



## Table of Contents

<b>I. BACKGROUND</b> .....	<b>3</b>
<b>II. LEGAL PRINCIPLES</b> .....	<b>7</b>
<b>III. CONSTRUCTION OF AGREED TERMS</b> .....	<b>11</b>
<b>IV. CONSTRUCTION OF DISPUTED TERMS</b> .....	<b>12</b>
A. “bus master” (‘906 Patent, Claim 9).....	12
B. “secondary memory” (‘906 Patent, Claim 9).....	12
C. “first cache memory” (‘906 Patent, Claim 9).....	16
D. “at a constant rate” (‘906 Patent, Claim 9) .....	20
E. “means for sequentially transferring . . .” (‘906 Patent, Claim 26).....	26
F. “means for . . . determining” (‘906 Patent, Claim 26).....	32
G. “initiating one and only one snoop access of said cache memory, said snoop accesses each specifying the respective N+1’th L-byte line” (‘291 Patent, Claims 73 & 88) .....	39
H. ‘291 Patent, Claims 88 and 89 .....	43
<b>V. CONCLUSION</b> .....	<b>49</b>

## **I. BACKGROUND**

Plaintiff brings suit alleging infringement of United States Patents No. 5,710,906 (“the ‘906 Patent”) and 6,405,291 (“the ‘291 Patent”) (collectively, the “patents-in-suit”). The ‘906 Patent issued on January 20, 1998, and bears a filing date of July 7, 1995. The ‘291 Patent issued on June 11, 2002, and is a continuation of a continuation of a continuation of a divisional of the ‘906 Patent, so the ‘291 Patent has the same specification as the ‘906 Patent. For convenience, references to the specification shall be to only the ‘906 Patent unless otherwise indicated.

The patents-in-suit are sometimes referred to by Plaintiff as the “Pre-Snoop Patents.” In general, the patents-in-suit relate to cache memory, which is a special, temporary memory that can be used, for example, with a central processing unit (“CPU”). Reading data from the cache is faster than reading data from main memory. The cache can thereby improve the performance of the CPU and other devices. Some devices can access main memory without passing the data through the CPU, however, using a feature known as Direct Memory Access (“DMA”). The terms “inquire” and “snoop” both refer to checking for consistency between data in the cache and corresponding data in main memory. If the data in the cache has been modified, then the corresponding data in the main memory should be updated before that data in main memory is accessed, such as by a DMA device. Devices communicate with each other, with memory, and with the processor over one or more buses, such as the Peripheral Component Interconnect (“PCI”) bus.

The Abstracts of the ‘906 Patent and the ‘291 Patent are the same and state:

When a PCI-bus controller receives a request from a PCI-bus master to transfer data with an address in secondary memory, the controller performs an initial

inquire cycle and withholds TRDY#<sup>1</sup> to the PCI-bus master until any write-back cycle completes. The controller then allows the burst access to take place between secondary memory and the PCI-bus master, and simultaneously and predictively, performs an inquire cycle of the L1 [(level one)] cache for the next cache line. In this manner, if the PCI burst continues past the cache line boundary, the new inquire cycle will already have taken place, or will already be in progress, thereby allowing the burst to proceed with, at most, a short delay. Predictive snoop cycles are not performed if the first transfer of a PCI-bus master access would be the last transfer before a cache line boundary is reached.

The patents-in-suit have been construed in each of the following three cases: *OPTi Inc. v. NVIDIA Corp.*, No. 2:04-CV-377, Dkt. No. 96 (E.D. Tex. Apr. 24, 2006) (“*NVIDIA Markman*”), *OPTi Inc. v. Advanced Micro Devices, Inc.*, No. 2:06-CV-477, Dkt. No. 81 (E.D. Tex. July 16, 2008) (“*AMD Markman*”), and *OPTi Inc. v. Apple, Inc.*, No. 2:07-CV-21, Dkt. No. 62 (E.D. Tex. Dec. 4, 2008) (“*Apple Markman*”). These prior claim construction orders are attached to Plaintiff’s opening brief as Exhibits 3, 4, and 5, respectively.

The *NVIDIA Markman* set forth additional background regarding the technology of the patents-in-suit:

The OPTi patents relate to “core logic” chipsets, the processors that direct traffic between the central processor, memory, input/output devices, graphics cards, video cards, and various other devices that are contained within, or connected to, a computer.

In the earliest days of computer processing, there were no core logic chipsets. The central processor communicated directly with peripheral devices that made up the computer. As computers got more complicated, chipsets were introduced as a way of coordinating the burgeoning array of functionality and relieving central processors of that administrative burden. This freed more CPU resources for the fundamental mission of computing.

Broadly speaking, a typical chipset operates as an input/output (I/O) hub for the CPU, memory, peripherals, etc.

---

<sup>1</sup> “TRDY#” is a “Target Ready” signal. (Dkt. No. 135, Ex. 10, 12/10/2001 Ghosh Decl., at p. 3.) The “#” indicates that it is an “active low” signal, meaning that the signal is considered active or “asserted” when the signal voltage is low. (*Id.*, at p. 2 & 3; ‘906 Patent at 13:44-55.) “IRDY#” is a similar “Initiator Ready” signal that is controlled by the bus master.

\* \* \*

The links between the various devices comprising the computer, known as interfaces, consist of conductors on which the devices transmit signals to one another, communicating address, command, and data information. The most common type of interface is known as a bus. The buses and interfaces allow the various computer devices to exchange data and to operate in coordination with one another.

*NVIDIA Markman* at 5-6.

The Pre-Snoop patents addressed a[n] . . . issue that arose with the introduction of the PCI (“Peripheral Component Interconnect”) bus and the subsequent development of the Pentium and Pentium-compatible processors. One of the advantages of the PCI was its ability to transfer data from one device to another by a particular method called “burst” transfers. The Pre-Snoop patents disclose a technique for optimizing such burst transfers with Pentium processors.

Data is stored, created or used at a lot of places in a computer. Each such location is known as an address. For example, a memory storage device containing data to be read (the “target”) cannot know that it is being asked to transfer data or what data to transfer unless and until the requesting device (the “master”) puts an address onto the bus notifying the target that it is the object of a request and notifying the target what data is being requested. In a “burst” transfer, this information is all that the target needs to figure out which data to transfer, as the target dispatches data until the target is told to stop by the master or elects to stop the transaction itself.

A complication arises in this scenario because much of the memory can be stored in two places: main memory or cache memory. Cache memory is memory that stores copies of information expected to be used by the CPU at addresses that correspond to addresses for that information in secondary memory. This memory typically operates at particularly high speed and is typically positioned adjacent to the CPU. Access to the cache is thus generally quicker than access to the main memory. This speeds up the CPU’s ability to access and process the data that it needs.

As the CPU processes data, it saves that data to the cache for continued convenient access. The problem is that the CPU may well change the data that it is processing. If that modified data is stored only in the cache, it will not be identical to the data stored on, for example, the disk drive from which it was initially read. Thus, if some other device – a CD drive, for example – accesses the main memory to read data, it may get data that is no longer current.

In Intel’s X86 line of CPU’s, the system solved this problem by using a “write-through” cache. Basically, as data was modified by the CPU, it was written to

both the main and cache memories, thereby assuring constant “cache consistency.” In the Pentium processors, the cache was a “write-back” cache. This meant that the CPU did not take the time to write every modification through to main memory. Instead, modified data was stored in the cache with a flag to indicate its modified state. Thus, to read data from memory in a Pentium system, it was necessary to adopt some mechanism to check for the presence of this flag, to assure that the data in main memory was valid, and that the cache did not contain a version of data that had been modified by the CPU.

