THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

VANTAGE POINT TECHNOLOGY, INC.,	§	
	§	
Plaintiff,	§	
	§	
v.	§	CASE NO. 2:13-CV-909-JRG
	§	
AMAZON.COM, INC., et al.,	§	
	§	
Defendants.	§	
	§	
	§	

MEMORANDUM OPINION AND ORDER

Before the Court are Plaintiff Vantage Point Technology, Inc.'s Opening Brief on Claim Construction (Dkt. No. 192), Defendants' response (Dkt. No. 197), and Plaintiff's reply (Dkt. No. 201). The Court held a claim construction hearing on January 29, 2015.

Table of Contents

I. BACKGROUND 3

II. LEGAL PRINCIPLES 4

III. CONSTRUCTION OF AGREED TERMS 8

IV. CONSTRUCTION OF DISPUTED TERMS IN THE '329 PATENT 9

A. "private processor cache, the contents of each such private cache being unknown externally to its associated microprocessor" 9

B. "external tag memory" 14

C. "external tag memory non-hardwired to the private processor cache" 17

D. "deriving the status of the private processor cache associated with the first processor cluster by tracking evicted cache lines and data entering and exiting the first processor cluster" 23

E. "tag controller" 26

F. "when the processor bus is idling" 28

G. "to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" 29

H. "a snooper for snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" 31

I. "processor cluster" 37

V. CONSTRUCTION OF DISPUTED TERMS IN THE '750 PATENT 40

A. "[first/second] instruction pipeline" 40

B. "master translation memory" 44

C. "direct address translation unit" 47

D. "storing the translation data for the [first/second/third] virtual address from the master translation memory into the [first/second/third] translation buffer" 51

E. "whenever translation data for the [first/second] virtual address from the master translation memory is stored into the [first/second] translation buffer" 54

VI. CONCLUSION 55

I. BACKGROUND

Plaintiff brings suit alleging infringement of United States Patents No. 5,463,750 ("the '750 Patent") and 6,374,329 ("the '329 Patent"). Dkt. No. 192, Exs. A & C. The '750 Patent is asserted against all Defendants. The '329 Patent is asserted against Defendants Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., Samsung Telecommunications America, LLC, and MediaTek USA Inc. Dkt. No. 197 at 1 n.1.

The '750 Patent is titled "Method and Apparatus for Translating Virtual Addresses in a Data Processing System Having Multiple Instruction Pipelines and Separate TLB's for Each Pipeline." The '750 Patent issued on October 31, 1995, and bears a filing date of November 2, 1993. The Abstract states:

> A computing system has multiple instruction pipelines, wherein one or more pipelines require translating virtual addresses to real addresses. A TLB is provided for each pipeline requiring address translation services, and an ad[d]ress translator is provided for each such pipeline for translating a virtual address rec[ei]ved from its associated pipeline into corresponding real addresses. Each address translator comprises a translation buffer accessing circuit for accessing the TLB, a translation indicating circuit for indicating whether translation data for the virtual address is stored in the translation buffer, and an update control circuit for activating the direct address translation circuit when the translation data for the virtual address is not stored in the TLB. The update control circuit also stores the translation data retrieved from the main memory into the TLB. If it is desired to have the same translation information available for all the pipelines in a group, then the update control circuit also updates all the other TLB's in the group.

The '329 Patent is titled "High-Availability Super Server." The '329 Patent issued on April 16, 2002, and bears a priority date of February 20, 1996. The Abstract states:

> A high-availability parallel processing server has multiple processors that are grouped into processor clusters and a plurality of memory segments. Each cluster may have up to four processors, and there may be up to five clusters of processors. Each of the processor clusters has dedicated memory buses for communicating with each of the memory segments. The server may be designed to maintain coherent interaction between all processor clusters and the memory segments.

The Court construed various terms in the '329 Patent in *Intergraph Hardware Techs. Co., Inc. v. Hewlett-Packard Co.*, No. 6:04-CV-214, Dkt. No. 93 (E.D. Tex. Dec. 22, 2004) (Davis, J.) (attached to Defendants' response brief as Ex. 1) (hereinafter, *Intergraph*). The *Intergraph* case ended in a settlement approximately one month after the Court entered its claim construction order. *See id.*, Dkt. Nos. 101–103 in 6:04-CV-214.

II. LEGAL PRINCIPLES

It is understood that "[a] claim in a patent provides the metes and bounds of the right which the patent confers on the patentee to exclude others from making, using or selling the protected invention." *Burke, Inc. v. Bruno Indep. Living Aids, Inc.*, 183 F.3d 1334, 1340 (Fed. Cir. 1999). Claim construction is clearly an issue of law for the court to decide. *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 970–71 (Fed. Cir. 1995) (en banc), *aff'd*, 517 U.S. 370 (1996).

To ascertain the meaning of claims, courts look to three primary sources: the claims, the specification, and the prosecution history. *Markman*, 52 F.3d at 979. The specification must contain a written description of the invention that enables one of ordinary skill in the art to make and use the invention. *Id.* A patent's claims must be read in view of the specification, of which they are a part. *Id.* For claim construction purposes, the description may act as a sort of dictionary, which explains the invention and may define terms used in the claims. *Id.* "One purpose for examining the specification is to determine if the patentee has limited the scope of the claims." *Watts v. XL Sys., Inc.*, 232 F.3d 877, 882 (Fed. Cir. 2000).

Nonetheless, it is the function of the claims, not the specification, to set forth the limits of the patentee's invention. Otherwise, there would be no need for claims. *SRI Int'l v. Matsushita*

Elec. Corp., 775 F.2d 1107, 1121 (Fed. Cir. 1985) (en banc). The patentee is free to be his own lexicographer, but any special definition given to a word must be clearly set forth in the specification. *Intellicall, Inc. v. Phonometrics, Inc.*, 952 F.2d 1384, 1388 (Fed. Cir. 1992). Although the specification may indicate that certain embodiments are preferred, particular embodiments appearing in the specification will not be read into the claims when the claim language is broader than the embodiments. *Electro Med. Sys., S.A. v. Cooper Life Sciences, Inc.*, 34 F.3d 1048, 1054 (Fed. Cir. 1994).

This Court's claim construction analysis is substantially guided by the Federal Circuit's decision in *Phillips v. AWH Corporation*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). In *Phillips*, the court set forth several guideposts that courts should follow when construing claims. In particular, the court reiterated that "the claims of a patent define the invention to which the patentee is entitled the right to exclude." *Id.* at 1312 (quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1115 (Fed. Cir. 2004)). To that end, the words used in a claim are generally given their ordinary and customary meaning. *Id.* The ordinary and customary meaning of a claim term "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. This principle of patent law flows naturally from the recognition that inventors are usually persons who are skilled in the field of the invention and that patents are addressed to, and intended to be read by, others skilled in the particular art. *Id.*

Despite the importance of claim terms, *Phillips* made clear that "the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification." *Id.* Although the claims themselves may provide guidance as to the meaning of

particular terms, those terms are part of "a fully integrated written instrument." *Id.* at 1315 (quoting *Markman*, 52 F.3d at 978). Thus, the *Phillips* court emphasized the specification as being the primary basis for construing the claims. *Id.* at 1314–17. As the Supreme Court stated long ago, "in case of doubt or ambiguity it is proper in all cases to refer back to the descriptive portions of the specification to aid in solving the doubt or in ascertaining the true intent and meaning of the language employed in the claims." *Bates v. Coe*, 98 U.S. 31, 38 (1878). In addressing the role of the specification, the *Phillips* court quoted with approval its earlier observations from *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1250 (Fed. Cir. 1998):

> Ultimately, the interpretation to be given a term can only be determined and confirmed with a full understanding of what the inventors actually invented and intended to envelop with the claim. The construction that stays true to the claim language and most naturally aligns with the patent's description of the invention will be, in the end, the correct construction.

Phillips, 415 F.3d at 1316. Consequently, *Phillips* emphasized the important role the specification plays in the claim construction process.

The prosecution history also continues to play an important role in claim interpretation. Like the specification, the prosecution history helps to demonstrate how the inventor and the United States Patent and Trademark Office ("PTO") understood the patent. *Id.* at 1317. Because the file history, however, "represents an ongoing negotiation between the PTO and the applicant," it may lack the clarity of the specification and thus be less useful in claim construction proceedings. *Id.* Nevertheless, the prosecution history is intrinsic evidence that is relevant to the determination of how the inventor understood the invention and whether the inventor limited the invention during prosecution by narrowing the scope of the claims. *Id.*; *see Microsoft Corp. v. Multi-Tech Sys., Inc.*, 357 F.3d 1340, 1350 (Fed. Cir. 2004) (noting that "a

patentee's statements during prosecution, whether relied on by the examiner or not, are relevant to claim interpretation").

Phillips rejected any claim construction approach that sacrificed the intrinsic record in favor of extrinsic evidence, such as dictionary definitions or expert testimony. The *en banc* court condemned the suggestion made by *Texas Digital Systems, Inc. v. Telegenix, Inc.*, 308 F.3d 1193 (Fed. Cir. 2002), that a court should discern the ordinary meaning of the claim terms (through dictionaries or otherwise) before resorting to the specification for certain limited purposes. *Phillips*, 415 F.3d at 1319–24. According to *Phillips*, reliance on dictionary definitions at the expense of the specification had the effect of "focus[ing] the inquiry on the abstract meaning of words rather than on the meaning of claim terms within the context of the patent." *Id.* at 1321. *Phillips* emphasized that the patent system is based on the proposition that the claims cover only the invented subject matter. *Id.*

Phillips does not preclude all uses of dictionaries in claim construction proceedings. Instead, the court assigned dictionaries a role subordinate to the intrinsic record. In doing so, the court emphasized that claim construction issues are not resolved by any magic formula. The court did not impose any particular sequence of steps for a court to follow when it considers disputed claim language. *Id.* at 1323–25. Rather, *Phillips* held that a court must attach the appropriate weight to the intrinsic sources offered in support of a proposed claim construction, bearing in mind the general rule that the claims measure the scope of the patent grant.

In general, prior claim construction proceedings involving the same patents-in-suit are "entitled to reasoned deference under the broad principals of *stare decisis* and the goals articulated by the Supreme Court in *Markman*, even though *stare decisis* may not be applicable *per se*." *Maurice Mitchell Innovations, LP v. Intel Corp.*, No. 2:04-CV-450, 2006 WL 1751779,

at *4 (E.D. Tex. June 21, 2006) (Davis, J.); *see TQP Development, LLC v. Inuit Inc.*, No. 2:12-CV-180, 2014 WL 2810016, at *6 (E.D. Tex. June 20, 2014) (Bryson, J.) ("[P]revious claim constructions in cases involving the same patent are entitled to substantial weight, and the Court has determined that it will not depart from those constructions absent a strong reason for doing so.").

