should not resolve questions that do not go to claim scope, but instead go to infringement . . . or improper attorney argument”); *PPG Indus. v. Guardian Indus. Corp.*, 156 F.3d 1351, 1355 (Fed. Cir. 1998) (“[A]fter the court has defined the claim with whatever specificity and precision is warranted by the language of the claim and the evidence bearing on

the proper construction, the task of determining whether the construed claim reads on the accused product is for the finder of fact.”).

Memory Integrity’s revised proposal, on the other hand, avoids expressly conflating “request” and “probe” and addresses that Court’s concern that Intel’s proposed construction may suggest that a “request” that results in a cache miss is not truly a “request.” Memory Integrity’s revised proposal does no more than take Intel’s original proposal and explain what Intel admits in its reply brief: that “the desired response is not guaranteed. If a cache access request is sent to a cache without the desired data, then there will be a cache miss.” Dkt. 151 at 25. Because Memory Integrity’s revised proposed construction both differentiates between a “request” and a “probe” and allows for cache misses, the Court adopts this construction.

2. Construction

Based on the above analysis, the Court construes “a cache access request” as follows: “a request for access to data stored in cache, including a request that may result in a cache miss.”

C. Term Three: “the cache access request”

Memory Integrity’s Proposed Construction	Intel’s Proposed Construction
This term is governed by its plain and ordinary meaning. No further express construction is needed.	’409 claims 1, 6: “the cache access request that originates from the first plurality of processors and is received by the first cache coherence controller” ⁵

The term “the cache access request” appears in the asserted claims of the ’409 patent.⁶

Initially, the parties appear to disagree on whether “*the* cache access request” is *the* request

⁵ Intel originally proposed constructing “the cache access request” in these claims as “the cache access request that the first cache coherence controller receives from the first plurality of processors.” Dkt. 125 at 5-6. Intel amended its proposed construction in response to the Court’s tentative opinion, submitted to the parties on March 9, 2016. *See* Dkt. 192.

⁶ The parties initially disputed the meaning of the term “the cache [access] request” in asserted claims of both the ’636 and ’409 patents, and Intel proposed five different constructions.

referred to earlier in the claims or can be applied to requests other than the one referred to earlier in the claims—*i.e.*, whether the term should be construed as incorporating the “antecedent basis for each claim.” Dkt. 135 at 26. At a deeper level, however, the disagreement actually concerns whether “the cache access request” may undergo any sort of alteration between when a cache coherence controller first receives the request and when the controller forwards it.

1. The Antecedent Basis for the Claims

The parties agree that the term “the cache access request” has a clear antecedent basis in the patents. Intel, however, argues that Memory Integrity claims in its infringement contentions that different requests can be both “*a* cache access request” and “*the* cache access request.” According to Intel, Memory Integrity should not be allowed to use purported plain meaning to avoid a construction that makes clear that “the cache access request” refers to the same request referenced earlier in the claims.

’409 patent claim one describes a “first cache coherence controller . . . configured to receive a cache access request originating from the first plurality of processors and send a probe to the first plurality of processors in the first cluster.” ’409 patent at 17:52-54. In claim one, the term “the cache access request” follows this description once. *Id.* at 17:54-55. Claim six of the ’409 patent describes a “first cache coherence controller coupled to the second cache coherence controller and constructed to receive a cache access request originating from the first plurality of processors and send a probe to the first plurality of processors in the first cluster.” *Id.* at 18:16-18. In claim six, the term “the cache access request” follows this description once. *Id.* at 18:18-19.

After the parties’ claim construction briefing, however, Memory Integrity amended its asserted claims. Because “the cache access request” now appears only in asserted claims one and six of the ’409 patent, Intel proposes only one construction for the term.

As the Federal Circuit has noted, “[i]t is a rule of law well established that the definite article ‘the’ particularizes the subject which it precedes. It is a word of limitation as opposed to the indefinite or generalizing force of ‘a’ or ‘an.’” *Warner-Lambert Co. v. Apotex Corp.*, 316 F.3d 1348, 1356 (Fed. Cir. 2003) (quoting *Am. Bus. Ass’n v. Slater*, 231 F.3d 1, 4-5 (D.C. Cir. 2000)) (quotation marks omitted). Where separate clauses in a patent claim use identical language and contain the same term, the term carries the same meaning in both clauses despite being preceded by “a” in the first clause and “the” in the second. *See Process Control Corp. v. HydReclaim Corp.*, 190 F.3d 1350, 1356 (Fed. Cir. 1999) (“It is clear from the language of the claim itself that the term ‘a discharge rate’ in clause [b] is referring to the same rate as the term ‘the discharge rate’ in clause [d]. This conclusion necessarily results from the identical language associated with the term ‘discharge rate’ . . .”).

In the patents-in-suit, the term “the cache access request” closely follows the term “a cache access request.” Nothing in the claims’ text indicates that the definite article “the” denotes a request separate and distinct from the earlier request preceded by the indefinite article “a.” In fact, Memory Integrity concedes that “‘the request’ is the request referenced earlier in the claims.” Dkt. 146 at 17. Memory Integrity argues, however, that after the Court has construed “a cache access request,” the latter term “the cache access request” has a plain meaning that a jury will readily understand based upon the antecedent usage. Memory Integrity further argues that Intel’s construction adds structural requirements not present in the claims by suggesting that the request that the cache coherence controller receives must be identical to the cache access request that the cache coherence controller sends.