The mechanism adopted was a “snoop.” For example, a bus master seeking to initiate a transaction would first initiate an “inquire” or “snoop” cycle to the CPU to find out whether the data being sought had been stored in cache memory in a modified state and to “write back” the most current version of the data to the main memory if a modification had been made. The PCI standard required one “line” of cache memory to be snooped at a time, and then permitted transfer of that line if the snoop showed it to be either absent from the cache or in the cache but unmodified. The PCI protocol required that a transfer stop at the end of each line transferred, snoop the next line, transfer the line just snooped, and then stop again to snoop the succeeding line. Because burst transfers could encompass any amount of data, including data stored in multiple lines of memory, this practice resulted in a non-uniform transfer in which the bus spent more time sitting idle than it did carrying data.

The sole exception to this rule was that multiple lines of memory could be read without stopping when the data to be transferred was not cacheable. In this instance, the snoop operation could be ignored because there was no risk of stale data being accessed. Thus, the entire burst of data could be sequentially pre-fetched and read to the master without interruption. The Pre-Snoop patents embody the idea that cacheable memory could be transferred nearly as rapidly as non-cacheable memory if the snoop of a line were conducted while the preceding line was transferring. In that way, as in the case of non-cacheable memory, the system would know that the line was not stale and there would be no need to stop the transfer at the end of the first line to snoop the second line because that snoop would already be completed. The “snoop ahead” process could be repeated as long as the burst transfer was underway so that all of the data comprising the burst could be sent without interruption and at a constant rate.

*Id.* at 8-10.

The asserted claims, according to Defendants, are Claims 9 and 26 of the '906 Patent and Claims 73, 74, 88, and 89 of the '291 Patent. (Dkt. No. 139, at 1 n.1.)

## 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." 415 F.3d at 1312 (emphasis added) (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 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.

Indefiniteness is a “legal conclusion that is drawn from the court’s performance of its duty as the construer of patent claims.” *Exxon Research & Eng’g Co. v. U.S.*, 265 F.3d 1371, 1376 (Fed. Cir. 2001) (citation omitted). A finding of indefiniteness must overcome the statutory presumption of validity. *See* 35 U.S.C. § 282. That is, the “standard [for finding indefiniteness] is met where an accused infringer shows by clear and convincing evidence that a skilled artisan could not discern the boundaries of the claim based on the claim language, the specification, and the prosecution history, as well as her knowledge of the relevant art area.”

*Halliburton Energy Servs., Inc. v. M-I LLC*, 514 F.3d 1244, 1249-50 (Fed. Cir. 2008).

In determining whether that standard is met, i.e., whether the claims at issue are sufficiently precise to permit a potential competitor to determine whether or not he is infringing, we have not held that a claim is indefinite merely because it poses a difficult issue of claim construction. We engage in claim construction every day, and cases frequently present close questions of claim construction on which expert witnesses, trial courts, and even the judges of this court may

disagree. Under a broad concept of indefiniteness, all but the clearest claim construction issues could be regarded as giving rise to invalidating indefiniteness in the claims at issue. But we have not adopted that approach to the law of indefiniteness. We have not insisted that claims be plain on their face in order to avoid condemnation for indefiniteness; rather, what we have asked is that the claims be amenable to construction, however difficult that task may be. If a claim is insolubly ambiguous, and no narrowing construction can properly be adopted, we have held the claim indefinite. If the meaning of the claim is discernible, even though the task may be formidable and the conclusion may be one over which reasonable persons will disagree, we have held the claim sufficiently clear to avoid invalidity on indefiniteness grounds. . . . By finding claims indefinite only if reasonable efforts at claim construction prove futile, we accord respect to the statutory presumption of patent validity . . . and we protect the inventive contribution of patentees, even when the drafting of their patents has been less than ideal.

*Exxon*, 265 F.3d at 1375 (citations and internal quotation marks omitted).

### III. CONSTRUCTION OF AGREED TERMS

The Court hereby adopts the following agreed-upon constructions:

<b>Term</b>	<b>Patent / Claims</b>	<b>Agreed Construction</b>
“sequentially transferring”	‘906 Pat., Cl. 9, 26; ‘291 Pat., Cl. 73	“transferring in the address sequence of their arrangement in memory”
“determining whether an N+1’th 1-byte line of said secondary memory is cached in a modified state in said first cache memory”	‘906 Pat., Cl. 9, 26	“snooping the next line of secondary memory beyond the line currently being transferred to determine if it is cached in a modified state”
“said step of sequentially transferring”	‘906 Pat., Cl. 9; ‘291 Pat., Cl. 74	Same as “sequentially transferring,” above.
“PCI-bus burst transaction”	‘291 Pat., Cl. 73, 88	“a burst transaction in accordance with any version of the PCI Local Bus Specification”
“said step of transferring”	‘291 Pat., Cl. 73	Same as “sequentially transferring,” above.
“PCI-bus burst read transaction”	‘291 Pat., Cl. 88	“a burst transaction in accordance with any version of the PCI Local Bus Specification”
“sequentially transfers”	‘291 Pat., Cl. 88	Same as “sequentially transferring,” above.

(Dkt. No. 130, 10/15/2012 Joint Claim Construction and Prehearing Statement, at Ex. A.)

#### IV. CONSTRUCTION OF DISPUTED TERMS

##### A. “bus master” (‘906 Patent, Claim 9)

Plaintiff’s Proposed Construction	Defendants’ Proposed Construction
“An I/O-bus device that initiates a data transfer on an I/O bus, including a PCI bus device that initiates a data transfer on a PCI bus”	“A PCI I/O device that initiates a data transfer on a PCI bus”

(Dkt. No. 135, at 12.)

“After further analysis, and without waiver to its defenses of enablement and written description, [Defendants] do[] not contest [Plaintiff’s] construction of ‘bus master.’” (Dkt. No. 139, at 7 (footnote omitted); *see* Dkt. No. 142, 12/7/2012 P.R. 4-5(d) Joint Claim Construction Chart, Ex. A, at 1.)

The Court therefore hereby adopts Plaintiff’s unopposed proposal that **“bus master”** means **“an I/O-bus device that initiates a data transfer on an I/O bus, including a PCI bus device that initiates a data transfer on a PCI bus.”**

##### B. “secondary memory” (‘906 Patent, Claim 9)

Plaintiff’s Proposed Construction	Defendants’ Proposed Construction
“memory located logically behind the first level cache memory, i.e., DRAM memory and, if present, L2 and L3 cache memory”	“memory located logically behind the first cache memory”  Previously: “main system memory, e.g., DRAM memory”

(Dkt. No. 135, at 14; Dkt. No. 139, at 7.)

##### (1) The Parties’ Positions

Plaintiff argues that its proposed construction was the construction reached in *AMD* and *Apple* and that “[t]he Pre-Snoop Patents are clear throughout that ‘secondary memory’ is any kind of memory logically behind the first level cache (L1) memory and specifically recognizes

that ‘secondary memory’ can include not only the system’s DRAM, but also L2, L3 and other caches.” (Dkt. No. 135, at 14.) Plaintiff further explains that from the perspective of the L1 cache, any additional level of cache is treated as part of the same structure as the main memory. (*Id.*, at 15.)

Defendants submit that “[a]fter further analysis, [Defendants] agree[] to the first portion of [Plaintiff’s] construction of ‘secondary memory’ as ‘memory located logically behind the first cache memory’ with the slight revision, the elimination of the word ‘level.’” (Dkt. No. 139, at 8.) Defendants argue that Plaintiff’s proposed use of “level” is “vague in the context of the claims” and that Plaintiff’s proposed “i.e.” phrase is improper because it is redundant and “because other types of memory should not be excluded.” (*Id.*, at 8.)