The Court nonetheless conducts an independent evaluation during claim construction proceedings. *See, e.g.*, *Texas Instruments, Inc. v. Linear Techs. Corp.*, 182 F. Supp. 2d 580, 589–90 (E.D. Tex. 2002); *Burns, Morris & Stewart Ltd. P'ship v. Masonite Int'l Corp.*, 401 F. Supp. 2d 692, 697 (E.D. Tex. 2005); *Negotiated Data Solutions, Inc. v. Apple, Inc.*, No. 2:11-CV-390, 2012 WL 6494240, at *5 (E.D. Tex. Dec. 13, 2012).

III. CONSTRUCTION OF AGREED TERMS

The Court hereby adopts the following agreed constructions:

Term	Agreed Construction
"associated with [the/a] first processor cluster"	**"associated with only that processor cluster, not any other cluster"**
"checking the tag memory to determine if a response to the request corresponds to data in the private processor cache associated with the first processor cluster"	**"the tag controller determines whether the information in a request from a second cluster matches a tag stored in the tag memory. If there is a match, the tag controller concludes that a response to the request would require data that is only stored in a private processor cache that is in the same cluster as the tag controller and tag memory"**
"external cache controller"	**"hardware that is situated apart from the private processor cache, which includes, but is not limited to, an external tag memory, a data-request monitor, and a tag controller"**
Preamble of Claim 8 of the '750 Patent	**Limiting**

Joint Claim Construction and Prehearing Statement at 1, Dkt. No. 173, Ex. A; Dkt. No. 197 at 2 n.2; *see id.* at Ex. 3; *see also* Joint Claim Construction Chart at 1–2, 4 & 6, Dkt. No. 203, Ex. A.

IV. CONSTRUCTION OF DISPUTED TERMS IN THE '329 PATENT

A. "private processor cache, the contents of each such private cache being unknown externally to its associated microprocessor"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"a cache within and dedicated to a processor whose contents cannot be read directly by any device external to that processor, including at least the external cache controller and external tag memory" Dkt. No. 192 at 5.	"a cache within and dedicated to a processor whose contents are unknown to any device external to that processor, including at least the external cache controller and external tag memory" Dkt. No. 197 at 2.

This disputed term appears in Claim 1 of the '329 Patent. Dkt. No. 197 at 2.

> *Intergraph* found as follows, in full, regarding this disputed term:
>
> The Court construes "cache" to mean "a small portion of high-speed memory used for temporary storage of frequently-used data, instructions, or operands." *IEEE Standard Dictionary of Elec. & Elecs. Terms*, at 124 (6th ed. 1996). The Court declines to construe the remainder of the term because its meaning is clear on its face.
>
> HP's proposed construction requires that the cache be dedicated to one single processor. HP supports this construction by arguing that Intergraph acted as its own lexicographer during prosecution. Contrary to HP's argument, Intergraph did not make the statements HP points to in the prosecution history to overcome the examiner's rejection because the language was "indefinite." Instead, Intergraph was responding to an objection that the written description did not convey to one skilled in the art that the claim's subject matter was in the inventor's possession. Thus, Intergraph was not acting as its own lexicographer and was not attempting to limit the claim beyond the claim terms' plain meaning.

Intergraph at 10.

(1) The Parties' Positions

Plaintiff argues that the patentee "has acted as his own lexicographer in assigning a meaning to the term 'being unknown externally' that is not just the ordinary meaning of those words." Dkt. No. 192 at 5. Plaintiff explains that its proposal is consistent with the "stated purpose of the invention," namely "to provide a mechanism whereby the contents of the private processor cache of one processor cluster can be known by other processor clusters." *Id.* at 5–6. Plaintiff concludes that "[i]t is not accurate to say that that the contents of the cache are unknown to any external device when the point of the invention is in part to make those contents known." *Id.* at 7.

Defendants respond: "the word 'unknown' requires no construction. [Plaintiff] contends it should be rewritten in a way that contradicts its plain meaning and conflicts directly with the patentee's own statements." Dkt. No. 197 at 2. Defendants emphasize that "unknown" does not appear in the written description, and Defendants submit that "[a]lthough used during prosecution, the patentee neither expressed an intent to define 'unknown' nor equated it with 'cannot be read directly.'" *Id.* at 3. Defendants urge that Plaintiff's proposal must be rejected because "[t]he meaning of 'unknown' is unequivocal; it does not encompass indirect knowledge." *Id.*

As to Plaintiff's reliance on the "stated purpose of the invention," Dkt. No. 192 at 6, Defendants respond that "the purpose of claim 1 is to use an 'external-cache controller' to derive the *status* of private processor cache (*e.g.*, whether a cache line is in a 'modified' state), the *contents* of which is unknown." Dkt. No. 197 at 3. Finally, Defendants argue that the patentee relied upon "unknown" to distinguish prior art during prosecution. *Id.* at 4.

Plaintiff replies: "Defendants selectively rely on the language of claim 1 describing the external-cache controller as being used to derive the status of the private processor cache and assert that 'does not mean that the cache content becomes known.' Dkt. No. 197 at 3. Defendants ignore the remaining claim language in claim 1 that explicitly states that such derivation is achieved by 'tracking . . . data entering and exiting the private processor cache.'" Dkt. No. 201 at 1. Plaintiff also argues that the prosecution history is consistent with Plaintiff's proposal. *Id.* at 2.

At the January 29, 2015 hearing, Defendants submitted that although the specification discloses tracking "contents," such tracking refers to knowing the status of cache contents rather than knowing the contents itself. Defendants also cited the statement by Plaintiff during the claim construction hearing in *Intergraph* that: "The contents of the private cache, the memory that is within that microprocessor chip or associated to the microprocessor, the contents of that memory is unknown externally to anything but its associated microprocessor." *See* Dkt. No. 197, Ex. 2 at 55:22–56:9.

(2) Analysis

Claim 1 of the '329 Patent recites (emphasis added):

> 1. An external-cache controller for a plurality of processor clusters within a computer having a main memory, each cluster being in inter-cluster communication over a common communication bus, and each cluster having at least one microprocessor having a processor bus and a *private processor cache, the contents of each such private cache being unknown externally to its associated microprocessor*, the controller comprising:
>
> a. an external tag memory for storing a status of the private processor cache associated with a first processor [cl]uster, the external tag memory non-hardwired to the private processor cache associated with the first processor cluster;
>
> b. a data-request monitor for monitoring a data request from a second processor cluster; and

> c. a tag controller associated with the first processor cluster for managing the external tag memory, the tag controller deriving the status of the private processor cache associated with the first processor cluster by tracking evicted cache lines and data entering and exiting the first processor cluster, the tag controller checking the tag memory to determine if a response to the request corresponds to data in the private processor cache associated with the first processor [cl]uster.

The specification discloses: "In preferred embodiments of the invention, external cache contents will be inclusive of L2/L1 contents." '329 Patent at 16:26–27.

During prosecution, the patentee stated:

> Claim 26 discloses an external-cache controller which maintains coherent cache interaction between multiple processor clusters. The contents of each cache serving a processor cluster is *hidden*, and not made public to any other cluster. Therefore, to maintain cache coherency, the controller must externally track data entering and exiting the processor cluster. In contrast, the processor caches in Fletcher [U.S. Pat. No. 4,484,267)] *are not hidden*. A copy of each processor's cache directory is provided to an external storage controller, allowing the controller to monitor each processor's cache so that coherency can be maintained.
>
> Claim 26 requires "that each cluster have at least one microprocessor having a private processor cache, the contents of each such *private* cache being *unknown externally* to its associated microprocessor." (emphasis added). . . . In contrast, Fletcher discloses a cache controller for a plurality of central processors (CP), where each CP's cache contents are *public* and made available to the controller, as described at col. 5, lines 18-19, referencing Flus[c]he et al. [(United States Patent No. 4,394,731).] The controller is provided copy directories (CD), which duplicate and have the same logical address as each processor's cache directory, as described by Flus[c]he at col. 5, lines 39-44. External tracking of data entering and exiting the cache is not taught or suggested by either reference. Fletcher's and Flus[c]he's teachings are contrary, and utilize *public* cache directories. Thus, claim 26 is not obvious in view of Fletcher.

Response of May, 26 2000 at 4–5, Dkt. No. 197, Ex. 4 (emphasis in original); *see* November 15, 2000 Amendment Under 37 C.F.R. § 1.116 at 5, Dkt. No 197, Ex. 20 ("The contents of the L2 cache are not known externally, and must be tracked. Only after such tracking is a reflected copy of the local P6 cache properly defined in the external cache." (emphasis omitted))

"To act as its own lexicographer, a patentee must *clearly* set forth a definition of the disputed claim term, and *clearly* express an intent to define the term." *GE Lighting Solutions, LLC v. AgiLight, Inc.*, 750 F.3d 1304, 1309 (Fed. Cir. 2014) (emphasis added) (citation and internal quotation marks omitted).

On balance, Plaintiff has failed to demonstrate that the plain meaning of "unknown" should be set aside and replaced with the vague concept of data not being "read directly." Of particular note, above-quoted Claim 1 recites that what is "deriv[ed]" by "tracking evicted cache lines and data entering and exiting the first processor cluster" is the "status" of cache lines, not necessarily the data contained therein. *Intergraph* noted this general concept as well. *Intergraph* at 14 ("the state or status of a cache line is . . . distinct from the data itself").

Particularly given that the patentee relied upon the "unknown" limitation to distinguish a prior art reference, as quoted above, the Court rejects Plaintiff's attempt to narrow the meaning of that limitation. Further, to whatever extent the patentee also distinguished the Fletcher reference on other grounds, Plaintiff's reliance on the "unknown" limitation is nonetheless binding. *See Andersen Corp. v. Fiber Composites, LLC*, 474 F.3d 1361, 1374 (Fed. Cir. 2007) ("An applicant's invocation of multiple grounds for distinguishing a prior art reference does not immunize each of them from being used to construe the claim language. Rather, as we have made clear, an applicant's argument that a prior art reference is distinguishable on a particular ground can serve as a disclaimer of claim scope even if the applicant distinguishes the reference on other grounds as well."); *see also Typhoon Touch Techs., Inc. v. Dell, Inc.*, 659 F.3d 1376, 1381 (Fed. Cir. 2011) ("The patentee is bound by representations made and actions that were taken in order to obtain the patent."); *Southwall Tech., Inc. v. Cardinal IG Co.*, 54 F.3d 1570,

1576 (Fed. Cir. 1995) ("Claims may not be construed one way in order to obtain their allowance and in a different way against accused infringers.").

The Court accordingly hereby construes **"private processor cache, the contents of each such private cache being unknown externally to its associated microprocessor"** to mean **"a cache within and dedicated to a processor whose contents are unknown to any device external to that processor, including at least the external cache controller and external tag memory."**

B. "external tag memory"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"storage that is situated apart from the main memory that includes information about the status of the private processor cache" Dkt. No. 192 at 7.	"storage that is situated apart from the processor and main memory that includes information about the status of the private processor cache and which has a larger number of entries than the number of cache lines in the private processor caches associated with the first processor cluster" Dkt. No. 197 at 5–6.

This disputed term appears in Claim 1 of the '329 Patent. Dkt. No. 197 at 5.