Intel counters that a determination that “the cache access request” has a plain and ordinary meaning is insufficient to resolve the parties’ dispute. *See O2 Micro Int’l Ltd. v. Beyond*

Innovation Tech. Co., 521 F.3d 1351, 1361 (Fed. Cir. 2008) (“A determination that a claim term ‘needs no construction’ or has the ‘plain and ordinary meaning’ may be inadequate when . . . reliance on a term’s ‘ordinary’ meaning does not resolve the parties’ dispute.”). According to Intel, Memory Integrity’s infringement contentions show that Memory Integrity will attempt, among other things, to argue that “a request” refers to a “Request For Ownership” while “the request” refers to a “RdInvOwn” request. *See, e.g.*, ’409 Patent Contentions, Dkt. 136-16 at 3. Additionally, Intel argues that the specification’s use of the word “forwarding” shows that “the cache access request” that is received from a request cluster is the same as the request that is sent on to a non-local cluster. *See, e.g.*, ’409 patent at 3:25-32 (describing how “[a] cache access request is received from a local processor” and the possibility of speculative probing is assessed “before forwarding the cache request to a non-local cache coherence controller”).

2. Possible Changes to “the Cache Access Request”

Although the parties agree that, in theory, “a request” and “the request” should carry the same meaning when the context of the claim indicates that this is the case, the parties continue to disagree on whether “a request” may be altered before a cache coherence controller sends “the request.” As highlighted by Memory Integrity, the text of the specifications call into question Intel’s argument that “the request” sent is always identical to the request received. The patent specification indicates that there may be some variation between the request received and the request sent because the patent specifications state, “[A]though messages associated with requests, probes, responses, and data are described as forwarded from one node to another, the messages themselves may contain variations. In one embodiment, alterations are made to messages to allow the multiple cluster architecture to be transparent to various local nodes.” *Id.* at 10:54-59.

Intel responds that the altered messages are just “associated” with the requests and that the requests themselves are not actually altered. Intel does not, however, deny that the specification describes “a request” as a “message.” *Id.* at 9:14-17 (“In one example, a coherence protocol contains four types of messages; data or cache access requests, probes, responses or probe responses, and data packets.”).

Additionally, ordinary grammatical concepts indicate that there need not be an exact correspondence between a noun preceded by the definite article “the” and the immediately antecedent use of the same noun preceded by the indefinite article “a.” An anaphor depends on an antecedent for its meaning.⁷ Laure Gardelle provides an example in her article “*Anaphora*,” “*Anaphor*” and “*Antecedent*” in *Nominal Anaphora: Definitions and Theoretical Implications*, 22 *Cercles* 2012, at 25, available at <http://www.cercles.com/n22/gardelle.pdf> (last visited March 22, 2016): “The cat jumped out of the armchair onto the table.” Here, “[t]he cat denotes that the referent is a cat, but only the previous mention of the animal can saturate the phrase and determine that the referent is the specific cat mentioned ahead in the depiction of the scene.” *Id.* Thus, an anaphor, might not simply relate back to the antecedent on a one-to-one basis because the anaphor “includes the information given before, but also that given in the clause that contains the anaphor.” *Id.* at 34.

To capture this concept, Francis Cornish proposes the term “antecedent-trigger,” which would include “no associated assumption that this formal device wholly determines the in-context sense and reference of a given anaphor.” *Anaphora, Discourse and Understanding* 42

⁷ “Anaphora” is “use of a grammatical substitute to refer to a preceding word or group of words.” *Webster’s Third New International Dictionary of the English Language Unabridged* 78 (2002). “Anaphora” also refers to “the relation between an anaphor and an antecedent, where the interpretation of the anaphor is determined via that of the antecedent.” Rodney Huddleston & Geoffrey Pullum, *The Cambridge Grammar of the English Language* 1453 (2002) (emphasis omitted).

(Clarendon Press 1999). As Cornish explains, there is “a dynamic relationship between the antecedent-trigger and anaphor” characterized by “the lack of any necessary formal parallelism between antecedent-trigger and anaphor.” *Id.* Cornish gives the following example: “Joe ate an apple last night, but *it* was much too acidic for his liking.” *Id.* at 44. (emphasis in original). “It” in the example “accesses the apple in question in its ‘eaten’ (or partially eaten), not pristine, state.” *Id.* According to Cornish, the human mind interprets “the anaphoric predication as a whole” in a way that “is integrated into the context provided by the processing of the antecedent-trigger clause, as well as earlier ones, in addition to the inferences introduced by such processing.” *Id.*

The term “a cache access request originating from the first plurality of processors” is an antecedent-trigger for “the cache access request.” That the anaphor derives meaning from the antecedent-trigger does not, however, mean that the anaphor is not also independently modified by the context in which it appears. For instance, claim one states that “the cache access request is received by a serialization point in the second cluster,” indicating that the request is both temporally and spatially removed from “a cache access request originating from the first plurality of processors.” ’409 patent at 17:52-56. Thus, “the cache access request” need not be exactly identical to the antecedent referred to by “a cache access request,” *i.e.*, the request could undergo changes after it originates from the first plurality of processors. The specifications show that variations may occur, and the concept of an antecedent-trigger allows for such variations.