Plaintiff replies that “[Defendants’] position is that any memory system which is logically behind some undefined ‘first cache’ will do.” (Dkt. No. 143, at 2.) Plaintiff submits that Defendants intend to argue that their VT82C505 prior art chip, referred to as the “505 Chip,” practiced the claims-at-issue when it snooped the L2 instead of the L1 cache. (*Id.*, at 2-3.) Plaintiff reiterates that Figure 1 and the specification explain that an L2 cache, if present, is part of the secondary memory. (*Id.*, at 3.)

## (2) Analysis

Claim 9 of the ‘906 Patent recites (emphasis added):

9. A method for transferring data between a bus master and a plurality of memory locations at respective addresses in an address space of a *secondary memory*, for use with a host processing unit and a first cache memory which caches memory locations of said *secondary memory* for said host processing unit, said first cache memory having a line size of 1 bytes, comprising the steps of:

sequentially transferring at least three data units between said bus master and said *secondary memory* beginning at a first starting memory location address in said *secondary memory* address space and continuing sequentially beyond an 1-byte boundary of said *secondary memory* address space; and

prior to completion of the transfer of the first data unit beyond said 1-byte boundary, determining whether an N+1'th 1-byte line of said *secondary memory* is cached in a modified state in said first cache memory, said N+1'th 1-byte line being the line of said *secondary memory* which includes said first data unit beyond said 1-byte boundary,

all of said transfers of data units in said step of sequentially transferring, occurring at a constant rate.

The specification discloses:

The second level (L2) cache is logically behind the first level cache, and DRAM memory (which in this case can be referred to as tertiary memory) is located logically behind the second level cache.

\* \* \*

Note that different embodiments can have a wide variety of different kinds of host processing subsystems. For example, they can include a "level 0" cache between the CPU and the L1 cache; they can include one or multiple processors; they can include bridges between the host bus 112 and a bus protocol expected by a CPU in the host processing subsystem, and so on. As a group, however, all the components of the host processing subsystem use an L1 cache to cache at least some lines of the secondary memory address space.

('906 Patent at 3:2-3 & 9:11-20.)

In *AMD* and *Apple*, the parties agreed upon the construction of "secondary memory," so the Court did not consider any arguments as to that term. *See AMD Markman* at 2; *Apple Markman* at 1-2.

Figure 1 is reproduced here and illustrates "DRAM" and "L2 cache" within a box labeled "Secondary Memory":

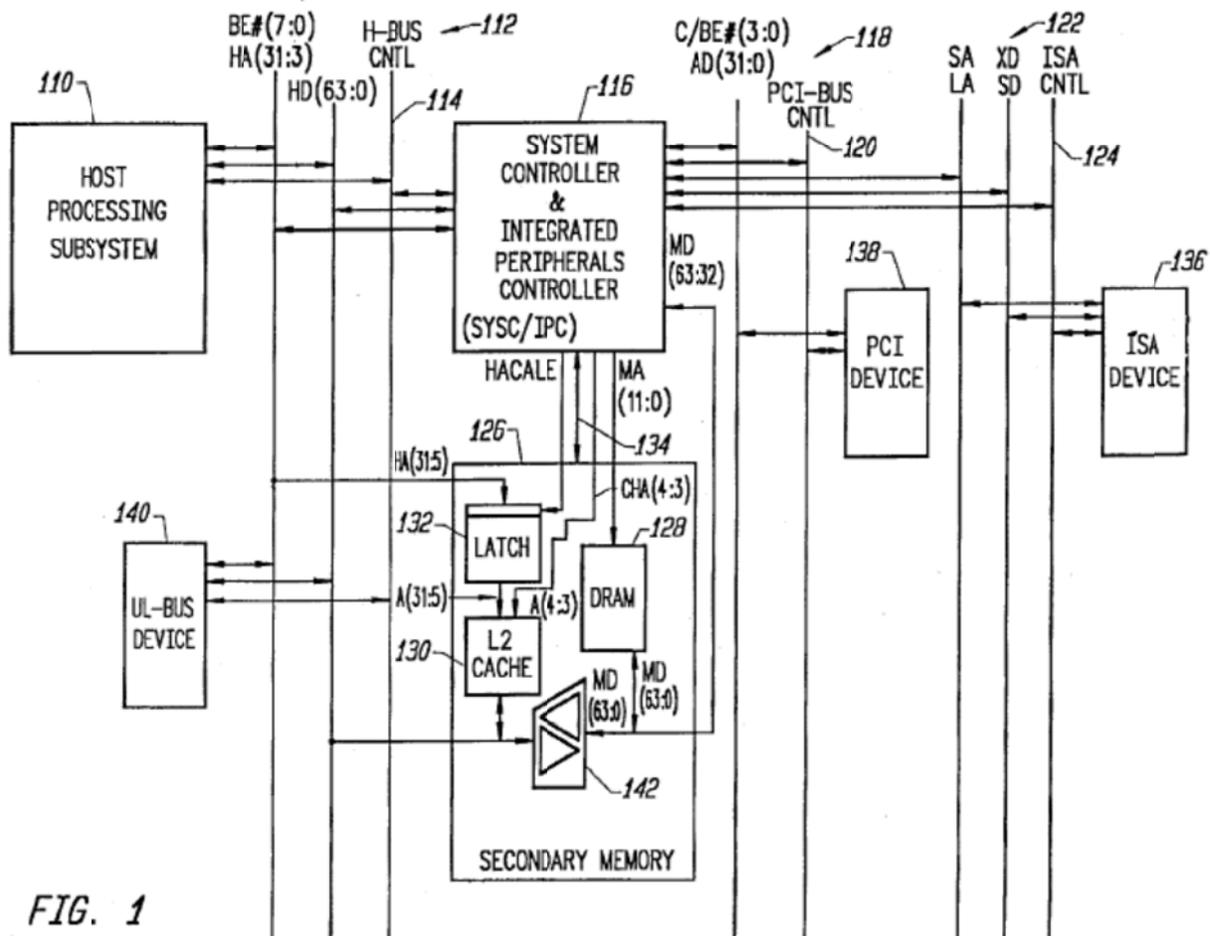


FIG. 1

The phrase “i.e.,” in Plaintiff’s proposed construction, is an abbreviation for the Latin *id est*, meaning “that is.” (Dkt. No. 139, Ex. D, *Webster’s Third New International Dictionary* 1124 (2002).) Plaintiff has failed to support limiting the construction to require “DRAM memory and, if present, L2 and L3 cache memory.” The general term “secondary memory” should not be limited to the embodiment depicted in Figure 1 and described in the specification. *See Electro Med.*, 34 F.3d at 1054 (“[A]lthough the specifications may well indicate that certain embodiments are preferred, particular embodiments appearing in a specification will not be read into the claims when the claim language is broader than such embodiments.”). Finally, even if

Plaintiff’s proposal were modified to replace “i.e.” with “such as,” the listed examples are unnecessary and redundant and would risk being improperly perceived as limiting.

The Court therefore hereby construes “**secondary memory**” to mean “**memory located logically behind the first cache memory.**”

**C. “first cache memory” (‘906 Patent, Claim 9)**

Plaintiff’s Proposed Construction	Defendants’ Proposed Construction
“the first level of cache memory, commonly referred to as L1 cache memory”	<p>“the first level of cache memory that uses a write-back protocol”</p> <p>Previously: “any memory logically located between the host processing unit and secondary memory that holds recently accessed code or data”</p>

(Dkt. No. 135, at 16; Dkt. No. 139, at 8.)