> *Intergraph* found as follows, in full, regarding this disputed term:
>
> The Court adopts Intergraph's construction and construes "external tag memory" to mean "storage that is situated apart from the main memory that includes information about the status of the private processor cache." HP argues that "external" must mean external to the processor, as it did in the first claim term construed. However, "external" is an adjective, and its meaning depends upon its context. Here, "external" modifies the noun "memory," thus, "external tag memory" must be situated apart from the main memory. The Court also rejects HP's construction of "tag memory" because it imports limitations from the specification. The Court finds Intergraph's construction comports with the plain meanings of "tag" and "memory." *See IEEE Standard Dictionary of Elec. & Elecs. Terms*, at 685, 1083 (6th ed. 1996). For these reasons, the Court adopts Intergraph's construction.

Intergraph at 10–11.

(1) The Parties' Positions

Plaintiff argues that "Defendants propose to import a limitation from a preferred embodiment in the specification." Dkt. No. 192 at 7.

Defendants respond that their proposed construction "clarifies the scope of the limitation by describing the *only* manner in which the claimed invention can possibly operate." Dkt. No. 197 at 6. Specifically, Defendants argue that "[t]he *only* way for the disclosed invention to 'forcibly track' the cache contents is to have an external tag memory that has a larger number of entries than the private processor cache." *Id.* (citing '329 Patent at 15:27–30)

Plaintiff replies, in full: "As explained in [Plaintiff's] opening brief, Defendants propose to import a limitation from a preferred embodiment in the specification. Defendants provide no evidence beyond their own speculation as to why the invention should be so limited." Dkt. No. 201 at 2.

At the January 29, 2015 hearing, Defendants were agreeable to construing this disputed term to require "storage that is situated apart from main memory" rather than "storage that is situated apart from the *processor* and main memory." Defendants maintained, however, that the number of tag entries must be greater than the number of cache lines because otherwise the tag memory would not be able to track the status of all of the cache lines.

(2) Analysis

Claim 1 of the '329 Patent is reproduced in the discussion of the "private processor cache . . ." term, above. Claim 1 recites, in relevant part (emphasis added): "an *external tag memory* for storing a status of the private processor cache associated with a first processor [cl]uster."

As quoted above, *Intergraph* did not address whether the external tag memory "has a larger number of entries than the number of cache lines in the private processor caches associated with the first processor cluster" as Defendants have proposed here.

The specification discloses:

> A problem that a preferred embodiment addresses is that no Intel specification for the P6 processor discloses how to make public the contents of the level 2 (L2) cache serving a processor cluster. No method is disclosed for determining 30 [*sic*] which internal cache line is being replaced when a new one is fetched. And, it may be inferred from Intel disclosure that the Bus Read and Invalidate (BRIL) and Bus Invalidate Line (BIL) functions may or may not modify the requested line. That is, Intel disclosure indicates that the requests intend to modify the lines, but does not say they will; therefore the data associated with a requested line may still be maintained within a processor cache, and may not become stale unless the requested line is actually modified. Consequently, the invention needs to forcibly track the L2 cache's contents. Towards this end, the XAP 204 tracks evicted P6 cache lines in producing a correct reflection of the P6 internal caches.

'329 Patent at 7:18–34.

> In order to minimize snoop cycles on the P6 bus, while also maintaining coherency in preferred embodiments, the following facilities are provided by the tag controller 905, as shown in FIG. 9. The first is that *a large external tag be able to accommodate up to 16 to 32 times the number of cache lines that may be stored by a single quad-CPU processor segment*.

'329 Patent at 15:24–30 (emphasis added).

Defendants explain:

> Because the claimed invention works with processors that do not disclose the contents of their private caches, the external cache controller must "derive" the status of those private caches. *Id.* at 7:19-21; claim 1. It does so by keeping track of what *might* be in the private cache via tracking data entering and exiting the processor and, then, by "snooping" to eliminate those items that are no longer in the cache. *Id.* at 15:42-44, 15:51-60, claim 1. Because snooping cannot occur after every moment that the cache contents may change, the external tag memory must have capacity to store statuses of items that *might* be in each cache line. *Id.* at 7:14-18, 15:45-60.

Dkt. No. 197 at 6.

On balance, Defendants' argument is not adequately supported by the specification. Instead, Defendants' proposed construction would improperly import limitations from particular disclosed embodiments. *See Electro Med.*, 34 F.3d at 1054; *see also Hill-Rom Services, Inc. v. Stryker Corp.*, 755 F.3d 1367, 1373 (Fed. Cir. 2014) (although the preferred embodiment "undisputedly uses a cable to convey data, and the patent does not disclose an alternative embodiment that uses a wireless datalink," the court found "no basis to narrow the plain and ordinary meaning of the term datalink to only one type of datalink—a cable"). Defendants' proposed construction is therefore hereby expressly rejected.

The Court accordingly hereby construes **"external tag memory"** to mean **"storage that is situated apart from the main memory and that includes information about the status of the private processor cache."**

C. "external tag memory non-hardwired to the private processor cache"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"the private processor cache does not duplicate its status in the external tag memory or the private processor cache does not use the same logical address as the external tag memory" Dkt. No. 192 at 8.	"external tag memory not permanently connected to the private processor cache such that, for example, the private processor cache does not duplicate its status in the external tag memory or the private processor cache does not use the same address as the external memory" Dkt. No. 197 at 7.

This disputed term appears in Claim 1 of the '329 Patent. Dkt. No. 197 at 7.

Intergraph found as follows, in full, regarding this disputed term:

> The Court adopts Intergraph's construction and construes "external tag memory non-hardwired to the private processor cache" to mean "the private processor cache does not duplicate its status in the external tag memory or the private processor cache does not use the same logical address as the external tag memory." The parties agree that Intergraph acted as its lexicographer in using the

> term "non-hardwired." The parties also seem to agree on the substantive meaning of "non-hardwired," but HP's proposed construction defines "hardwired" and states the external tag memory is not hardwired. The Court adopts Intergraph's construction as it is more straightforward.

Intergraph at 11.

(1) The Parties' Positions

Plaintiff argues that it "proposes adopting Judge Davis' construction, which reflects the inventors' definition of non-hardwired discussed in the prosecution history." Dkt. No. 192 at 8. Further, Plaintiff argues, "[i]t is entirely unclear how one would determine whether or not an external tag memory was permanently connected to a private processor cache under Defendants' proposal." *Id.* at 9.

Defendants respond that the parties here have presented a dispute that was not presented in *Intergraph*. Dkt. No. 197 at 7.

Defendants argue that the prosecution history contains no lexicography, and "[b]ecause the lexicography exception is inapplicable, the claim construction analysis must begin with the plain and ordinary meaning of the term." *Id.* at 8. Defendants submit that "'[h]ardwired' is a common term used to describe permanent electrical connections, as opposed to temporary ones." *Id.* Defendants also submit that "[a]ll processors of the preferred embodiments are detachably connected to a motherboard and other processors through sockets that allow users to easily change processors." *Id.* at 9.

Plaintiff replies that "the law does not require any particular magical definitional language," and Plaintiff argues that "[i]n responding to the office action rejection over [the] Fletcher [reference], the inventors explained how this added ['non-hardwired'] language distinguished the invention." Dkt. No. 201 at 3–4.

At the January 29, 2015 hearing, Plaintiff urged that tracking cache status has nothing to do with, for example, whether the tag memory is permanently soldered onto a motherboard or instead is removably installed in a socket. Plaintiff also argued that by omitting "logical" from their proposed construction, Defendants are attempting to eliminate the patentee's distinction between virtual addresses and physical addresses.

(2) Analysis

As a threshold matter, as quoted above, the parties in *Intergraph* "agree[d] on the substantive meaning of 'non-hardwired.'" *See Intergraph* at 11. That agreement is not binding upon Defendants here. *See, e.g., Fuji Photo Film Co. v. Int'l Trade Comm'n*, 386 F.3d 1095, 1101 (Fed. Cir. 2004) ("The infringement analysis in the initial determination was made pursuant to a stipulation with respect to the meaning of the claim terms by one of the respondents in that proceeding. Since the respondents who are affected by the claim construction in the present proceedings were not parties to that stipulation, they are not bound by it, nor does the administrative law judge's acceptance of the stipulation constitute a formal claim construction.").

Claim 1 of the '329 Patent is reproduced in the discussion of the "private processor cache . . ." term, above.

Neither "non-hardwired" nor "hardwired" appears in the written description. The specification discloses various connectors and configuration options:

> The following sections illustrate board layouts that may be used in high-end and mid[-]size servers. In some embodiments of the invention, the XBus and its associated ASICs may reside on the mother board, a no mid-plane design. Such a configuration may reduce packaging size for these products. Embodiments configured with mother boards may cost more, and force customers to purchase more complete products up front, as all ASICs will be on the mother board and the P6 CPUs will become the main option. The invention's mid-plane based products will allow customers to purchase a minimum chassis and then configure their product over time.

> FIG. 5 shows the processor segments in a preferred embodiment of the four XBus system configuration. This embodiment has SMT straddle mounted connectors for XABus 140 and XDBus 142. Immediately above each connector will be a data path 202 or address path 204 ASIC. Preferably the connection length between the ASIC and the connectors is minimized, and preferred embodiments will utilize stub connections. Further, each side of the ASIC must be used for a specific interface: XBus (XDP 202/XAP 204), tag 210 or cache 212, P6 bus, and XAP 204 or XDP 204 [*sic*, 202] interconnect.
>
> The tag RAM 210 is fast access, preferably 5-8 n sec, static RAMs mounted on a SIMM or its equivalent (e.g. DINIM). The SIMM will house as many as twelve RAMs for tracking 512 K cache lines. A preferred minimum tag configuration will contain six RAMs on one side of the SIMM for tracking 256 K cache lines. The tag RAMs are not optional and one of the two possible options must be resident. In order to enable a high speed design, each SIMM will service only one XAP. A more complex embodiment may be configured with 1,024K cache lines per XBus, through use of higher density RAM technology.
>
> In a preferred embodiment, the L3 Cache 212 will be designed around SDRAM technology to allow fast cache burst access. The SDRAMs will be mounted on a SIMM and provide a 72 bit data path. Due to available memory size, the SIMMs may contain as much as 1,024K cache lines. TAGs presently track only up to 512K cache lines. As TAGs are a critical component of the system, every effort is made to minimize cache line access time.
>
> Each processor segment will preferably house up to four CPUs 500 and associated power conditioning circuitry 502. Although preferred embodiments of the invention will utilize CPUs with a low profile heat sink that extends beyond the CPUs, other embodiments may use heat pipes, flanges, fanned enclosures, or other cooling devices.

‘329 Patent at 17:15–64; *see id.* at Fig. 5.

As for the prosecution history, Plaintiff submits that “[t]he term was added in a Response to Office Action dated August 8, 2001 in order to distinguish a prior art reference,” Dkt. No. 197 at 8–9:

> Claim 26, as amended, requires “an external tag memory for storing a status of the private processor cache associated with a first processor cluster, the external tag memory *non-hardwired* to the private processor cache associated with the first processor cluster.” (emphasis added) In contrast, Fletcher discloses a storage controller that includes copy directories at col. 6, lines 19 and 20, referencing

> Flus[c]he et al. The copy directories duplicate and have the same logical address as each processor's cache directory, as described by Flus[c]he at col. 5, lines 39-44. *Hence*, in Fletcher, the copy directories are hardwired to processor cache, unlike claim 26, which requires that the external tag memory be non-hardwired to the private processor cache.