Accordingly, the Court construes the term “the cache access request” as having the plain and ordinary meaning given to it by the claims. A jury can understand the scope of “the cache access request” from the text of the claims and construction of the term “a cache access request.” *See Activevideo Networks, Inc. v. Verizon Commc’ns, Inc.*, 694 F.3d 1312, 1326 (Fed. Cir. 2012)

(holding that the a district court commits no error and properly resolves disputes between parties by giving terms “plain meanings that do not require additional construction” when one party’s “proposed construction erroneously reads limitations into the claims”). Additionally, the jury, as the finder of fact, must determine precisely what sorts of alterations would so materially modify a cache access request that it could no longer properly be considered *the* cache access request that derives meaning from the antecedent-trigger. *See Eon Corp.*, 2016 WL 766661, at *3 (“[C]ourts should not resolve questions that do not go to claim scope, but instead go to infringement . . . or improper attorney argument . . .”).

3. Construction

Based on the above analysis, the Court construes “the cache [access] request” as having its plain and ordinary meaning in light of the construction of “a cache access request,” where the noun following the definite article “the” refers back to the noun, or its material equivalent, that follows the immediately antecedent use of the indefinite article “a.”

D. Term Four: “states associated with selected ones of the cache memories”

Memory Integrity’s Proposed Construction	Intel’s Proposed Construction
“cache coherence protocol states associated with selected ones of the cache memories”	Plain meaning. To the extent a construction is necessary: “status of data stored in selected ones of the cache memories”

The term “states associated with selected ones of the cache memories” appears in asserted claims of patent ’121. The parties dispute whether the term has the plain meaning of the “status of data stored in selected ones of the cache memories,” including whether the data is present in the cache, or whether the term should be limited to “cache coherence protocol states” that do not include whether the data is present in cache memory. Memory Integrity argues that a person of ordinary skill in the art would understand “state” as referring to cache coherence

protocol states. Intel responds that Memory Integrity’s construction is too narrow and that the patent never limits “state” to “cache coherence protocol states.”

1. Whether “States” Is Limited to Cache Coherence Protocol States

The ’121 patent claims a computer system that uses a probe filtering unit to receive probes “and to transmit the probes only to selected ones of the processing nodes with reference to probe filtering information representative of states associated with selected ones of the cache memories.” ’121 patent at 31: 3-7, claim 1, 32:12-15, claim 16. The patent further claims a computer system that evaluates probes to determine whether a valid copy of data is in any cache memories using “probe filtering information associated with the probe filtering unit and representative of states associated with selected ones of the cache memories.” *Id.* at 32:51-54. The claims offer no other definition of “state.”

On the other hand, the specification discusses the meaning of “state” several times. The specification describes “a coherence directory that can be used to allow management and filtering of probes.” *Id.* at 13:44-45. In an embodiment depicted by Figure 7, the directory “includes state information 713, dirty data owner information 715, and an occupancy vector 717 associated with memory lines 711.” *Id.* at 13:55-57.⁸ The specification notes, “In some embodiments, the memory line states are modified, owned, shared, and invalid.” *Id.* at 13:58-59. The states of “modified, owned, shared, and invalid” appear in column 713 of Figure 7. The specification defines “occupancy vector” as “[a]ny mechanism for tracking what clusters hold a copy of the relevant memory line in cache.” *Id.* at 14:2-4. The specification further explains that “an occupancy vector 717 can be checked to determine what caches share the relevant data.” *Id.* at 13:64-66.

⁸ The numerals 711, 713, 715, and 717 all represent separate columns in the Figure 7 table.

Memory Integrity argues that Figure 7 establishes that the term “states” refers to cache coherence protocol states, such as “modified, owned, shared, and invalid,” and does not include presence information. According to Memory Integrity, Figure 7 would not have shown a separate “occupancy vector” column if “states” conveyed information about where data is located. The Federal Circuit, however, emphasizes that courts must not confine claims to specific embodiments without evidence that “the patentee instead intends for the claims and the embodiments in the specification to be strictly coextensive.” *Phillips*, 415 F.3d at 1323. Here, the specifications do not indicate any intention to limit the claims to the embodiment in Figure 7.

The specification shows that the separate columns in Figure 7 are not mutually exclusive. Although Figure 7 also includes a separate column for “dirty data owner information,” the specification describes the “dirty” status of data as one of the claimed “states.” The specification explains how a memory controller “generates probes to all of the nodes in the cluster (including the cache coherence controller) asking whether any of the nodes have the requested memory line in their corresponding caches in either a ‘dirty’ (i.e., modified) or ‘clean’ (unmodified) state.” ’121 patent at 19:21-25. The specification goes on to identify a “‘dirty’ state” of data and clarify how the directory uses “the indicated state of the corresponding memory line, e.g., ‘clean’ vs. ‘dirty.’” *Id.* at 20:38, 22:45-46.

Moreover, the specification expressly states that invention is not limited to any particular set of “states.” Although Figure 7 “includes the four states of modified, owned, shared, and invalid,” the patent stresses that “it should be noted that particular implementations may use a different set of states. In one example, a system may have the five states of modified, exclusive, owned, shared, and invalid. The techniques of the present invention can be used with a variety of different possible memory line states.” *Id.* at 14:30-36. Memory Integrity argues that a person of

ordinary skill in the art would understand that this passage merely establishes that “states” is not limited to any particular set of *cache coherence protocol states*. Intel points out, however, that the patent uses the term “cache coherence protocol states” in other places, *see id.* at 19:1-12 (“This information corresponds to the standard coherence protocol states . . .”), but does not use that term in the passage discussing the possible use of a different set of states. This omission, argues Intel, shows that the patentees differentiated between “states” and “coherence protocol states” and that they did not intend to limit the former broad term to the latter narrow term.