(1) The Parties’ Positions

Plaintiff argues that its proposed construction was the construction reached in *AMD* and *Apple*, that “‘first cache memory’ is the first level of cache memory disclosed in the Patent,” and that “the Patent is clear that ‘L1 cache memory’ and ‘first cache memory’ are one and the same, such that it uses them interchangeably.” (Dkt. No. 135, at 16.) Plaintiff further explains that “it is the L1 cache memory that is the subject of the ‘snoops’ or ‘inquiries’ that are called for in the claims, and which are carried out predictively by ‘the invention.’” (*Id.* (citing ‘906 Patent at 4:33-46, 5:17-25 & 7:64-8:13).)

Defendants note that in *AMD* and *Apple*, the parties agreed upon the construction of “first cache memory,” so the Court did not consider any arguments as to that term. (Dkt. No. 139, at 2.) Defendants argue that their proposed “restriction on the type of cache is appropriate because the other common type of cache protocol, write through protocol, does not require

snooping in order to maintain data integrity.” (*Id.*, at 9 (citing ‘291 Patent at 2:54-58).) As to Plaintiff’s proposal of “L1 cache memory,” Defendants argue that the patentees’ “decision to claim something broader than L1 cache memory by using the more generic phrase ‘first cache memory’ or ‘a cache memory’ should be honored.” (*Id.*, at 9.) Defendants submit that Plaintiff seeks to add additional, unsupported limitations in order to avoid Defendants’ prior art 505 Chip, which in at least one instance was implemented with “write-through protocol for the first level of cache memory and write-back protocol for the second level of cache memory.” (*Id.*, at 10.)

Plaintiff replies that “the first cache memory is simply the first level of cache memory—namely, the L1 cache—as has been stipulated in all prior cases, and as the Patent makes clear.” (Dkt. No. 143, at 3.) Plaintiff also highlights its position, discussed in subsection IV.B., above, that a second or higher level cache is part of the “secondary memory” and therefore cannot be a “first cache memory.” (*Id.*, at 4.) Plaintiff further argues that when the specification states that “[t]he invention is useful whenever an L1 cache is present which can use a write back protocol,” “[i]t means that the invention is useful whenever the L1 cache can use a write-back protocol.” (*Id.* (quoting ‘906 Patent at 6:37-38).) Finally, Plaintiff argues that “for a variety of reasons having to do with the ways in which memory transfers are performed, . . . presnooping the L2 cache is of much less value.” (*Id.*, at 4 n.3.)

## (2) Analysis

Claim 9 of the ‘906 Patent recites (emphasis added):

9. A method for transferring data between a bus master and a plurality of memory locations at respective addresses in an address space of a secondary memory, for use with a host processing unit and a *first cache memory* which caches memory locations of said secondary memory for said host processing unit, said *first cache memory* having a line size of 1 bytes, comprising the steps of:  
sequentially transferring at least three data units between said bus master and said secondary memory beginning at a first starting memory location address

in said secondary memory address space and continuing sequentially beyond an l-byte boundary of said secondary memory address space; and

prior to completion of the transfer of the first data unit beyond said l-byte boundary, determining whether an N+1'th l-byte line of said secondary memory is cached in a modified state in said *first cache memory*, said N+1'th l-byte line being the line of said secondary memory which includes said first data unit beyond said l-byte boundary,

all of said transfers of data units in said step of sequentially transferring, occurring at a constant rate.

The disputed term “first cache memory” does not appear outside of the claims, but the specification repeatedly refers to “cache memory” in levels such as “L1” and “L2”:

Many IBM PC AT-compatible computers today include one, and usually two, levels of cache memory. A cache memory is a high-speed memory that is positioned between a microprocessor and main memory in a computer system in order to improve system performance. Cache memories (or caches) store copies of portions of main memory data that are actively being used by the central processing unit (CPU) while a program is running. Since the access time of a cache can be faster than that of main memory, the overall access time can be reduced.

\* \* \*

A computer system can have more than one level of cache memory for a given address space. For example, in a two-level cache system, the “level one” (L1) cache is logically adjacent to the host processor. The second level (L2) cache is logically behind the first level cache, and DRAM memory (which in this case can be referred to as tertiary memory) is located logically behind the second level cache. When the host processor performs an access to an address in the memory address space, the first level cache responds if possible. If the first level cache cannot respond (for example, because of an L1 cache miss), then the second level cache responds if possible. If the second level cache also cannot respond, then the access is made to DRAM itself. The host processor does not need to know how many levels of caching are present in the system or indeed that any caching exists at all. Similarly, the first level cache does not need to know whether a second level of caching exists prior to the DRAM. Thus, to the host processing unit, the combination of both caches and DRAM is considered merely as a single main memory structure. Similarly, to the L1 cache, the combination of the L2 cache and DRAM is considered simply as a single main memory structure. In fact, a third level of caching could be included between the L2 cache and the actual DRAM, and the L2 cache would still consider the combination of L3 and DRAM as a single main memory structure.

(‘906 Patent at 1:48-57 & 2:66-3:24.)

The parties dispute whether the term “first cache memory” refers to the first level of “write-back” cache or refers to the first level of cache that is logically adjacent to the host processor. The specification explains the difference between “write-through” and “write-back”:

When the CPU executes instructions that modify the contents of the cache, these modifications must also be made in the main memory or the data in main memory will become “stale.” There are two conventional techniques for keeping the contents of the main memory consistent with that of the cache—(1) the write-through method and (2) the write-back or copy-back method. In the *write-through method*, on a cache write hit, data is written to the main memory immediately after or while data is written into the cache. This enables the contents of the main memory always to be valid and consistent with that of the cache. In the *write-back method*, on a cache write hit, the system writes data into the cache and sets a “dirty bit” which indicates that a data word has been written into the cache but not into the main memory. A cache controller checks for a dirty bit before overwriting any line of data in the cache, and if set, writes the line of data out to main memory before loading the cache with new data.

(‘906 Patent at 2:48-65 (emphasis added).)

The Summary of the Invention states:

The invention is useful whenever an L1 cache is present which can use a write back protocol, and which supports inquire cycles, and whenever an I/O bus is present which has a linear-incrementing capability or mode which can continue beyond an L1 cache line boundary.

(‘906 Patent at 6:37-41.)

Sometimes, “the specification may reveal an intentional disclaimer, or disavowal, of claim scope by the inventor,” in which case “the inventor has dictated the correct claim scope, and the inventor’s intention, as expressed in the specification, is regarded as dispositive.”

*Phillips*, 415 F.3d at 1316. The above-quoted disclosure of when the invention is “useful,” however, does not amount to a disclaimer. *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 909 (Fed. Cir. 2004) (“Absent a clear disclaimer of particular subject matter, the fact that the inventor may have anticipated that the invention would be used in a particular way does not mean that the scope of the invention is limited to that context.”) (quoting *Northrop Grumman*

*Corp. v. Intel Corp.*, 325 F.3d 1346, 1355 (Fed. Cir. 2003).) The term “first cache memory” therefore refers to the first level of cache.

The Court hereby construes “**first cache memory**” to mean “**the first level of cache memory, commonly referred to as L1 cache memory.**”

**D. “at a constant rate” (‘906 Patent, Claim 9)**

<b>Plaintiff’s Proposed Construction</b>	<b>Defendants’ Proposed Construction</b>
“a uniform rate”	“with at most a short delay”

(Dkt. No. 135, at 17; Dkt. No. 139, at 10.)