August, 8 2001 Response to Office Action at 7, Dkt. No. 197, Ex. 8 (emphasis added); *see id.* at 9 (amending application claim 26, which issued as Claim 1, so as to recite a "non-hardwired" limitation).

On balance, the patentee's use of "[h]ence" in the above-quoted passage, in context, signals a lexicography. *See GE Lighting Solutions*, 750 F.3d at 1309. Alternatively, even if this does not rise to the level of a lexicography, this prosecution history nonetheless demonstrates how the patentee used the term "non-hardwired." *See, e.g.*, *Typhoon Touch*, 659 F.3d at 1381; *Southwall Tech.*, 54 F.3d at 1576. That is, the patentee characterized Fletcher as "hardwired" and asserted that the claimed invention is not hardwired.

As to extrinsic evidence, Defendants have submitted dictionary and treatise definitions of "hardwired" as meaning: "Of, pertaining to, or effected by means of logic circuitry that is permanently connected within a computer or calculator," *Webster's II New College Dictionary* 505 (1995), Dkt. No. 197, Ex. 9; "Some personal computers have cache memory chips hardwired onto the motherboard," Ashok Arora, *Foundations of Computer Science* 42 (2006), Dkt. No. 197, Ex. 12; and "changes and modifications that are made to computer equipment that actually permanently change the machine as opposed to plug[-]in changes with accessory boards" (John V. Lombardi, *Computer Literacy* 100 (1983), Dkt. No. 197, Ex. 13).[1]

[1] Defendants have also cited two "online" dictionaries that set forth definitions similar to those reproduced here. *See* Dkt. No. 197, Exs. 10 & 11.

Because these definitions relate to the physical attachment of computer boards and chips, for example, these definitions are not germane to the context in which the term "non-hardwired" is used in the claim and in the prosecution history, as set forth above. *See Phillips*, 415 F.3d at 1321 ("[H]eavy reliance on the dictionary divorced from the intrinsic evidence risks transforming the meaning of the claim term to the artisan into the meaning of the term in the abstract, out of its particular context, which is the specification.").

This extrinsic evidence thus is not of sufficient weight to overcome the intrinsic evidence set forth above. Finally, to enhance clarity consistent with the above-quoted prosecution history, the Court modifies the Intergraph construction by replacing "or" with "and."

The Court accordingly hereby construes **"external tag memory non-hardwired to the private processor cache"** to mean **"the private processor cache does not duplicate its status in the external tag memory and the private processor cache does not use the same logical address as the external tag memory."**

D. "deriving the status of the private processor cache associated with the first processor cluster by tracking evicted cache lines and data entering and exiting the first processor cluster"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"to get or obtain the status of the private processor cache associated with the first processor cluster by tracking: (1) all modified lines[2] within the private processor cache of that same cluster, (2) all data that enters that same cluster, and (3) all data that exits that same cluster, and using at least that information to determine the status of each private processor cache in only that same cluster" Dkt. No. 192 at 10.	"to get or obtain the status of the private processor cache associated with the first processor cluster by tracking: (1) all modified lines within the private processor cache of that same cluster, (2) all data that enters that same cluster, and (3) all data that exits that same cluster, and using at least that information to determine the status of each private processor cache in only that same cluster. 'Evicted' does not mean what was in the cache but is no longer there (i.e., a line that has been ejected), but rather something that is now in the cache that is different from what was previously there (i.e., a modified line)" Dkt. No. 197 at 10.

This disputed term appears in Claim 1 of the '329 Patent. Dkt. No. 197 at 10.

In *Intergraph*, the Court construed this disputed term to mean:

> To get or obtain the status of the private processor cache associated with the first processor cluster by tracking: (1) all modified lines within the private processor cache of that same cluster, (2) all data that enters that same cluster, and (3) all data that exits that same cluster, and using at least that information to determine the status of each private processor cache in only that same cluster.

Intergraph at 12 & 15.

(1) The Parties' Positions

Plaintiff argues: "The second sentence proposed by Defendants also appears in Judge Davis's Order but was not part of his construction and should not be part of this construction for

[2] Plaintiff previously proposed "evicted lines" instead of "modified lines." Dkt. No. 173, Ex. A at 6.

multiple reasons." Dkt. No. 192 at 10. Plaintiff explains that "evicted" does not appear in the agreed-upon portion of the construction and, moreover, "Judge Davis' conclusion as to the meaning of 'evicted' is contrary to its ordinary meaning of 'ejected,' which in context includes all cache lines that have exited the private processor cache of a particular private processor." *Id.* at 11. Plaintiff urges that "the specification clearly discloses tracking of unmodified lines." *Id.* (citing '329 Patent at 7:18–34).)

Defendants respond that even though Plaintiff "repeats the *exact* same argument that Judge Davis considered and rejected" in *Intergraph*, Plaintiff fails to explain how *Intergraph* erred. Dkt. No. 197 at 10–11. Defendants submit: "Defendants' proposed construction incorporates Judge Davis' explicit definition of 'evicted' to clarify the scope of the claim limitation at issue. *See* [*Intergraph*] at 13. This definition of 'evicted' was the heart of Judge Davis' construction for the disputed claim term. Its inclusion would avoid jury confusion by expressly defining an important word within the disputed claim term." *Id.* at 11–12.

Plaintiff replies: "Judge Davis expressly did not include the definition of evicted in his claim construction. (Intergraph Order at Appendix A, element [c][ii]). Because the construction language removes the term 'evicted' from the language to be used by the jury in interpreting the claims, providing some additional instruction about 'evicted' is superfluous at best and confusing at worst." Dkt. No. 201 at 5–6.

(2) Analysis

Claim 1 of the '329 Patent is reproduced in the discussion of the "private processor cache . . ." term, above. Claim 1 recites, in relevant part (emphasis added):

> . . .
> a tag controller associated with the first processor cluster for managing the external tag memory, the tag controller *deriving the status of the private*

> *processor cache associated with the first processor cluster by tracking evicted cache lines and data entering and exiting the first processor cluster*, the tag controller checking the tag memory to determine if a response to the request corresponds to data in the private processor cache associated with the first processor [cl]uster . . . .

The specification discloses:

> A problem that a preferred embodiment addresses is that no Intel specification for the P6 processor discloses how to make public the contents of the level 2 (L2) cache serving a processor cluster. No method is disclosed for determining 30 [*sic*] which internal cache line is being replaced when a new one is fetched. And, it may be inferred from Intel disclosure that the Bus Read and Invalidate (BRIL) and Bus Invalidate Line (BIL) functions may or may not modify the requested line. That is, Intel disclosure indicates that the requests intend to modify the lines, but does not say they will; therefore the data associated with a requested line may still be maintained within a processor cache, and may not become stale unless the requested line is actually modified. Consequently, the invention needs to forcibly track the L2 cache's contents. Towards this end, the XAP 204 tracks *evicted P6 cache lines* in producing a correct reflection of the P6 internal caches. If external TAGs indicate that a line is in the modified state, then XAP 204 must not include that line as part of the cleansing process.

'329 Patent at 7:18–36 (emphasis added).

In *Intergraph*, "both parties focus[ed] on the specification to determine whether 'evicted cache lines' refers to modified or unmodified cache lines. Intergraph argue[d] the term means 'tracking the replacement of unmodified cache lines.' HP contend[ed] that the proper construction is 'observing write cycles initiated by the processor to determine which dirty cache lines are being evicted.'" *Intergraph* at 13. *Intergraph* found:

> This passage indicates to one skilled in the art that an "evicted" cache line is one that has been "modified." The term "evicted" is clearly a poor choice of words. The word usually carries the connotation of something that has been ejected. But, here, the patentee is using the term in a way other than its ordinary meaning as the parties agree. Rather than referring to what was in the cache but is no longer there (*i.e.*, a line that has been ejected), the term refers to something that is now in the cache that is different from what was previously there (*i.e.*, a modified line).

Id. Intergraph also noted that "the state or status of a cache line is therefore distinct from the data itself." *Id.* at 14.

Because the parties here agree that an "evicted line" is a "modified line," and because neither party's proposed construction includes the term "evicted," Defendants' proposed explanatory language should not be included in the Court's construction. The parties are otherwise in agreement as to the proper construction.

The Court accordingly hereby construes **"deriving the status of the private processor cache associated with the first processor cluster by tracking evicted cache lines and data entering and exiting the first processor cluster"** to mean **"to get or obtain the status of the private processor cache associated with the first processor cluster by tracking: (1) all modified lines within the private processor cache of that same cluster, (2) all data that enters that same cluster, and (3) all data that exits that same cluster; and using at least that information to determine the status of each private processor cache in only that same cluster."**

E. "tag controller"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"a controller that manages the external tag memory" Dkt. No. 192 at 12.	"a controller that manages the external tag memory for processors that do not disclose which internal cache line is being replaced when a new one is fetched" Dkt. No. 197 at 12.

This disputed term appears in Claim 1 of the '329 Patent. Dkt. No. 197 at 12. In *Intergraph*, the Court did not address this term apart from addressing "tag controller associated with the first processor cluster." *See Intergraph* at 12.

(1) The Parties' Positions

Plaintiff argues: "Although the specification describes the cache management function as a solution to the problem of the example processor (Intel P6) not disclosing a method for determin[ing] which internal cache line is being replaced when a new one is fetched, the P6 is merely the example being used. Nothing in the claims or the prosecution history indicates an intent by the inventor to limit the invention to application only to a P6 processor or a processor with identical characteristics." Dkt. No. 192 at 12.

Defendants respond that "Defendants' construction does not conflict with Judge Davis's construction of 'tag controller associated with the first processor cluster' because the dispute in the *Intergraph* case centered on the meaning of the word 'associated' and not on the scope of 'tag controller.'" Dkt. No. 197 at 13 (citing *Intergraph* at 12); *see Intergraph* at 4–9 (discussing "associated with").

Plaintiff replies that "Defendants cite no evidence to support their attempt to limit the claim language to a preferred embodiment." Dkt. No. 201 at 6.

(2) Analysis

Claim 1 of the '329 Patent is reproduced in the discussion of the "private processor cache . . ." term, above.

The specification discloses that "the invention" can track cache lines:

> A problem that a preferred embodiment addresses is that no Intel specification for the P6 processor discloses how to make public the contents of the level 2 (L2) cache serving a processor cluster. No method is disclosed for determining 30 [*sic*] which internal cache line is being replaced when a new one is fetched. And, it may be inferred from Intel disclosure that the Bus Read and Invalidate (BRIL) and Bus Invalidate Line (BIL) functions may or may not modify the requested line. That is, Intel disclosure indicates that the requests intend to modify the lines, but does not say they will; therefore the data associated with a requested line may still be maintained within a processor cache, and may not become stale unless the

> requested line is actually modified. Consequently, *the invention* needs to forcibly track the L2 cache's contents. Towards this end, the XAP 204 tracks evicted P6 cache lines in producing a correct reflection of the P6 internal caches.