The specification also contains evidence that the patentees intended “states” to include presence information. For example, the description of Figure 7 and Figure 8 reads, “According to a specific embodiment, the directory of shared states may be implemented as described above with reference to FIGS. 7 and 8, and indicates where particular memory lines are cached within the cluster.” *Id.* at 28:31-34. In addition, the specification explains that information regarding “the memory address corresponding to the cached memory line, *the remote cache location*, whether the line is ‘clean’ or ‘dirty,’ and whether the associated processor has read-only access or read/write access . . . *corresponds to the standard coherence protocol states.*” *Id.* at 18:67-19:5. These descriptions suggest “states” may include information about whether data is present in a particular cache.

Nevertheless, the claims and specification do contain some ambiguity about the breadth of the term “states.” Because “probe filtering information” is only “representative of states associated with selected ones of the cache memories,” “states” and other information, such as presence information, could be separate and distinct. *See id.* at 31: 3-7, claim 1, 32:12-15, claim 16. In addition, in one place in the specification, the patent explains that “because the cache coherence directory provides information about where memory lines are cached *as well as their*

states, probes only need be directed toward the clusters in which the requested memory line is cached.” *Id.* at 19:36-40. Information about the location of cached data could differ from the “states” of the data. The Court thus turns to the extrinsic evidence presented.

In his declaration, Professor Methal asserts that a person of ordinary skill in the art would understand “states” as referring to several different types of information “helpful for performing cache coherence functions.” Dkt. 147 ¶ 57. According to Professor Mithal, the goal of a cache coherence protocol—reducing probe traffic—could be served by “state information that consists of mere presence and non-presence information.” *Id.* ¶ 58. Professor Mithal concludes, “Thus, ‘state’ information refers to any type of information on the status of data in cache memories—including whether it is present in certain cache.” *Id.* ¶ 60.

In Professor Jones’s declaration, he disagrees with Professor Mithal, asserting that although the term “state” may have many broad meanings in the English language, “state” is understood to refer to “cache coherence protocol state” in the field of cache coherency. Dkt. 139 ¶ 31. Professor Jones cites multiple texts in support of his assertion. These texts, however, actually support Professor Mithal’s position that “states” may have many different meanings in the field of cache coherence and may include presence information. For example, Professor Jones cites Sorin, et al., *A Primer on Memory Consistency and Cache Coherence* (2011), Dkt. 138-9. The primer explains, “Many coherence protocols use a subset of the classic five state MOESI model [Modified, Owned, Exclusive, Shared, Invalid].” Dkt. 138-9 at 106-7. The invalid state may indicate that a cache does not contain a block of data, a situation that may also “be denoted as the ‘Not Present’ state.” *Id.* The primer also clarifies, “There are many possible coherence states, but we focus our attention in this primer on the well-known MOESI states.” *Id.* at 108. The primer therefore indicates that although the term “states” includes well-known sets of

coherence protocol states, the term may also generally refer to information about the status of data in a cache.⁹

Moreover, Intel submitted a supplemental declaration of Dr. Daniel J. Sorin, the author of the primer in question. In his declaration, Dr. Sorin states, “A person of ordinary skill in the art would understand the term ‘states associated with selected ones of the cache memories’ to not be limited to cache coherence protocol states, and be broad enough to include the condition of presence—i.e., what is stored in cache memory.” Dkt. 196-1 ¶ 17. Additionally, Dr. Sorin clarified that the primer “does not use the term ‘state’ to mean only a cache coherence protocol state. As examples, the book uses the terms ‘final states of the memory’ and ‘state of a register.’” *Id.* ¶ 19. Considering the extrinsic and intrinsic evidence, the Court finds that one skilled in the art would have an understanding of “states” as the status of data in cache memory. The Court construes the term beyond its plain meaning to resolve the parties’ dispute. *See O2 Micro*, 521 F.3d at 1360-61.

⁹ The Court notes the U.S. Patent and Trademark Office’s Patent Trial and Appeal Board (“PTAB”) has undertaken *inter partes* review (“IPR”) of the ’121 patent. PTAB determined “that ‘states associated with selected ones of the cache memories’ is broad enough to include the condition of presence.” IPR 2015-00159, Institution Decision, Dkt. 136-21 at 9-10. Memory Integrity argues that the Court should give the PTAB decision little weight for two reasons: (1) the decision is not binding on the Court; and (2) PTAB generally applies claim construction standards that differ from those used by district courts. *See In re Cuozzo Speed Techs., LLC*, 793 F.3d 1268, 1276 n.6 (Fed. Cir. 2015), *cert. granted sub nom. Cuozzo Speed Techs., LLC v. Lee*, 136 S. Ct. 890 (2016) (noting that “[t]he claims of an expired patent are the one exception where the broadest reasonable interpretation is not used [by PTAB and the construction principles in *Phillips*, 415 F.3d 1303, are used instead] because the patentee is unable to amend the claims.”). In its reply brief, Intel does not attempt to refute these points. The Court agrees with Memory Integrity that PTAB’s decision has little weight for purposes of construing terms in a district court, but the Court nonetheless notes that PTAB reached the same conclusion about the construction of “states associated with selected ones of the cache memories” that the Court now reaches. Courts have held that PTAB decisions may at least provide a district court with guidance. *See, e.g., Evolutionary Intelligence, LLC v. Sprint Nextel Corp.*, 2014 WL 4802426, at *4 (N.D. Cal. Sept. 26, 2014) (“While the PTAB’s constructions will not be binding on this court, the IPR will inform this court’s ultimate reasoning.”). This Court uses the PTAB decision on this issue not for guidance, but for comfort.