(1) The Parties’ Positions

Plaintiff submits that this disputed term was “one of the more hotly contested terms” in *AMD*. (Dkt. No. 135, at 17.) In *AMD*, the defendant proposed that “constant rate” be construed to mean that “the same number of wait states is inserted between the transfer of each data unit, without inserting any additional wait states.” (No. 2:06-CV-477, Dkt. No. 75, 6/30/2008 P.R. 4-5(d) Joint Claim Construction Chart, Ex. A, at 9.) The *AMD Markman* reached the construction that Plaintiff now proposes here. (Dkt. No. 135, at 17.) Plaintiff also notes that in *Apple*, the parties agreed upon the construction that Plaintiff here proposes. (*Id.*)

Plaintiff also argues that Defendants have shown no reason to depart from the ordinary meaning of the disputed term. (*Id.*) Plaintiff submits that “the ordinary language meaning of ‘constant’ that we are familiar with in everyday discourse is no different from the definition of the term to be found in technical dictionaries.” (*Id.*) Plaintiff further argues that Defendants have shown no basis in the claims, the specification, or the prosecution history for departing from the ordinary meaning. (*Id.*, at 17-18 (citing *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994) (“[W]hen interpreting a claim, words of the claim are generally given their ordinary and

accustomed meaning, unless it appears from the specification or the file history that they were used differently by the inventor.”)).)

Defendants respond that they here present arguments not raised in the *AMD* case. (Dkt. No. 139, at 2.) Defendants argue that Plaintiff’s proposal of “[u]niform’ is no more or less clear than the word ‘constant.’” (*Id.*, at 10.) Defendants then provide technical background regarding PCI burst transfer and submit that “Figure 4 provides conclusive evidence that ‘constant rate’ does not mean without any wait states” because “the only embodiments shown in the Patents have a wait state (a short delay) at each cache line boundary.” (*Id.*, at 14.) Defendants argue that Plaintiff’s proposal is based on dictionary definitions and on portions of the specification that Plaintiff has taken out of context. (*Id.*, at 15.) Defendants submit that “the inventors did not consider a one cycle wait state to be a ‘delay’ at all” and that “the goal of the inventors was not to necessarily eliminate delays, but to minimize delays.” (*Id.*, at 15-16.)

Finally, Defendants cite prosecution history in which the patentee distinguished Defendants’ 505 Chip (in an inaccurate description of the 505 Chip, according to Defendants) as having to “insert at least one wait state between the second and third cache lines. The [505 Chip] diagram would then fail to teach that the VT82C505 can sustain a constant rate into the third cache line.” (*Id.*, at 16 (quoting Ex. E, 12/10/2001 Request for Further Consideration of IDS Documents D2-1 through D2-7, at p. 31 (OPTI-SIS\_VIA016016))).)

Plaintiff replies that the “short delay” referenced in the specification “occurs only under certain circumstances—namely, when a transfer begins near the end of a cache line, and a ‘short delay’ is necessary for the predictive snoop to complete.” (Dkt. No. 143, at 6 (citing ‘906 Patent at 17:40-18:35, 17:41-42 & 18:26-31).) Plaintiff also replies that “in Fig. 4 a wait state is inserted before every data transfer, so that, as the text discussing the figure makes clear ‘in the

situation illustrated in FIG 4, all of the data transfers take place at a constant rate, specifically, one Dword in every two PCICLK cycles, even as the burst continues beyond the cache line boundary.” (*Id.*, at 7 (quoting ‘906 Patent at 14:31-33).)

## (2) Analysis

Claim 9 of the ‘906 Patent recites (emphasis added):

9. A method for transferring data between a bus master and a plurality of memory locations at respective addresses in an address space of a secondary memory, for use with a host processing unit and a first cache memory which caches memory locations of said secondary memory for said host processing unit, said first cache memory having a line size of 1 bytes, comprising the steps of:

    sequentially transferring at least three data units between said bus master and said secondary memory beginning at a first starting memory location address in said secondary memory address space and continuing sequentially beyond an 1-byte boundary of said secondary memory address space; and

    prior to completion of the transfer of the first data unit beyond said 1-byte boundary, determining whether an N+1’th 1-byte line of said secondary memory is cached in a modified state in said first cache memory, said N+1’th 1-byte line being the line of said secondary memory which includes said first data unit beyond said 1-byte boundary,

    all of said transfers of data units in said step of sequentially transferring, occurring *at a constant rate*.

Plaintiff has cited dictionary definitions of “constant” as meaning uniform: *Webster’s Ninth New Collegiate Dictionary* at 281 (1988) (“something invariable or unchanging”); *Merriam-Webster Online Dictionary* (2012) (“invariable, uniform”); *Oxford Dictionary of Computing*, (5th ed. 2004), at p. 111 (“constant 1. a quantity or data item whose value does not change.”). (Dkt. No. 135, at Exs. 8 & 9.) Extrinsic dictionary definitions, however, are generally not a reliable starting point for construction because “heavy 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.” *Phillips*, 415 F.3d at 1321. Also of note, the above-quoted technical definition of “constant” defines a noun whereas the disputed term uses “constant” as an adjective.

Turning then to the intrinsic evidence, the specification discloses “at most a short delay” or, more specifically, “no delay is incurred” by moving from one cache line to the next:

The controller then allows the burst access to take place between secondary memory and the PCI-bus master, and simultaneously and predictively, performs an inquire cycle of the L1 cache for the next cache line. In this manner, if the PCI burst does in fact continue past the cache line boundary, the new inquire cycle will already have taken place (or will already be in progress), thereby allowing the burst to proceed *with at most a short delay* absent a hit-modified condition. This avoids the need to incur the penalty of stopping the transfer on the PCI bus and restarting it anew at a later time, every time a linear burst transaction crosses a cache line boundary.

(‘906 Patent at 6:13-24 (emphasis added).)

The last Dword<sup>2</sup> in the cache line-sized block of DRAM 28, Dword 1C, is transferred to the PCI device 138 on the rising edge of PCICLK which begins PCICLK cycle 54/55. Note, however, that *no delay is incurred before the transfer of Dword 20, which is the first Dword of the next cache line address*. In fact, in the situation illustrated in FIG. 4, all of the data transfers in the burst take place at a *constant rate*, specifically one Dword in every two PCICLK cycles, even as the burst continues beyond the cache line boundary. This is a consequence of the features of the present embodiment of the invention.

In order to minimize or eliminate delays at cache line boundaries, as previously described, the system controller 116 performs a predictive snoop (“pre-snoop”) of the second cache line address of the burst, prior to completion of the last PCI-bus data transfer from the initial cache line address of the burst. . . .

(*Id.* at 14:26-41 (emphasis added).)