'329 Patent at 7:18–32 (emphasis added).

In some circumstances, discussion of "the invention" can be limiting. *See Regents of the Univ. of Minn. v. AGA Med. Corp.*, 717 F.3d 929, 936 (Fed. Cir. 2013) ("When a patent . . . describes the features of the 'present invention' as a whole, this description limits the scope of the invention." (quoting *Verizon Servs. Corp. v. Vonage Holdings Corp.*, 503 F.3d 1295, 1308 (Fed. Cir. 2007)).

Here, however, the above-quoted reference to "the invention" appears in the context of the description of a preferred embodiment and does not purport to describe the claimed invention as a whole. In other words, this disclosure refers to applying "the invention" in a particular situation. On balance, Defendants' proposed construction would improperly import a limitation from particular disclosed embodiments and is therefore rejected. *See Electro Med.*, 34 F.3d at 1054.

The Court accordingly hereby construes **"tag controller"** to mean **"a controller that manages the external tag memory."**

F. "when the processor bus is idling"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"when the processor bus is not being used by the processor" Dkt. No. 173, Ex. A at 7	"when the processor bus is not being used by the processor or other connected device" Dkt. No. 197 at 13.

This term appears in Claim 2 of the '329 Patent. Dkt. No. 197 at 13. *Intergraph* did not address this term.

Plaintiff's opening brief does not address this term. *See* Dkt. No. 192.

Defendants argue that "[t]o be idle, the bus cannot be used by any connected device." Dkt. No. 197 at 13 (citing '329 Patent 7:14–18). Defendants further argue that "[Plaintiff's] construction not only contradicts the plain meaning of the term, but would rewrite the limitation from the bus being idle to the processor being idle." Dkt. No. 197 at 14.

Plaintiff's reply brief does not address this term. *See* Dkt. No. 201.

In their January 13, 2015 Joint Claim Construction Chart, the parties submitted that Plaintiff has agreed to Defendants' proposed construction. *See* Dkt. No. 203, Ex. A at 5. At the January 29, 2015 hearing, the parties confirmed that they have reached agreement in this regard.

The Court accordingly hereby construes **"when the processor bus is idling"** to mean **"when the processor bus is not being used by the processor or other connected device."**

G. "to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
Plain and ordinary meaning Dkt. No. 192 at 13.	"asking the private processor cache if particular data is present to determine whether a cache line has been modified if the tag controller indicates a copy is loaded in one of the first processor cluster's private processor caches" Dkt. No. 197 at 14.

This term appears in Claim 2 of the '329 Patent. Dkt. No. 197 at 14. *Intergraph* did not address this term.

Plaintiff argued that "[i]t is unclear how Defendants' proposal is different from the plain and ordinary meaning of the actual words of the claims." Dkt. No. 192 at 13.

Defendants responded that "[i]n the context of the claim language, it is clear that the 'snooping' function requires making an inquiry into the contents of the private processor cache." Dkt. No. 197 at 14. Defendants concluded: "Defendants' construction clarifies claim scope and puts the technical language of the claim limitation in a more jury-friendly form. Because [Plaintiff] does not disagree with the substance of Defendants' clarification of the claim language, Defendants' construction should be adopted." *Id.*

Plaintiff replied that "[t]he meaning of Defendants' proposal is still unclear even with the explanation in their brief, although it appears that yet again Defendants seek to define this term by importing examples from the specification." Dkt. No. 201 at 7.

At the January 29, 2015 hearing, the parties submitted that they have reached agreement that this term should be construed to have its plain meaning.

The Court accordingly hereby construes **"to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster"** to have its **plain meaning.**

H. "a snooper for snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
Plain and ordinary meaning (not subject to 35 U.S.C. § 112 ¶ 6) Dkt. No. 192 at 13	Governed by 35 U.S.C. § 112 ¶6. Function: "snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" Structure: "There is insufficient disclosure of structure to perform this function. Alternatively, the closest corresponding structure is XAP 204 in Figure 9 as referenced at 4:51-65." Dkt. No. 197 at 15.

This disputed term appears in Claim 2 of the '329 Patent. Dkt. No. 197 at 15. *Intergraph* did not address this term.

(1) The Parties' Positions

Plaintiff emphasizes that this term does not use the word "means," and Plaintiff argues that "[o]ne of ordinary skill would read the '329 Patent disclosure and understand what is meant by a snooper, which is described repeatedly throughout the written description." Dkt. No. 192 at 14.

Defendants argue that "[t]he phrase 'snooper for snooping' is written as a purely functional term." Dkt. No. 197 at 15. Further, Defendants argue, Plaintiff "points to passages in

the specification that describe[] the snooping function, but fails to identify any structure corresponding to that function." *Id.* Defendants conclude that Claim 2 of the '329 Patent is invalid as indefinite. *Id.* at 16.

Plaintiff replies, in full: "As explained in [Plaintiff's] opening brief, Defendants have failed to prove that Section 112, paragraph 6, should apply to this term." Dkt. No. 201 at 8.

At the January 29, 2015 hearing, Plaintiff argued that by failing to present any evidence, Defendants have failed to meet their burden to overcome the presumption against means-plus-function treatment of a term that does not use the word "means." Defendants responded that they need not present any evidence, such as expert opinion, because the patentee simply rewrote the claimed function as a noun. Defendants urged that, under such circumstances, means-plus-function treatment is appropriate as a matter of law.

(2) Analysis

The Supreme Court of the United States "read[s] [35 U.S.C.] § 112, ¶ 2 to require that a patent's claims, viewed in light of the specification and prosecution history, inform those skilled in the art about the scope of the invention with reasonable certainty." *Nautilus, Inc. v. Biosig Instruments, Inc.*, 134 S. Ct. 2120, 2129 (2014). "A determination of claim indefiniteness is a legal conclusion that is drawn from the court's performance of its duty as the construer of patent claims." *Datamize, LLC v. Plumtree Software, Inc.*, 417 F.3d 1342, 1347 (Fed. Cir. 2005) (citations and internal quotation marks omitted), *abrogated on other grounds by Nautilus,* 134 S. Ct. 2120.

> It is well settled that [a] claim limitation that actually uses the word "means" invokes a rebuttable presumption that § 112, ¶ 6 applies. By contrast, a claim term that does not use "means" will trigger the rebuttable presumption that § 112, ¶ 6 does not apply. The term "means" is central to the analysis.

Apex Inc. v. Raritan Computer, Inc., 325 F.3d 1364, 1371–72 (Fed. Cir. 2003) (citations and internal quotation marks omitted); *see Inventio AG v. Thyssenkrupp Elevator Ams. Corp.*, 649 F.3d 1350, 1356 (Fed. Cir. 2011) ("[T]he presumption flowing from the absence of the term 'means' is a strong one that is not readily overcome."); *see also Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1297–98 (Fed. Cir. 2014) ("The strong presumption created by not including means in a claim limitation provides clarity and predictability for the public and the patentee alike.").

The presumption against applying 35 U.S.C. § 112, ¶ 6 can be overcome. Thus, "[m]eans-plus-function claiming applies only to purely functional limitations that do not provide the structure that performs the recited function." *Welker Bearing Co., v. PHD, Inc.*, 550 F.3d 1090, 1095–96 (Fed. Cir. 2008) (quoting *Phillips*, 415 F.3d at 1311).

The threshold issue, then, is whether the constituent term "snooper" connotes structure or, instead, fails to connote structure such that the presumption against means-plus-function treatment is rebutted. On its face, "snooper" appears to be merely a noun form of the claimed "snooping" function.

Plaintiff submits that "the patent examiner rejected the original claims concerning the snooper ([application] claim 27) but never on [35 U.S.C.] Section 112 grounds with respect to the snooper itself, as opposed to its functionality." Dkt. No. 192 at 14. The Court nonetheless has an independent duty to evaluate indefiniteness.

Claim 2 of the '329 Patent recites:

> 2. An external-cache controller according to claim 1, further comprising:
> d. a snooper for snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster, the snooper snooping when the processor bus is idling.

Plaintiff has cited various passages in the specification as support for its argument that "[o]ne of ordinary skill would read the '329 Patent disclosure and understand what is meant by a snooper." Dkt. No. 192 at 14 (citing '329 Patent at 6:19–48, 7:44–68, 8:18–36, 9:6–48, 12:54–67, 13:25–31, 13:54–64, 15:15–23, 15:34–38, 15:45–60 & 19:25–20:05).

The passages cited by Plaintiff refer to "snooping" and to, for example, a "snoop capability," '329 Patent at 12:62, but none of these passages sets forth a "snooper" structure or demonstrates that a "snooper" is a well-known structure in the relevant art. Indeed, one of Plaintiff's citations is to the disputed claim language itself. '329 Patent at 19:25-20:05.

On balance, Plaintiff has failed to identify any significantly probative evidence that "snooper" has a structural meaning. The Court therefore finds that the presumption against means-plus-function treatment has been rebutted. *See, e.g.*, *Mas-Hamilton Grp. v. LaGard, Inc.*, 156 F.3d 1206, 1214 (Fed. Cir. 1998) (finding that "lever moving element" "cannot be construed so broadly to cover every conceivable way or means to perform the function of moving a lever, and there is no structure recited in the limitation that would save it from application of section 112, ¶ 6"); *Lighting World, Inc. v. Birchwood Lighting, Inc.*, 382 F.3d 1354, 1360 (Fed. Cir. 2004) ("What is important is whether the term is one that is understood to describe structure, as opposed to a term that is simply a nonce word or a verbal construct that is not recognized as the name of structure and is simply a substitute for the term 'means for.'").

The parties do not appear to have any dispute regarding the claimed function, which is "snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster."

The remaining issue is whether the specification discloses sufficient corresponding structure for performing this claimed function. "[S]tructure disclosed in the specification is 'corresponding' structure only if the specification or prosecution history clearly links or associates that structure to the function recited in the claim." *Med. Instrumentation & Diagnostics Corp. v. Elekta AB*, 344 F.3d 1205, 1210 (Fed. Cir. 2003) (quoting *B. Braun Med. Inc. v. Abbott Labs.*, 124 F.3d 1419, 1424 (Fed. Cir. 1997)). A "black box" illustration that represents function without any mention of structure does not provide sufficient corresponding structure to satisfy 35 U.S.C. § 112, ¶ 6. *See ePlus, Inc. v. Lawson Software, Inc.*, 700 F.3d 509, 518 (Fed. Cir. 2012). Further, "[t]he indefiniteness inquiry is concerned with whether the bounds of the invention are sufficiently demarcated, not with whether one of ordinary skill in the art may find a way to practice the invention." *See id.* at 519 (rejecting argument that the structure of a means-plus-function term was known in the art).