2. Construction

Based on the above analysis, the Court construes “states associated with selected ones of the cache memories” as follows: “status of data stored in selected ones of the cache memories.”

E. Term Five: “protocol engine”

Memory Integrity’s Proposed Construction	Intel’s Proposed Construction
“a block of hardware within one or more interconnection controllers, that manages transaction flows, examines received packet(s) and makes decisions regarding the appropriate handling of the packet and any actions/response which need to be taken”	“a component within an interconnection controller that processes transactions” ¹⁰

The term “protocol engine” appears in the asserted claims of the ’206 and ’254 patents. The parties primarily dispute whether “protocol engine” should be limited to a specific structure with specific capabilities. Memory Integrity argues that Intel’s construction of “protocol engine” is vague and fails to make any distinction between “protocol engine” and “processor.” According to Memory Integrity, Intel’s construction ignores how protocol engines manage the flow of transactions, such as by examining received packets. Intel responds that its construction is more consistent with the plain teachings of the specifications. Additionally, argues Intel, Intel’s proposed construction does not improperly conflate “processor” and “protocol engine” because a protocol engine can be implemented as a processor that processes packets.

¹⁰ Intel revised its proposed construction in response to the Court’s tentative opinion. Intel originally proposed giving the term “protocol engine” its plain meaning or, to the extent a construction is necessary, construing the term as “a component for processing transactions.” Dkt. 125 at 9.

1. Whether “Protocol Engine” Is Limited to a Specific Structure with Specific Capabilities

The '206 patent claims a computer system including “a plurality of protocol engines for processing transactions in accordance with a cache coherence protocol.” '206 patent at 18:9-11, claim 1. The term “protocol engine” also appears throughout the '206 patent claims. The '254 patent claims “a plurality of protocol engines configured to process memory transactions in accordance with a cache coherence protocol, wherein each protocol engine of said plurality of protocol engines is configured to be assigned a distinct subset of global memory space.” '254 patent at 18:15-21, claim 1. As in the '206 patent claims, the term “protocol engine” also appears throughout the '254 patent claims. The claims of neither patent offer a precise definition of “protocol engine,” but the claims repeatedly refer to “protocol engines for processing transactions” or “protocol engines configured to process [the] memory transactions.” *See, e.g.*, '206 patent at 18:12-13, claim one, 20:8-9, claim 21, 20:46-47, claim 30.

The specifications provide additional meanings for “protocol engine.” The description of Figure 3, contained in both patents, explains that in “various embodiments, the interconnection controller includes a protocol engine 305 configured to handle packets such as probes and requests received from processors in various clusters of a multi-processor system.” '206 patent at 4:47-50; '254 patent at 4:50-54. The specifications also state:

According to various embodiments of the invention, protocol engines are blocks of hardware on the interconnection controller ASIC chip. The functionality of these engines are governed by associated microcode and relate generally to managing transaction flows. Generally speaking, a protocol engine looks at each packet received by the interconnection controller and makes decisions regarding the appropriate handling of the packet and any actions/response which need to be taken.

'206 patent at 12:11-20; '254 patent at 12:14-22. Additionally, in Figure 9 of both patents, “each of the protocol engines comprises substantially identical hardware blocks, each being

programmed by the associated microcode to perform a particular type of transaction processing.”
'206 patent at 12:26-31; '254 patent at 12:28-32.

Intel raises several issues regarding Memory Integrity's proposed construction. First, Intel argues that neither the claims nor the specifications require that protocol engines be particular blocks of hardware. The specifications establish that “the interconnection controller [in which the claimed protocol engine appears] may be represented (without limitation) in software (object code or machine code)” and that protocol engine functionality is “governed by associated microcode.” '206 at 12:11-15, 17:49-52; '254 patent at 12:14-17, 17:51-53. The specifications also emphasize, “It should be understood that the various embodiments of the invention may be implemented or represented in a wide variety of ways without departing from the scope of the invention.” '206 patent at 17:46-49; '254 patent at 17:48-51.

Second, the specifications do not support Memory Integrity's assertion that a protocol engine can exist “within *one or more* interconnection controllers.” The claims repeatedly refer to “the interconnection controller” that contains “a plurality of protocol engines” or “a protocol engine.” '206 patent at 20:8-7, 21:17-18; '206 patent 18:16-17, 18:54-55. Additionally, in these portions of the specifications, the embodiments describe only “the interconnection controller.”

Finally, Intel argues that Memory Integrity's proposed construction arbitrarily adds and removes language from the descriptions of protocol engines in the specifications. For example, Memory Integrity includes “blocks of hardware” from one passage but omits the description that follows, “on the interconnection controller ASIC chip.” The only consistent description in the specifications is that protocol engines “process transactions” or are “for processing transactions.” '206 patent at 1:47-65; '254 patent at 1:50-67.

Intel's points are well taken. The claims and specifications indicate that the protocol engines process transactions and that a "protocol engine" is located within a single interconnection controller. The claims and specifications never require that a "protocol engine" be represented as "a block of hardware." To the contrary, the specifications expressly give "various alternatives" that are within the scope of the invention, such as an entire interconnection controller represented "in software (object code or machine code)," "a hardware description language," or "partially or completely realized semiconductor devices." '206 at 17: 49-57; '254 patent at 51-59. In response to Intel, Memory Integrity argues that Intel's proposed construction is vague and conflates "protocol engine" and "processor." Memory Integrity does not, however, identify any claim or written description specifying that a "protocol engine" may never be implemented as a processor. *See In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994) ("Although an inventor is indeed free to define the specific terms used to describe his or her invention, this must be done with reasonable clarity, deliberateness, and precision.").