In order to accomplish pre-snoop, the system controller 116 detects the first PCI-bus data transfer by sampling IRDY# and TRDY# asserted at the beginning of PCICLK cycle 26/27. It then increments the cache line address on HA(31:5) at the beginning of PCICLK cycle 28/29, to refer to the next sequential cache line address (line address 20). System controller 116 then, in HCLK cycle 32, asserts EADS# to initiate an inquire cycle of the L1 cache 212 in the host processing subsystem 110. Two HCLK cycles later, at the beginning of HCLK cycle 35, the system controller 116 samples HITM# negated. Thus, the inquiry cycle for the second cache line has been completed before the last data transfer takes place in

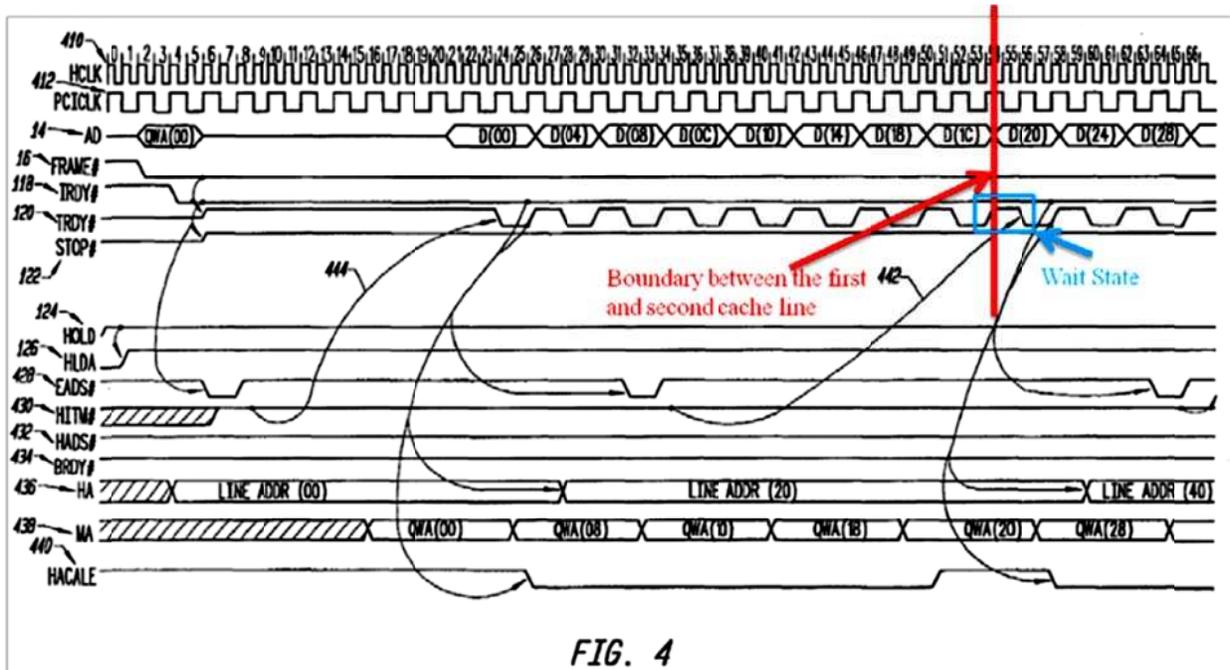
---

<sup>2</sup> In the patents-in-suit, a “word” is two bytes. A “DWord” is a “double word,” which is four bytes. (See ‘906 Patent at 13:56-59; see also Dkt. No. 135, Ex. 7, OPTI-SIS\_VIA011048, *PCI Local Bus Specification* 188 (rev. 2.0, Apr. 30, 1993).) A “quad word,” in turn, is eight bytes. A 64-byte line therefore contains eight quad words or, alternatively, 16 Dwords.

the first cache line. Assuming the first transfer does in fact proceed beyond the cache line boundary, *the first data transfer (Dword 20) of the second line of data can take place without stopping the burst and without inserting any additional PCI-bus wait states* (see arrow 442).

(*Id.* at 14:50-67 (emphasis added).)

On one hand, Plaintiff’s proposal of “uniform” should be rejected as vague and as ineffective in resolving whether there can be any wait states during transfer. Also, Figure 4, as annotated by Defendants in their response brief, is illustrative that the presence of “wait states,” or at least some period of waiting time, is not inconsistent with the absence of interruption between cache lines during a burst of sequential data transfers:



To whatever extent Plaintiff’s proposal seeks to exclude this embodiment, Plaintiff’s proposal is disfavored. *C.R. Bard, Inc. v. U.S. Surgical Corp.*, 388 F.3d 858, 865 (Fed. Cir. 2004) (citing *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1583 (Fed. Cir. 1996)). Although Plaintiff’s proposal is the construction reached in the *AMD* case, the Court there provided no analysis that

would dissuade the Court from more clearly resolving the parties' present dispute in the case at bar. *See AMD Markman* at 5.

On the other hand, although Defendants' proposal of "with at most a short delay" is supported by the specification, as quoted above (*see* '906 Patent at 6:13-24), the phrase "a short delay" is also vague as to what a "delay" is and how long it can be before it is no longer "short." The phrase also fails to effectively resolve the parties' dispute.

Instead, in the context of the claim and the above-quoted passages from the specification, the disclosure that "the first data transfer (Dword 20) of the second line of data can take place *without stopping the burst* and without inserting any additional PCI-bus wait states" is more accurate and more precise and will be more helpful to the finder of fact. (*See id.* at 14:50-67.) Construction in accordance with this above-quoted disclosure is also consistent with the prosecution history cited by Defendants. (*See* Dkt. No. 139, at 16.) The reference in the specification to "without inserting any additional PCI-bus wait states" should be omitted, however, because the claim is not limited to a PCI bus and because a reference to "additional wait states" would be duplicative of "without stopping the burst" and might serve to confuse rather than clarify.

To properly explain the word "burst," the construction should explain that the "burst" is a burst transfer, which comprises a series of sequential transfers as set forth in the claim. Finally, to avoid any misreading of "stopping" to refer to a "STOP#" signal (which is used to terminate a PCI transaction), the construction should use the word "delaying." (*See* '906 Patent at 11:23-25 & 14:26-41.)

The Court therefore hereby construes "**at a constant rate**" to mean "**without delaying the transfer of sequential data within a burst transfer.**"

**E. “means for sequentially transferring . . .” (‘906 Patent, Claim 26)**

The full disputed term is “means for sequentially transferring at least three data units between said bus master and said secondary memory beginning at a first starting memory location address in said secondary memory address space and continuing sequentially beyond an l-byte boundary of said secondary memory address space.”

Plaintiff's Proposed Construction	Defendants' Proposed Construction
<p>Construction under 35 U.S.C. §112, par. 6</p> <p>[Plaintiff] contends that the disclosed structure which meets limitation 26.2 [(“means for sequentially transferring . . .”)] is a core logic device or devices such as the system controller (SYSC) and integrated peripherals controller (IPC) that moves data between the bus connected to the secondary memory and the bus connected to the Bus Master and known equivalents thereof</p>	<p>Construction under 35 U.S.C. §112, par. 6</p> <p>Indefinite<sup>3</sup></p> <p>In the parties' P.R. 4-5(d) Joint Claim Construction Chart, Defendants propose:</p> <p>To the extent this term is not indefinite the structure associated with it should include at least the following structures:</p> <p>the partially disclosed circuitry inside the System Controller &amp; Integrated Peripherals Controller (SYSC/IPC) 116 in Fig. 1 that moves data between secondary memory and the PCI Bus, the SYSC/IPC further comprising:</p> <p>the following integrated circuit chips[:]</p> <p>[OPTI, Inc. 82C557 (SYS) and 82C558 (IPC) as described in] “Viper-M 82C556M/82C557M/82C558M, Data Book, Version 1.0” (April 1995), and an OPTi, Inc. 82C556 data buffer controller (DBC), also described in the above-incorporated data book, which includes some additional buffers. ‘906 Patent, col. 9, ll. 30-38; ‘291 Patent, col. 9, ll. 34-42.<sup>4</sup></p>

<sup>3</sup> In the parties' pre-briefing claim chart, Defendants proposed (Dkt. No. 130, Ex. A, at 4):

This term is indefinite because the specification does not disclose sufficient structure for performing the claimed function.