Here, the "snooping" functionality is disclosed as part of a general-purpose computer system. "[A] means-plus-function claim element for which the only disclosed structure is a general purpose computer is invalid if the specification fails to disclose an algorithm for performing the claimed function." *Net MoneyIN Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1367 (Fed. Cir. 2008); *see WMS Gaming, Inc. v. Int'l Game Tech.*, 184 F.3d 1339, 1349 (Fed. Cir. 1999) ("In a means-plus-function claim in which the disclosed structure is a computer, or microprocessor, programmed to carry out an algorithm, the disclosed structure is not the general purpose computer, but rather the special purpose computer programmed to perform the disclosed algorithm.").

If an algorithm is required, that algorithm may be disclosed in any understandable form. *See Typhoon Touch*, 659 F.3d at 1386 ("Indeed, the mathematical algorithm of the programmer

is not included in the specification. However, as precedent establishes, it suffices if the specification recites in prose the algorithm to be implemented by the programmer."); *see also Finisar Corp. v. DirecTV Group, Inc.*, 523 F.3d 1323, 1340 (Fed. Cir. 2008) (noting that "a patentee [may] express th[e] algorithm in any understandable terms including as a mathematical formula, in prose, or as a flow chart, or in any other manner that provides sufficient structure" (citation omitted)); *TecSec, Inc. v. Int'l Bus. Machs.*, 731 F.3d 1336, 1348 (Fed. Cir. 2013) (quoting *Finisar*).

Nonetheless, the purported algorithm cannot "merely provide[] functional language" and must provide a "step-by-step procedure" for accomplishing the claimed function. *Ergo Licensing, LLC v. CareFusion 303, Inc.*, 673 F.3d 1361, 1365 (Fed. Cir. 2012). Further, "[i]t is well settled that simply disclosing software, however, without providing some detail about the means to accomplish the function, is not enough.'" *Function Media, L.L.C. v. Google, Inc.*, 708 F.3d 1310, 1318 (Fed. Cir. 2013) (citation and internal quotations and alterations omitted). Finally, when citing sections of the specification, a patentee should demonstrate "how these sections explain to one of ordinary skill in the art the manner in which the claimed functions are implemented." *Personalized Media Commc'n, LLC v. Motorola, Inc.*, No. 2:08-CV-70, 2011 WL 4591898, at *38 (E.D. Tex. Sept. 30, 2011); *see Function Media*, 708 F.3d at 1318 ("These citations all explain that the software automatically transmits, but they contain no explanation of how the PGP software performs the transmission function.").

Plaintiff has not proposed any corresponding structure, let alone explained how the cited passages of the specification set forth any algorithm for performing the "snooping" function.

The Court, finding no corresponding structure, concludes that the term **"a snooper for snooping the private processor cache of the at least one microprocessor associated with the**

first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" is **indefinite**.

I. "processor cluster"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"cluster having at least one of each of a microprocessor, a processor bus, and a private processor cache" Dkt. No. 192 at 15.	"a set of one or more processors in a multiprocessor system that share a common CPU bus" Dkt. No. 197 at 16.

This disputed term appears in Claim 1 of the '329 Patent. Dkt. No. 197 at 16. In *Intergraph*, the Court did not separately address this term.

(1) The Parties' Positions

Plaintiff submits that its proposal "uses the same terminology from the Preamble and is proposed only to clarify the clause 'each [processor] cluster having at least one microprocessor having a processor bus and a private processor cache.'" Dkt. No. 192 at 15. Plaintiff argues that Defendants' proposal improperly introduces "multiprocessor system" and "CPU bus" without support. *Id.* In particular, Plaintiff argues that the term "CPU bus" "is only used in two places in the specification" and "is also only used in connection with preferred embodiments and should not replace the claim language." *Id.* at 15–16.

Defendants respond that because "[t]he claim already includes nearly the exact language with which [Plaintiff] is attempting to define 'processor cluster[,]'" "[Plaintiff's] construction renders 'processor cluster' entirely superfluous." Dkt. No. 197 at 16. Defendants submit that the specification explains that "the maximum number of processors allowed in each cluster is constrained by the capacity of their shared bus. [Plaintiff's] construction ignores this important

consideration and would allow any arbitrary group of processors to be considered a 'processor cluster' even when they do not share a common bus." *Id.* at 17.

Plaintiff replies by reiterating its opening arguments and by emphasizing that "[t]here is absolutely no evidence, and none is cited by Defendants, to indicate an intent by the inventor to limit the invention to a circumstance in which multiple processors share a single CPU bus." Dkt. No. 201 at 9.

(2) Analysis

Claim 1 of the '329 Patent recites, in relevant part (emphasis added):

> 1. An external-cache controller for a plurality of *processor clusters* within a computer having a main memory, *each cluster* being in *inter-cluster* communication over a common communication bus, and *each cluster* having at least one microprocessor having a processor bus and a private processor cache, the contents of each such private cache being unknown externally to its associated microprocessor, the controller comprising: . . . .

The "Summary" section of the specification states:

> *The present invention* provides a high-availability parallel processing server that is a multi-processor computer with a segmented memory architecture. The processors are grouped into *processor clusters, with each cluster consisting of up to four processors in a preferred embodiment*, and there may be up to 5 clusters of processors in a preferred embodiment. *Each cluster of processors has dedicated memory buses for communicating with each of the memory segments.* The invention is designed to be able to maintain coherent interaction between all processors and memory segments within a preferred embodiment. A preferred embodiment uses Intel Pentium-Pro processors (hereinafter P6). The invention may be modified to utilize other processors, such as those produced by AMD or CYRIX. (Registered trademarks referenced herein belong to their respective owners.)
>
> *The present invention* comprises a plurality of processor segments (a *cluster of one or more CPUs*), memory segments (separate regions of memory), and memory communication buses (pathways to communicate with the memory segment). Each processor segment has a dedicated communication bus for interacting with each memory segment, allowing different processors parallel access to different memory segments while working in parallel.

> *The processors, in a preferred embodiment, may further include an internal cache* and flags associated with the cache to indicate when the data within the cache may be out of date, or if the data within the cache is data shared by other processors within the invention. Through use of setting the internal cache flag to a desired state, the contents of the internal cache of a processor may be effectively monitored from a vantage point external to the processor. This would allow for maintaining multi-processor cache coherency without requiring all processors to observe all other processors[’] memory traffic.

‘329 Patent at 1:17–50 (emphasis added).

> Although the invention may be implemented with various central processing units (CPUs), in a preferred embodiment, the Intel P6 (or P6 MMX) processor will be used. According to the Intel specification, up to four P6 processors may be clustered on an In-Order split-transaction *CPU bus* to create a four way symmetrical multi-processor (SMP) platform. As used in this specification and claims that follow, such clustering of CPUs is referenced as a “processor segment.”

Id. at 3:15–22; *see id.* at 4:18 (“CPU bus”); *see also id.* at 6:10 (“processor bus”).

On balance, Defendants have the better reading of the plain language of the claim, particularly in light of the above-quoted disclosures of a “CPU bus” or “processor bus.” Plaintiff’s proposal, by contrast, would potentially allow the term “processor cluster” to encompass any collection of at least one microprocessor, processor bus, and private processor cache without requiring any connection or relationship with one another.

At the January 29, 2015 hearing, Plaintiff expressed concern that referring to a “common” bus in the construction would imply that a cluster must include at least two processors. The phrase “a set of one or more processors,” however, adequately makes clear that a “cluster” could be comprised of only one processor.

The Court accordingly hereby construes **“processor cluster”** to mean **“a set of one or more processors that share a common processor bus in a multiprocessor system.”**

V. CONSTRUCTION OF DISPUTED TERMS IN THE '750 PATENT

A. "[first/second] instruction pipeline"

Plaintiff's Proposed Construction	Defendants' Proposed Construction
"[first/second] of at least two structures, each consisting of a sequence of stages" Dkt. No. 192 at 16 (square brackets in original).	"[first/second] of at least two structures, each consisting of a sequence of stages that execute instructions received from a single instruction issuing unit" Dkt. No. 197 at 18 (square brackets in original).

These disputed terms appear in Claim 8 of the '750 Patent. Dkt. No. 197 at 18.

(1) The Parties' Positions

Plaintiff argues that Defendants' proposal "is inconsistent with the specification, which expressly discloses that the instruction pipelines can have different units from which they receive instructions." Dkt. No. 192 at 16.

Defendants respond that "[b]y excluding the requirement of executing instructions, [Plaintiff's] construction would allow any 'structure,' including a single instruction pipeline, to be arbitrarily divided into multiple segments to meet the 'first' and 'second' pipeline limitation, even though the divided segments cannot each execute instructions." Dkt. No. 197 at 18–19. Defendants also argue: "The only disclosed embodiment includes multiple instruction pipelines receiving instructions from the same instruction issuing unit. *See* ['750 Patent] at Fig. 1; 4:59-63. The '750 Patent does not disclose, and does not claim, two unrelated and uncoordinated instruction pipelines that operate independently." Dkt. No. 197 at 20.

Plaintiff replies, in full: "Defendants' proposed addition of the requirement that the two pipelines execute instructions received from a single instruction issuing unit is an improper

importation of a limitation from a preferred embodiment, as discussed in [Plaintiff's] opening brief." Dkt. No. 201 at 9.

At the January 29, 2015 hearing, Plaintiff argued that Defendants' proposal of stages "that execute instructions" is too narrow because instruction pipeline stages perform more operations than merely executing. Plaintiff submitted it would be agreeable to a construction setting forth that pipelines include stages "that process instructions."

As to Defendants' proposal that instructions must be "received from a single instruction issuing unit," Defendants urged that, without such a limitation, Plaintiff could accuse different stages of the same instruction pipeline as being the "first instruction pipeline" and "second instruction pipeline." Plaintiff responded that it has no intention of doing so.

(2) Analysis

Claim 8 of the '750 Patent recites (emphasis added):

> 8. A method for translating virtual addresses in a computing system having at least *a first and a second instruction pipeline* and a direct address translation unit for translating virtual addresses into real addresses, the direct address translation unit including a master translation memory for storing translation data, the direct address translation unit for translating a virtual address into a corresponding real address, comprising the steps of:
>
> storing a first subset of translation data from the master translation memory into a first translation buffer associated with the *first instruction pipeline*;
>
> translating a first virtual address received from the *first instruction pipeline* into a corresponding first real address, wherein the first virtual address translating step comprises the steps of:
>
> accessing the first translation buffer;
>
> indicating whether translation data for the first virtual address is stored in the first translation buffer;
>
> activating the direct address translation unit to translate the first virtual address when the translation data for the first virtual address is not stored in the first translation buffer; and
>
> storing the translation data for the first virtual address from the master translation memory into the first translation buffer;

storing a second subset of translation data from the master translation memory into a second translation buffer associated with the *second instruction pipeline*; and

translating a second virtual address received from the *second instruction pipeline* into a corresponding second real address, wherein the second virtual address translating step comprises the steps of:

accessing the second translation buffer;

indicating whether translation data for the second virtual address is stored in the second translation buffer;

activating the direct address translation unit to translate the second virtual address when the translation data for the second virtual address is not stored in the second translation buffer; and

storing the translation data for the second virtual address from the master translation memory into the second translation buffer.