PTAB also considered Memory Integrity's proposed construction of "protocol engine" as part of the IPR of the '254 and '206 patents. Ultimately, PTAB construed the term as "a functional block for processing transactions." Dkt. 201-1 at 9-11; Dkt. 201-2 at 9-11. According to PTAB, "[Memory Integrity's] proposed construction is an amalgamation of descriptions and functions associated with exemplary protocol engines from various parts of the Specification." Dkt. 201-2 at 10. PTAB emphasized that it was "mindful that, although the claims are interpreted in light of the Specification—including the portions of the Specification cited by [Memory Integrity]—limitations from the Specification are not read into the claims." *Id.* As noted above, PTAB decisions are not binding on the Court, but they may nonetheless "inform" the Court's reasoning. *Evolutionary Intelligence*, 2014 WL 4802426, at *4. The Court finds PTAB's

interpretation of the specifications instructive and consistent with Intel’s proposed construction. Like PTAB, the Court uses the general explanation found in the specifications that protocol engines are “for processing transactions” to construe the term “protocol engine” as an ordinary person skilled in the art would have understood it at the time of the patents-in-suit.

2. Construction

Based on the above analysis, the Court construes “protocol engine” as follows: “a component within an interconnection controller that processes transactions.”

F. Terms Six through Nine: “remote protocol engine” / “local protocol engine” / “a first protocol engine . . .” / “a second protocol engine . . .”¹¹

Term	Memory Integrity’s Proposed Construction	Intel’s Proposed Construction
“remote protocol engine”	“a protocol engine responsible for processing transactions which target remote memory”	“a protocol engine responsible for processing transactions which target remote as opposed to local memory”
“local protocol engine”	“a protocol engine responsible for processing transactions which target local memory”	“a protocol engine responsible for processing transactions which target local as opposed to remote memory”
“a first protocol engine configured to be assigned a first subset of the global memory space, said first subset of the global memory space corresponding to one of local and remote memory”	This term is governed by its plain and ordinary meaning. No further express construction is needed.	“a first protocol engine configured to be assigned addresses for one of either local or remote memory in a global memory space”

¹¹ Although Intel addressed the “[remote/local] protocol engine” and “[first/second] engine” terms separately in its opening brief, Memory Integrity addressed all four terms together. Because Memory Integrity’s arguments in favor of its proposed constructions and against Intel’s cut across all four terms, Intel addressed the four terms together in its reply brief. For the same reasons, the Court addresses all four terms together here.

“a second protocol engine configured to be assigned a second subset of the global memory space, said second subset of the global memory space corresponding to one of local and remote memory”	This term is governed by its plain and ordinary meaning. No further express construction is needed.	“a second protocol engine configured to be assigned addresses for one of either local or remote memory in a global memory space”
--	---	--

The terms “remote protocol engine” and “local protocol engine” appear in the asserted claims of the ’206 patent. The terms “first protocol engine configured to be assigned a first subset of the global memory space, said first subset of the global memory space corresponding to one of local and remote memory” and “second protocol engine configured to be assigned a second subset of the global memory space, said second subset of the global memory space corresponding to one of local and remote memory” appear in the asserted claims of the ’254 patent. With regard to all four terms, the parties dispute whether the separate claimed protocol engines may each process *both* remote and local memory, as Memory Integrity proposes, or whether the claimed protocol engines may process either transactions that target remote memory *or* transactions that target local memory *but not both*, as Intel proposes.

1. Remote Memory vs. Local Memory

The ’206 patent claims “a computer system comprising one or more processor clusters” in which each cluster includes an interconnection controller “compris[ing] a plurality of protocol engines for processing transactions in accordance with a cache coherence protocol, wherein the plurality of protocol engines in each interconnection controller includes *at least one remote protocol engine for processing transactions targeting remote memory and at least one local protocol engine for processing transactions targeting local memory.*” ’206 patent at 18:5-15, claim 1 (emphasis added); *see also id.* at 19:5-10, claim 12, 20:10-13, claim 21, 21:20-22, claim 39.

The '254 patent claims a cluster comprising, among other things:

the interconnection controller including a plurality of protocol engines configured to process memory transactions in accordance with a cache coherence protocol, wherein each protocol engine of said plurality of protocol engines is configured to be assigned a distinct subset of a global memory space, the plurality of protocol engines further comprising: (a) ***a first protocol engine configured to be assigned a first subset of the global memory space, said first subset of the global memory space corresponding to one of local and remote memory***; and (b) ***a second protocol engine configured to be assigned a second subset of the global memory space, said second subset of the global memory space corresponding to one of local and remote memory***

'254 patent at 18:13-28, claim 1 (emphasis added). The '254 patent further claims “[t]he cluster of claim 1, wherein the interconnection controller is further configured such that the second subset of the global memory space is mutually exclusive of the first subset of the global memory space.” *Id.* at 18:57-60, claim 8.