To the extent this term is not indefinite the structure associated with it should include at least the following structures:

the partially disclosed circuitry inside the System Controller & Integrated Peripherals Controller 116 in Fig. 1 that moves data between the H-Bus and the PCI Bus

<sup>4</sup> Plaintiff objects that Defendants have “abandoned and waived all arguments not addressed in its Markman briefing; namely, the alternative structures for Claim 26 of the ‘906 Patent and Claims 88 and 89 of the ‘291 Patent that [Defendants] attempt[] to disclose in this chart.” (Dkt. No. 142, Ex. A, at 5 n.2.) At the December 18, 2012 hearing, however, Plaintiff was agreeable to

(Dkt. No. 135, at 18; Dkt. No. 139, at 22; Dkt. No. 142, Ex. A, at 5-6.)

(1) The Parties' Positions

Plaintiff submits that the *NVIDIA Markman* found the following similar term in Claim 21 was *not* indefinite: “means for sequentially transferring data units between said bus master and said secondary memory beginning at a starting memory location address in said secondary memory address space . . . said sequentially transferred data units including a last data unit before said 1-byte boundary and a first data unit beyond said 1-byte boundary.” (Dkt. No. 135, at 19.) Plaintiff argues that the same analysis applies to Claim 26, in which the presently disputed term appears, and that Figure 1 illustrates the corresponding structure. (*Id.*, at 19-20.)

Defendants respond that Plaintiff “has failed to meet [the] statutory requirements [of 35 U.S.C. § 112, ¶ 6] because it simply drew a ‘black box,’ gave it a nonce name, SYSC/IPC, and then ascribed all or virtually all of the claimed inventive attributes to this allegedly new black box.” (Dkt. No. 139, at 22.) Defendants urge that “[t]he term System Controller & Integrated Peripherals Controller (‘SYSC/IPC’) has no reasonably well understood meaning in the art.” (*Id.*, at 23.) Defendants cite Figure 1 and argue that in the data transfer illustrated in Figure 1, “[a]ccommodating th[e] speed disparity between the host bus and the PCI bus requires structure to implement[, but t]he ‘906 Patent does not disclose this structure.” (*Id.*, at 24.) Defendants also submit the Declaration of Joe McAlexander in support of their indefiniteness arguments. (*Id.*, Ex. F, 11/27/2012 McAlexander Decl.)

Plaintiff replies as to all of the means-plus-function terms collectively. Plaintiff argues that the Declaration of Joe McAlexander is untimely, and Plaintiff has therefore filed a motion to

---

including the 82C556 (DBC), 82C557 (SYSC), and 82C558 (IPC) as part of the corresponding structure.

strike the declaration. (Dkt. No. 143, at 9; *see* Dkt. No. 145.) That motion was not yet ripe at the time of the December 18, 2012 claim construction hearing. Alternatively, Plaintiff notes that the specification gives examples of “off the shelf” circuit chips for the SYSC/IPC 116. (Dkt. No. 143, at 10 (discussing ‘906 Patent at 9:30-39).) Finally, Plaintiff submits that “[i]n addition, of course, the Patent contains extensive diagrams of the logic circuitry that executed the claimed functions (sequentially transferring data and determining whether data was cached in a modified state), and multiple timing diagrams that illustrated the SYSC/IPC’s operation in performing these functions, as well as more than 29 columns of text explaining its figures.” (*Id.* (citing 3:32-4:46, 5:6-16, 11:57-12:5, 12:20-15:16, 20:64-24:35, 27:49-29:48 & Figs. 4, 8-9 & 12).)

## (2) Analysis

Title 35 U.S.C. § 112 ¶ 6, allows a patentee to express a claim limitation as “a means or step for performing a specified function without the recital of structure, material, or acts in support thereof.” *See Inventio AG v. Thyssenkrupp Elevator Ams.*, 649 F.3d 1350, 1355-56 (Fed. Cir. 2011). The Federal Circuit has further clarified what such functional claiming requires:

Thus, in return for generic claiming ability, the applicant must indicate in the specification what structure constitutes the means. If the specification is not clear as to the structure that the patentee intends to correspond to the claimed function, then the patentee has not paid the price but is rather attempting to claim in functional terms unbounded by any reference to structure in the specification. Thus, if an applicant fails to set forth an adequate disclosure, the applicant has in effect failed to particularly point out and distinctly claim the invention as required by the second paragraph of § 112.

*Biomedino, LLC v. Waters Techs. Corp.*, 490 F.3d 946, 948 (Fed. Cir. 2007) (citations and internal quotation marks omitted). Failure to disclose adequate structure corresponding to the claimed function results in the claim being invalid for indefiniteness. *See, e.g., Tech. Licensing Corp. v. Videotek, Inc.*, 545 F.3d 1316, 1338 (Fed. Cir. 2008).

“Although one of skill in the art may have been able to find a structure that would work, that does not satisfy § 112 ¶ 6. Under § 112 ¶ 6, a patentee is only entitled to ‘corresponding structure . . . described in the specification and equivalents thereof,’ not any device capable of performing the function.” *Ergo Licensing, LLC v. Carefusion 303, Inc.*, 673 F.3d 1361, 1364 (Fed. Cir. 2012) (citing *Blackboard, Inc. v. Desire2Learn Inc.*, 574 F.3d 1371, 1385 (Fed. Cir. 2009)) (emphasis in original).

Claim 26 of the ‘906 Patent recites:

26. Apparatus for transferring data between a bus master and a plurality of memory locations at respective addresses in an address space of a secondary memory, for use with a host processing unit and a first cache memory which caches memory locations of said secondary memory for said host processing unit, said first cache memory having a line size of 1 bytes, comprising:

*means for sequentially transferring at least three data units between said bus master and said secondary memory beginning at a first starting memory location address in said secondary memory address space and continuing sequentially beyond an 1-byte boundary of said secondary memory address space;*  
and

means for, prior to completion of the transfer of the first data unit beyond said 1-byte boundary, determining whether an N+1<sup>th</sup> 1-byte line of said secondary memory is cached in a modified state in said first cache memory, said N+1<sup>th</sup> 1-byte line being the line of said secondary memory which includes said first data unit beyond said 1-byte boundary,

said means for sequentially transferring, transferring all of said data units at a constant rate.

The parties have not disputed the claimed function for this means-plus-function term.

Plaintiff relies upon the following disclosures as evidence for definiteness and corresponding structure (in addition to the Abstract):

According to the invention, roughly described, when a PCI-bus controller receives a request from a PCI-bus master to transfer data with an address in secondary memory, the controller performs an initial inquire cycle and withholds TRDY# to the PCI-bus master until any write-back cycle completes. The controller then allows the burst access to take place between secondary memory and the PCI-bus master, and simultaneously and predictively, performs an inquire cycle of the L1 cache for the next cache line. In this manner, if the PCI burst does in fact continue past the cache line boundary, the new inquire cycle will already have

taken place (or will already be in progress), thereby allowing the burst to proceed with at most a short delay absent a hit-modified condition. This avoids the need to incur the penalty of stopping the transfer on the PCI bus and restarting it anew at a later time, every time a linear burst transaction crosses a cache line boundary.

\* \* \*

A core logic chipset in the system includes a system controller (SYSC) and an integrated peripherals controller (IPC), indicated generally as 116 [in Figure 1].

\* \* \*

Returning to FIG. 1, the SYSC/IPC 116 comprises the following integrated circuit chips available from OPTi, Inc., Santa Clara, Calif.: 82C557 (SYSC) and 82C558 (IPC). These chips are described in OPTi, Inc., “Viper-M 82C556M/82C557M/82C558M, Data Book, Version 1.0” (April 1995), incorporated by reference herein. The chipset also includes an OPTi, Inc. 82C556 data buffer controller (DBC), also described in the above-incorporated data book, which includes some buffers not shown in FIG. 1.