The specification discloses:

> FIG. 5 is a block diagram of a particular embodiment of an apparatus 200 according to the present invention for translating virtual addresses in a computing system such as computing system 10 shown in FIG. 1. Apparatus 200 includes, for example, a load *instruction pipeline* 210A, a load *instruction pipeline* 210B, and a store *instruction pipeline* 210C. These pipelines may be three of the pipelines 18A-H shown in FIG. 1.

'750 Patent at 4:56–64 (emphasis added). As Defendants have pointed out, all eight instruction pipelines in Figure 1 are apparently illustrated as receiving instructions from the same instruction issuing unit 14. *See id.* at Fig. 1.

On balance, however, requiring that instructions are received from a single instruction issuing unit would improperly limit the claims to a particular feature of an exemplary embodiment. *See MBO Labs. Inc. v. Becton, Dickinson & Co.*, 474 F.3d 1323, 1333 (Fed. Cir. 2007) ("[P]atent coverage is not necessarily limited to inventions that look like the ones in the figures. To hold otherwise would be to import limitations [i]nto the claim[s] from the specification, which is fraught with danger."); *see also Electro Med.*, 34 F.3d at 1054; *Trebro Mfg., Inc. v. Firefly Equipment, LLC*, 748 F.3d 1159, 1166 (Fed. Cir. 2014); *Flo Healthcare*

Solutions, LLC v. Kappos, 697 F.3d 1367, 1375 (Fed. Cir. 2012) ("It is true that all of the embodiments discussed in the patent have a length-adjustable vertical beam, but it is not proper to import from the patent's written description limitations that are not found in the claims themselves.").

In response to a petition for *Inter Partes* Review ("IPR") at the PTO, Plaintiff explained the concept of a "pipeline" as follows:

> The CPUs employ pipelining, in which instructions are partitioned into various execution phases to allow simultaneous operations on different phases of different instructions by different logical units.

Patent Owner's Preliminary Response to Petition for *Inter Partes* Review at 5, Dkt. No. 197, Ex. 26.

As to extrinsic evidence, Defendants have submitted a declaration by the named inventor, Howard Sachs, declaring that: "the ['750] patent assumes that the pipelines (multiple cores or processors) are all running the same thread. If each pipeline is executing a different thread then this type of system is not comprehended by U.S. Patent No. 5,463,750." Sachs Decl. ¶ 7, Dkt. No. 197, Ex. 31. The opinions of the named inventor, however, do not significantly affect the Court's analysis here. *See Howmedica Osteonics Corp. v. Wright Med. Tech., Inc.*, 540 F.3d 1337, 1346–47 (Fed. Cir. 2008) (noting that inventor testimony is "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"); *but see Phillips*, 415 F.3d at 1317 ("Although we have emphasized the importance of intrinsic evidence in claim construction, we have also authorized district courts to rely on extrinsic evidence, which consists of all evidence external to the patent and prosecution history, including expert and *inventor testimony*, dictionaries, and learned treatises." (emphasis added) (citations and internal quotation marks

omitted)); *Voice Techs. Grp., Inc. v. VMC Sys., Inc.*, 164 F.3d 605, 615 (Fed. Cir. 1999) ("An inventor is a competent witness to explain the invention and what was intended to be conveyed by the specification and covered by the claims. The testimony of the inventor may also provide background information, including explanation of the problems that existed at the time the invention was made and the inventor's solution to these problems.").

In light of all of the foregoing, Defendants' proposed construction is hereby expressly rejected as lacking adequate support in the intrinsic evidence. To provide proper context, however, the Court adopts Plaintiff's alternative proposal to specify that the pipeline stages "process instructions."

The Court accordingly hereby construes **"[first/second] instruction pipeline"** to mean **"[first/second] of at least two structures, each consisting of a sequence of stages that process instructions."**

B. "master translation memory"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
Plain and ordinary meaning Dkt. No. 192 at 17.	"page tables in main memory" Dkt. No. 197 at 21.

This disputed term appears in Claim 8 of the '750 Patent. Dkt. No. 197 at 21.

(1) The Parties' Positions

Plaintiff argues that "Defendants' proposal reads a limitation from the preferred embodiment into claim language that does not need to be construed." Dkt. No. 192 at 17.

Defendants respond that "'[m]aster translation memory' is a coined term that does not appear in the specification, and it has a specific technical meaning in the context of claim 8." Dkt. No. 197 at 21. Defendants explain: "[C]laim 8 requires the 'translation data' to be sourced

from the 'master translation memory' when that 'translation data' is not present in the 'translation buffer.' *See* ['750 Patent] at 3:33-47. That source—the 'master translation memory'—must be the master source of all 'translation data' for the 'translation buffer.' The master source of the translation data is the page tables in main memory." Dkt. No. 197 at 22. Defendants conclude that "translation buffers (TLBs) cannot be the '*master* translation memory' because they store only a subset of the translation data. . . . To give meaning to the word 'master,' the 'master translation memory' must mean those page tables stored in main memory." *Id.* at 23.

Plaintiff replies that whereas "Defendants focus on distinguishing between two storage locations of memory described in the specification: translation buffers (TLB) and address translation tables in main memory," "those storage locations are merely examples of the claimed invention." Dkt. No. 201 at 10.

At the January 29, 2015 hearing, Defendants argued that the authoritative, "master" translation memory must be in main memory rather than in a cache because a cache, by definition, holds only copies of data.

(2) Analysis

Claim 8 of the '750 Patent is reproduced in the discussion of the "[first/second] instruction pipeline . . ." term, above.

The specification discloses "a dynamic translation unit (DTU) 162 for accessing page tables in main memory 34." '750 Patent at 2:67–3:01.

> When the miss signal is generated, DTU 162 accesses the page tables in main memory 34 to determine whether in fact the requested data is currently stored in main memory 34. If not, then DTU 162 instructs data transfer unit 42 through a communication path 194 to fetch the page containing the requested data from

> mass storage device 30. In any event, TLB 158 is updated through a communication path 196, and instruction issuing resumes.

Id. at 3:40–47; *see id.* at 5:42–44 ("[T]he address translation tables in main memory are accessed to obtain address translation information . . . .").

Also, Figure 4 of the '750 Patent illustrates TLB 158 connected to DTU 162 shown as connected to "Main Memory." Figure 5 likewise shows three TLBs 222A, 222B, and 222C connected to a DTU 162 (through Update Control 240) that is connected to "Main Memory."

On balance, however, requiring "page tables in main memory" would improperly limit the claims to a particular feature of an exemplary embodiment. *See MBO Labs.*, 474 F.3d at 1333 ("[P]atent coverage is not necessarily limited to inventions that look like the ones in the figures. To hold otherwise would be to import limitations [i]nto the claim[s] from the specification, which is fraught with danger."); *see also Electro Med.*, 34 F.3d at 1054; *Trebro Mfg.*, 748 F.3d at 1166; *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 906 (Fed. Cir. 2004) ("[T]his court has expressly rejected the contention that if a patent describes only a single embodiment, the claims of the patent must be construed as being limited to that embodiment."); *GE Lighting Solutions*, 750 F.3d at 1309.

The Court therefore hereby expressly rejects Defendants' proposed construction. No further construction is necessary. *See U.S. Surgical Corp. v. Ethicon, Inc.*, 103 F.3d 1554, 1568 (Fed. Cir. 1997) ("Claim construction is a matter of resolution of disputed meanings and technical scope, to clarify and when necessary to explain what the patentee covered by the claims, for use in the determination of infringement. It is not an obligatory exercise in redundancy."); *see also O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1362 (Fed. Cir. 2008) ("[D]istrict courts are not (and should not be) required to construe every

limitation present in a patent's asserted claims."); *Finjan, Inc. v. Secure Computing Corp.*, 626 F.3d 1197, 1207 (Fed. Cir. 2010) ("Unlike *O2 Micro*, where the court failed to resolve the parties' quarrel, the district court rejected Defendants' construction.").

The Court accordingly hereby construes **"master translation memory"** to have its **plain meaning**.

C. "direct address translation unit"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"a common unit shared by all instruction execution pipelines that translates a virtual memory address and also comprising the master translation memory"[3] Dkt. No. 192 at 17. Dkt. No. 203, Ex. A at 6.	"a common unit shared by all instruction execution pipelines that translates a virtual memory address using data in the page table and also comprising [the master translation memory]" Dkt. No. 197 at 23. Dkt. No. 203, Ex. A at 6 (brackets in original).

This disputed term appears in Claim 8 of the '750 Patent. Dkt. No. 197 at 23.

(1) The Parties' Positions

Plaintiff argues that "page tables are merely an example used in the specification." Dkt. No. 192 at 18.

Defendants respond that "[t]he 'unit' that is common and shared cannot be any processing unit. Rather, it must be the unit that performs 'direct address translation.' The only unit that performs such 'direct address translation' is the circuitry (or software) that uses data in the page tables to perform address translations—converting virtual addresses into physical addresses." Dkt. No. 197 at 23–24. Defendants conclude that their proposed construction

[3] Plaintiff previously proposed: "Single direct address translation unit." Dkt. No. 173, Ex. A at 14.

"resolves the ambiguity of [Plaintiff's] construction by clarifying that the DATU term requires the unit that directly translates addresses using the page tables in main memory to be shared by the pipelines." *Id.* at 27.

Plaintiff replies that "[p]age tables are merely an example used in the specification and therefore should not be included in the construction." Dkt. No. 201 at 10–11.

At the January 29, 2015 hearing, Defendants reiterated that there must be one distinct data address translation unit, for all pipelines, as opposed to a group of translation units.

(2) Analysis

Claim 8 of the '750 Patent is reproduced in the discussion of the "[first/second] instruction pipeline . . ." term, above. Claim 8 recites, in relevant part (emphasis added):

> . . .
> wherein the first virtual address translating step comprises the steps of:
> accessing the first translation buffer;
> indicating whether translation data for the first virtual address is stored in the first translation buffer;
> activating the *direct address translation unit* to translate the first virtual address when the translation data for the first virtual address is not stored in the first translation buffer; and
> storing the translation data for the first virtual address from the master translation memory into the first translation buffer;
> . . . .

The Abstract of the '750 Patent refers to "an update control circuit for activating the direct address translation circuit when the translation data for the virtual address is not stored in the TLB." Also, Figure 4 of the '750 Patent illustrates TLB 158 connected to DTU 162 shown as connected to "Main Memory." Figure 5 likewise shows three TLBs 222A, 222B, and 222C connected to a DTU 162 (through Update Control 240) that is connected to "Main Memory."

Requiring that the direct address translation unit uses pages from main memory, however, would improperly limit the claims to a particular feature of an exemplary embodiment. *See MBO Labs.*, 474 F.3d at 1333 ("[P]atent coverage is not necessarily limited to inventions that look like the ones in the figures. To hold otherwise would be to import limitations [i]nto the claim[s] from the specification, which is fraught with danger."); *see also Electro Med.*, 34 F.3d at 1054; *Trebro Mfg.*, 748 F.3d at 1166.

As for the prosecution history, at one point all of the then-pending claims were rejected because:

> Applicant failed to teach how to make first and second direct address translating means in address translators coupled to respective pipeline [*sic*, pipelines], and it would require a person of ordinary skill in the art undue experimentation to develop such means.