The claims do not expressly clarify whether a “remote protocol engine” may process local memory in some cases or whether a “local protocol engine” may process remote memory in some cases. Memory Integrity argues that the claims’ use of the word “comprising” suggests that the claims may cover instrumentalities that perform additional functionalities, even if the additional functionalities are otherwise inconsistent with some limitations of the claim. *See, e.g., Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 811 (Fed. Cir. 1999) (“The signal ‘comprising’ implements the general rule that absent some special circumstance or estoppel which excludes the additional factor, infringement is not avoided by the presence of elements or steps in addition to those specifically recited in the claim.”). Intel responds that it does not seek to exclude different or additional functionalities. For example, local protocol engines could process local transactions as well as process interrupts. *See* '206 patent at 12:45-51; '254 patent at 12:46-52. According to Intel, it seeks only to exclude functionalities that the claims

themselves exclude, *i.e.*, “remote protocol engines” do not process local memory transactions, and “local protocol engines” do not process remote memory transactions.

The specifications provide some assistance in resolving this dispute. Although specific embodiments do not generally limit the claims, in every disclosed embodiment and every figure depicting multiple protocol engines in the patents, the protocol engines are assigned either local or remote memory. For example:

- Remote memory protocol engines (RMPEs) 1002 are responsible for processing transactions which target remote memory, *i.e.*, memory associated with another cluster, and all subsequent transactions in that particular transaction flow. Local memory protocol engines (LMPEs) 1004 are responsible for processing transactions which target local memory, *i.e.*, memory associated with the local cluster with which the interconnection controller is associated, and all subsequent transactions in the flow. ’206 patent at 12:41-49; ’254 patent at 12:42-50.
- Where the packet targets a memory address associated with a remote cluster, the packet is directed to one of the remote protocol engines associated with the local interconnection controller. Where the packet targets a memory address associated with the local cluster, the packet is directed (*e.g.*, in the case of a broadcast packet like a probe) to one of the local protocol engines associated with the local interconnection controller. ’206 patent at 19-26; ’254 patent at 16:23-30.
- As discussed above, if the target address corresponds to the local memory, the packet is mapped to the corresponding one of the local protocol engines. If, on the other hand, the target address corresponds to remote memory, the packet is mapped to the corresponding one of the remote protocol engines. ’206 patent at 16:47-52; ’254 patent at 16:51-56.

Memory Integrity cites *Linear Technology Corp. v. International Trade Commission*, 566 F.3d 1049, 1055 (Fed. Cir. 2009), in support of its conclusion that reciting a “remote protocol engine” separately from a “local protocol engine” and a “first protocol engine” separately from a “second protocol engine” does not require that corresponding structures be separate and distinct.

In *Linear Technology*, the Federal Circuit addressed the claim construction of a “second circuit” and a “third circuit” for controlling a switching voltage regulator. *Id.* The Federal Circuit found that the two terms did “not require entirely separate and distinct circuits.” *Id.* The patent required only that the “second” and “third” circuits “perform their stated functions.” *Id.* The Federal Circuit reached this holding because “there is nothing in the claim language or specification that supports narrowly construing the terms to require a specific structural requirement or entirely distinct ‘second’ and ‘third’ circuits.” *Id.* The “specification expressly disclose[d] that the ‘second circuit’ and ‘third circuit’ can share common components.” *Id.*

Here, in contrast to the facts of *Linear Technology*, the claim language and specifications repeatedly indicate that a “remote protocol engine” and a “first protocol engine” are separate and distinct from a “local protocol engine” and a “second protocol engine.” Use of both a remote and local processor was designed to solve a specific problem: “Having a single protocol engine to manage transactions . . . can be a transaction processing bottleneck.” ’206 patent at 12:20-22; ’254 patent at 12:22-24. The solution of assigning remote memory transactions to remote protocol engines and local memory transactions to local protocol engines only makes sense if the two different types of protocol engines each process a single type of memory transaction. The situation is akin to a city traffic planner deciding to solve the problem of a congested two-way street by creating two one-way streets, one northbound and one southbound. It would make little sense, in light of the problem to be solved, to say that the northbound street could also accommodate southbound traffic and the southbound street could also accommodate northbound traffic.

The prosecution history of the patents also indicates that the patentees intended a “remote protocol engine” and a “local protocol engine” to perform separate functions. *See Southwall*

Techs., Inc. v. Cardinal IG Co., 54 F.3d 1570, 1576 (Fed. Cir. 1995) (“The prosecution history limits the interpretation of claim terms so as to exclude any interpretation that was disclaimed during prosecution.”). The original claims of the ’206 patent recited “a plurality of protocol engines for processing transactions, and wherein at least one of the interconnection controller and the local nodes in each cluster is operable to map the transactions to the protocol engines according to destination information associated with the transactions.” ’206 File History, Application, Dkt. 136-27 at 34. The examiner rejected this claim as obvious in light of prior art computer systems comprising a plurality of processor clusters in which at least one of the interconnection controllers in each node mapped transactions according to destination data associated with the transactions. To overcome the rejection, the patentees amended the application to add “wherein the plurality of protocol engines includes at least one remote protocol engine for processing transactions targeting remote memory and at least one local protocol engine for processing transactions targeting local memory.” ’206 File History, Examiner-Initiated Interview Summary, Dkt. 136-29 at 1.

Similarly, in amending the ’254 patent application to overcome an identical rejection, the patentees highlighted “the unique nature of the multiple protocol engines of the interconnection controller.” ’254 File History, April 18, 2014 Non-Final Response, Dkt. 136-32 at 10. According to the patentees, the protocol engines were unique based on the protocol engines’ assignment to a subset of a global memory space that was “local or remote in nature” and that “may be referred to as disjoint or even ‘mutually exclusive’ to one another.” *Id.* The patentees explained that the prior art references “d[id] not teach the type of multiple protocol engines as clarifyingly claimed here.” *Id.* Intel argues that these statements in the prosecution history limit a “remote protocol

engine” and a “first protocol engine” to processing remote memory transactions and a “local protocol engine” and a “second protocol engine” to processing local memory transactions.