Office Action of July 28, 1994 at 2–3, Dkt. No. 197, Ex. 23. In response, the patentee relied on United States Patent No. 4,933,835 Patent ("the '835 Patent"), which the '750 Patent incorporates by reference, as disclosing "address translators." Amendment of September 29, 1994 at 1–2, Dkt. No. 197, Ex. 24. The '835 Patent, in turn, uses "direct address translation" to refer to translating an address using page tables after a TLB miss. '835 Patent at 22:6–10, Dkt. No. 197, Ex. 29 ("In the event of a TLB miss, a TLB miss signal 372 is coupled to the direct address translation unit 280. The DAT 280 provides *page table* access as illustrated at 374, and provides replacement of TLB lines as illustrated at 375." (emphasis added)); *see id.* at Fig. 9.

Finally, in response to an IPR petition, Plaintiff stated:

> [T]he VAX 8800 contains two processor cores. * * * The two cores do not *share* any 'direct address translation unit,' hence, the VAX 8800 cannot meet the requirements of claim 1.

(Patent Owner's Preliminary Response to Petition for *Inter Partes* Review at 6–7, Dkt. No. 197, Ex. 26 (emphasis added). Defendants urge that because "the DATU [(direct address translation unit)] term must refer to a common unit that performs direct address translations for both pipelines, not a collection of distinct units," "the DATU term requires the unit that directly translates addresses using the page tables in main memory to be shared by the pipelines." Dkt. No. 197 at 27 (citing September 5, 2014 Decision at 16–18, Dkt. No. 197, Ex. 33).

On balance, Defendants have failed to identify any definitive statement that would warrant limiting the disputed term to require the use of page tables. *See Omega Eng'g v. Raytek Corp.*, 334 F.3d 1314, 1324 (Fed. Cir. 2003) ("As a basic principle of claim interpretation, prosecution disclaimer promotes the public notice function of the intrinsic evidence and protects the public's reliance on *definitive* statements made during prosecution." (emphasis added)); *see also id.* at 1325–26 ("[F]or prosecution disclaimer to attach, our precedent requires that the alleged disavowing actions or statements made during prosecution be both *clear and unmistakable*." (emphasis added)); *Plantronics, Inc. v. Aliph, Inc.*, 724 F.3d 1343, 1349–53 (Fed. Cir. 2013) (reversing claim construction regarding certain "stabilizer support" terms, noting that "[t]he asserted claims are drafted broadly, without bounds to any particular structure," and neither the specification nor the prosecution history limited the structure).

The Court accordingly hereby construes **"direct address translation unit"** to mean **"a common unit shared by all instruction execution pipelines that translates a virtual memory address and that also comprises the master translation memory."**

D. "storing the translation data for the [first/second/third] virtual address from the master translation memory into the [first/second/third] translation buffer"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
Plain and ordinary meaning Dkt. No. 192 at 18.	"storing the translation data for the [first/second/third] virtual address from the master translation memory into the [first/second/third] translation buffer even when the translation data for the [first/second/third] virtual address is stored in the [first/second/third] translation buffer" Dkt. No. 197 at 27–28 (brackets in original)

These disputed terms appear in Claims 8 through 12 of the '750 Patent. Dkt. No. 197 at 27.

(1) The Parties' Positions

Plaintiff argues that the additional language proposed by Defendants "appears to repeat the claim language and/or may be attempting to read a limitation from the preferred embodiment into claim language that does not need to be construed." Dkt. No. 192 at 18.

Defendants respond that "while the '750 Patent specification discloses only updating the translation buffer when the comparator issues a miss signal—*i.e.*, only when translation data for the virtual address is not already stored in the translation buffer—Defendants' construction merely seeks to clarify that the claim language does not have any such required condition precedent." Dkt. No. 197 at 28. Further, Defendants explain, "[t]he 'storing' step comes next and is structurally and grammatically set forth as another step in the method independent of whether the 'indicating' step indicated a hit or miss, and independent of whether the conditional 'activating' step occurred." *Id.*

Plaintiff replies, in full: "[Plaintiff] believes that the language of this claim does not require construction. Defendants' brief confirms that they are improperly seeking to add an

additional requirement to the claim language, effectively importing a temporal negative limitation." Dkt. No. 201 at 11.

At the January 29, 2015 hearing, Plaintiff argued that the context set forth in Claim 8 of the '750 Patent makes clear that the "storing" steps occur only if the needed translation data is not found in the translation buffers (which the parties have referred to as a "TLB miss").

Defendants responded that although the embodiments disclosed in the specification require a "miss" before such storing occurs, no such condition appears in the claims.

(2) Analysis

The specification discloses updating TLBs when there is a TLB miss, that is, when requested information has not been found in a TLB:

> Comparators 238A-C compare the virtual address tags to the higher order bits of the respective virtual addresses and provide hit/miss signals on communication paths 238A-C to an update control circuit 240.
>
> Update control circuit 240 controls the operation of DTU 162 through a communication path 244 and updates TLB's 222A-C through respective update circuits 241-243 and communication paths 248A-C whenever there is a miss signal generated on one or more of communication paths 238A-C. That is, update control circuit 240 activates DTU 162 whenever a miss signal is received over communication path 238A and stores the desired translation information in TLB 222A through communication path 248A; update control circuit 240 activates DTU 162 whenever a miss signal is received over communication path 238B and stores the desired translation information in TLB 222B through communication path 248B; and update control circuit 240 activates DTU 162 whenever a miss signal is received over communication path 238C and stores the desired translation information in TLB 222C through communication path 248C.

'750 Patent at 5:06–26; *see id.* at 3:40–48 & 3:56–58 ("new virtual-to-real address translation information is stored in TLB 158 whenever a miss signal is generated by comparator 170"). Claim 8 of the '750 Patent is reproduced in the discussion of the "[first/second] instruction

pipeline . . .” term, above. Claims 9 through 12 depend from Claim 8 and recite (emphasis added):

> 9. The method according to claim 8 further comprising the step of:
> *storing the translation data for the first virtual address from the master translation memory into the second translation buffer* whenever translation data for the first virtual address from the master translating memory is stored into the first translation buffer.
>
> 10. The method according to claim 9 further comprising the step of:
> *storing the translation data for the second virtual address from the master translation memory into the first translation buffer* whenever translation data for the second virtual address from the master translation memory is stored into the second translation buffer.
>
> 11. The method according to claim 10 further comprising the steps of:
> storing a third subset of translation data from the master translation memory into a third translation buffer associated with the third instruction pipeline; and
> translating a third virtual address received from the third instruction pipeline into a corresponding third real address, where[]in the third virtual address translating step comprises the steps of:
> accessing the third translation buffer;
> indicating whether translation data for the third virtual address is stored in the third translation buffer;
> activating the direct address translation unit to translate the third virtual address when the translation data for the third virtual address is not stored in the third translation buffer; and
> *storing the translation data for the third virtual address from the master translation memory into the third translation buffer*.
>
> 12. The method according to claim 11,
> wherein the step of storing the translation data for the third virtual address comprises the step of storing translation data for only the third virtual address in the third translation buffer.

On balance, Defendants’ proposed construction would tend to confuse rather than clarify the scope of the claims and is hereby expressly rejected. *See Power-One, Inc. v. Artesyn Techs., Inc.*, 599 F.3d 1343, 1348 (Fed. Cir. 2010) (“The terms, as construed by the court, must ensure

that the jury fully understands the court's claim construction rulings and what the patentee covered by the claims." (citation and internal quotation marks omitted)).

Nonetheless, the Court hereby expressly rejects Plaintiff's argument that the "storing" steps are conditioned on the occurrence of a TLB miss. No such condition appears in the claims. *See K-2 Corp. v. Salomon S.A.*, 191 F.3d 1356, 1364 (Fed. Cir. 1999) ("Courts do not rewrite claims; instead, we give effect to the terms chosen by the patentee."); *see also Chef Am., Inc. v. Lamb-Weston, Inc.*, 358 F.3d 1371, 1374 (Fed. Cir. 2004) ("[C]ourts may not redraft claims, whether to make them operable or to sustain their validity.").

No further construction is required. *See U.S. Surgical*, 103 F.3d at 1568; *see also O2 Micro*, 521 F.3d at 1362; *Finjan*, 626 F.3d at 1207.

The Court accordingly hereby construes **"storing the translation data for the [first/second/third] virtual address from the master translation memory into the [first/second/third] translation buffer"** to have its **plain meaning**.

E. "whenever translation data for the [first/second] virtual address from the master translation memory is stored into the [first/second] translation buffer"

Plaintiff's Proposed Construction	**Defendants' Proposed Construction**
"whenever" should be construed to mean "at the same time when" Dkt. No. 173, Ex. A at 15.	"whenever" should be construed to mean "every time that" Dkt. No. 197 at 29.

These disputed terms appear in Claims 9 through 10 of the '750 Patent. Dkt. No. 197 at 29.

Plaintiff's opening brief does not address this term. *See* Dkt. No. 192.

Defendants argue that they "seek a construction using the plain and ordinary definition of 'whenever'—which is consistent with the intrinsic record and taken verbatim from the

specification." Dkt. No. 197 at 29. Defendants argue that Plaintiff's proposal improperly removes the "causal link" between the two actions. *Id.*

Plaintiff's reply brief does not address this term. *See* Dkt. No. 201.

In their January 13, 2015 Joint Claim Construction Chart, the parties submitted that Plaintiff has agreed to Defendants' proposed construction. *See* Dkt. No. 203, Ex. A at 10. At the January 29, 2015 hearing, the parties confirmed that they have reached agreement in this regard.

The Court accordingly hereby construes **"whenever"** to mean **"every time that."**

VI. CONCLUSION

The Court adopts the constructions set forth in this opinion for the disputed terms of the patents-in-suit. The parties are ORDERED that they may not refer, directly or indirectly, to each other's claim construction positions in the presence of the jury. Likewise, the parties are ORDERED to refrain from mentioning any portion of this opinion, other than the actual definitions adopted by the Court, in the presence of the jury. Any reference to claim construction proceedings is limited to informing the jury of the definitions adopted by the Court.

Having found that the term "a snooper for snooping the private processor cache of the at least one microprocessor associated with the first processor cluster to identify if the data has been modified if the tag controller indicates that data is held within the private processor cache associated with the first processor cluster" is indefinite, as discussed above, the Court hereby finds that Claim 2 of the '329 Patent is invalid.

Within thirty (30) days of the issuance of this Memorandum Opinion and Order, the parties are hereby ORDERED, in good faith, to mediate this case with the mediator agreed upon by the parties. As a part of such mediation, each party shall appear by counsel and by at least one corporate officer possessing sufficient authority and control to unilaterally make binding

decisions for the corporation adequate to address any good faith offer or counteroffer of settlement that might arise during such mediation. Failure to do so shall be deemed by the Court as a failure to mediate in good faith and may subject that party to such sanctions as the Court deems appropriate.

So ORDERED and SIGNED this 10th day of February, 2015.

RODNEY GILSTRAP
UNITED STATES DISTRICT JUDGE