Memory Integrity responds that the prosecution history does not require limiting the claims as Intel proposes. According to Memory Integrity, nothing about the examiner’s rejections, the prior art at issue, or Memory Integrity’s arguments in support of the amended claims supports limiting remote protocol engines to only processing remote memory transactions and local protocol engines to only processing local memory transactions. Memory Integrity points to a statement in the ’254 File History in which the patentees explained that each protocol engine “may be assigned its own distinct subset of global memory space (e.g. dedicated *local and/or remote memory*).” Dkt. 136-32 at 7 (emphasis added). Memory Integrity further emphasizes that in one of the passages quoted by Intel, the patentees stated only that the subsets of global memory space assigned to protocol engines “*may* be referred to as disjoint or even ‘mutually exclusive’ to one another.” *Id.* at 10 (emphasis added). In other words, argues Memory Integrity, the local or remote subsets may also overlap.

The prosecution history is conflicting. The patentees distinguished prior art by claiming a unique type of protocol engine assigned to “local or remote” global memory space. *Id.* at 10. Yet the patentees also described the memory space assigned to a protocol engine as “local and/or remote.” *Id.* at 7. Additionally, the statement that the global memory space subsets assigned to protocol engines *may* be disjoint or mutually exclusive could indicate that: (1) remote and local protocol engines are assigned overlapping subsets of both remote and local memory; or (2) remote protocol engines are assigned overlapping subsets of remote memory and local protocol engines are assigned overlapping subsets of local memory, but no protocol engine has a

subset of *both* remote and local memory. In light of these uncertainties in the intrinsic evidence, the Court turns to extrinsic evidence to determine the breadth of the claims.

In his deposition, Mr. Glasco, one of the named inventors of the '206 and '254 patents, testifies concerning the meaning of “remote protocol engine,” “local protocol engine,” and the “one of” first and second protocol engine terms. He testifies that the protocol engines of the invention process transactions targeting remote or local memory, but not both. According to Mr. Glasco, “The remote protocol engines and the local protocol engines do not process both protocol flows.” Dkt. 162-1 at 10. The inventors separated the protocol engines based on “protocol flows dealing with remote memory and protocol flows dealing with local memory” in order to “simplify the design of the protocol engine to only have to deal with one or the other.” *Id.* at 7. The separation simplified the design, in part, because “there were some secondary behaviors of the protocol engines that were . . . specific to remote and specific to local that we could also separate out.” *Id.* at 12. In addition, Mr. Glasco asserts that “[t]he decision to have multiple protocol engines and the decision to split remote and local protocol flows is somewhat robotical,” indicating that one skilled in the art would understand remote and local protocol engines to have separate functions. Mr. Glasco’s testimony clarifies the intrinsic evidence and shows that the ordinary meaning of the terms is that remote protocol engines correspond to remote—not local—memory transactions and local protocol engines correspond to local—not remote—memory transactions.

2. Construction

Based on the above analysis, the Court adopts the following constructions: “remote protocol engine” is construed as “a protocol engine responsible for processing transactions that target remote as opposed to local memory”; “local protocol engine” is construed as “a protocol engine responsible for processing transactions that target local as opposed to remote memory”;

“a first protocol engine configured to be assigned a first subset of the global memory space, said first subset of the global memory space corresponding to one of local and remote memory” is construed as “a first protocol engine configured to be assigned addresses for one of either local or remote memory in a global memory space”; and “a second protocol engine configured to be assigned a second subset of the global memory space, said second subset of the global memory space corresponding to one of local and remote memory” is construed as “a second protocol engine configured to be assigned addresses for one of either local or remote memory in a global memory space.”

ORDER

For the foregoing reasons, the Court adopts the following construction of the disputed terms in patents '636, '409, '121, '206, and '254.

Term	Construction
“point-to-point architecture” (all patents-in-suit)	“an architecture including multiple processors that are directly connected to each other through point-to-point links”
“a cache access request” ('636 and '409 patents)	“a request for access to data stored in cache, including a request that may result in a cache miss”
“the cache access request” ('409 patent)	“plain and ordinary meaning in light of the construction of ‘a cache access request,’ where the noun following the definite article ‘the’ refers back to the noun, or its material equivalent, that follows the immediately antecedent use of the indefinite article ‘a’”
“states associated with selected ones of the cache memories” ('121 patent)	“status of data stored in selected ones of the cache memories”
“protocol engine” ('206 and '254 patents)	“a component within an interconnection controller that processes transactions”
“remote protocol engine” ('206 and	“a protocol engine responsible for

'254 patents)	processing transactions that target remote as opposed to local memory”
“local protocol engine” (’206 and ’254 patents)	“a protocol engine responsible for processing transactions that target local as opposed to remote memory”
“a first protocol engine configured to be assigned a first subset of the global memory space, said first subset of the global memory space corresponding to one of local and remote memory” (’206 and ’254 patents)	“a first protocol engine configured to be assigned addresses for one of either local or remote memory in a global memory space”
“a second protocol engine configured to be assigned a second subset of the global memory space, said second subset of the global memory space corresponding to one of local and remote memory” (’206 and ’254 patents)	“a second protocol engine configured to be assigned addresses for one of either local or remote memory in a global memory space”

IT IS SO ORDERED.

DATED this 22nd day of March, 2016.

/s/ Michael H. Simon
Michael H. Simon
United States District Judge