Briefly, the SYSC provides the control functions for interfacing with host processing subsystem 110, the 64-bit-wide L2 cache 130, the 64-bit DRAM 128 data bus, an interface to VL-bus aspects of the host bus 112, and an interface to the PCI-bus 118. The SYSC also controls the data flow between the host bus 112, the DRAM bus, the local buses, and the 8/16-bit ISA bus. The SYSC interprets and translates cycles from the CPU, PCI-bus masters, ISA-bus masters, and DMA to the secondary memory subsystem 126, local bus slaves, PCI-bus slaves, or ISA-bus devices.

The IPC contains an ISA-bus controller and includes the equivalent of an industry standard 82C206, a real time clock interface, a DMA controller, and a power management unit.

(‘906 Patent at 6:7-24, 7:1-3 & 9:30-52.) Figure 1 is reproduced in subsection IV.B., above.

General legal principles regarding indefiniteness are discussed in Section II., above. As to means-plus-function terms, “[i]f there is no structure in the specification corresponding to the means-plus-function limitation in the claims, the claim will be found invalid as indefinite.” *Biomedino*, 490 F.3d at 950. Further, “the written description must clearly link or associate structure to the claimed function.” *Telcordia Techs., Inc. v. Cisco Sys., Inc.*, 612 F.3d 1365, 1376 (Fed. Cir. 2010).

On balance, the disclosure of a “system controller” and an “integrated peripherals controller” is sufficient corresponding structure to avoid indefiniteness. Disclosure of specific circuitry, such as for accommodating the difference in the speeds of a host bus and a PCI bus, is not required. *See Intel Corp. v. VIA Techs., Inc.*, 319 F.3d 1357, 1365-67 (Fed. Cir. 2003) (finding that “core logic” was sufficient corresponding structure despite absence of disclosure of circuitry); *see also Tech. Licensing*, 545 F.3d at 1338-39; *S3 Inc. v. nVidia Corp.*, 259 F.3d 1364, 1370-71 (Fed. Cir. 2001). Defendants have thus not met their burden to prove indefiniteness by clear and convincing evidence. *See Halliburton*, 514 F.3d at 1249-50.

Finally, as to Plaintiff’s proposal of “a core logic device or devices such as [the SYSC/IPC 116],” the SYSC/IPC 116 is the only disclosed structure that performs the claimed function. Specifically, the SYSC/IPC 116 includes Plaintiff’s “82C557 (SYSC) and 82C558 (IPC)” chips, as well as Plaintiff’s “82C556 data buffer controller (DBC)” chip. (*See* ‘906 Patent at 9:31-39.) The corresponding structure should include those identified chips. *Cf. Odetics, Inc. v. Storage Tech. Corp.*, 185 F.3d 1259, 1268 (Fed. Cir. 1999) (noting, in the context of a statutory equivalence analysis under 35 U.S.C. § 112, ¶ 6, that “the claim limitation is the overall structure corresponding to the claimed function” and that “[f]urther deconstruction or parsing is incorrect”).

The Court therefore finds that the “**means for sequentially transferring . . .**” term is not indefinite and that the corresponding structure is “**a system controller and an integrated peripherals controller (SYSC/IPC) 116, including Plaintiff’s 82C556 (DBC), 82C557 (SYSC), and 82C558 (IPC), and equivalents thereof.**”

**F. “means for . . . determining . . .” (‘906 Patent, Claim 26)**

The full disputed term is: “means for, prior to completion of the transfer of the first data unit beyond said l-byte boundary, determining whether an N+1’th l-byte line of said secondary

memory is cached in a modified state in said first cache memory, said N+1'th 1-byte line being the line of said secondary memory which includes said first data unit beyond said 1-byte boundary.”

Plaintiff's Proposed Construction	Defendants' Proposed Construction
<p>Construction under 35 U.S.C. §112, par. 6</p> <p>[Plaintiff] contends that the disclosed structure which meets limitation 26.3 [(“means for . . . determining . . .”)] is the circuitry shown in Figs. 8 &amp; 9 that generates the PSNSTR1 and EADS# signals and known equivalents thereof</p>	<p>Construction under 35 U.S.C. §112, par. 6</p> <p>Indefinite</p> <p>In the parties' P.R. 4-5(d) Joint Claim Construction Chart, Defendants propose:</p> <p>This term is indefinite because the specification does not disclose sufficient structure for performing the claimed function. To the extent this term is not indefinite the structure associated with it should include at least the following structures:</p> <p>The circuitry shown in Figs. 8 &amp; 9 that generates the EADS# signal. The EADS# signal causes the function of determining if the secondary memory is cached in a modified state in the cache memory to be performed by an external cache controller whose circuitry [the patentee] did not disclose.<sup>5</sup></p>

(Dkt. No. 135, at 18-19; Dkt. No. 139, at 26; Dkt. No. 130, Ex. A, at 5-6.)

(1) The Parties' Positions

Plaintiff submits that the parties in *NVIDIA* agreed that a similar limitation in Claim 21 was not indefinite, and the Court then adopted Plaintiff's proposal that the corresponding structure for the term in Claim 21 was “the logic circuitry of the SYSC/IPC schematically illustrated in Figure 9 of the patent.” (Dkt. No. 135, at 21; *see NVIDIA Markman* at 28-29.) Plaintiff cites disclosure in the patent that “PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer” and “PSNSTR1 is

---

<sup>5</sup> This proposal by Defendants also appears in the parties' pre-briefing claim chart. (Dkt. No. 130, Ex. A, at 5-6.)

provided to an input of NAND gate 822 in FIG. 8 and, like LT2, initiates an L1 cache inquiry cycle.” (‘906 Patent at 23:34-36 & 24:32-35.)

Defendants respond that Plaintiff “identifies the circuitry shown in Figs. 8 & 9 that generate two signals PSNSTR1 and EADS# that only accomplish a portion of the claimed function.” (Dkt. No. 139, at 26.) Defendants submit that “[t]he disclosure of the ‘906 Patent contains no structure for generating an address of the next cache line.” (*Id.*) Defendants also submit the Declaration of Joe McAlexander in support of their indefiniteness arguments. (*Id.*, Ex. F, 11/27/2012 McAlexander Decl.)

Plaintiff replies to all of the means-plus-function terms collectively, as discussed in subsection IV.E.(1), above.

## (2) Analysis

26. Apparatus for transferring data between a bus master and a plurality of memory locations at respective addresses in an address space of a secondary memory, for use with a host processing unit and a first cache memory which caches memory locations of said secondary memory for said host processing unit, said first cache memory having a line size of  $l$  bytes, comprising:

means for sequentially transferring at least three data units between said bus master and said secondary memory beginning at a first starting memory location address in said secondary memory address space and continuing sequentially beyond an  $l$ -byte boundary of said secondary memory address space; and

*means for, prior to completion of the transfer of the first data unit beyond said  $l$ -byte boundary, determining whether an  $N+1$ 'th  $l$ -byte line of said secondary memory is cached in a modified state in said first cache memory, said  $N+1$ 'th  $l$ -byte line being the line of said secondary memory which includes said first data unit beyond said  $l$ -byte boundary,*

said means for sequentially transferring, transferring all of said data units at a constant rate.

The specification discloses:

FIG. 9 is a schematic diagram of circuitry in the system controller 116 which produces the PSNSTR1 signal used in FIG. 8. As previously mentioned, PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer.

\* \* \*

As previously described, PSNSTR1 is provided to an input of NAND gate 822 in FIG. 8 and, like LT2, initiates an L1 cache inquiry cycle.

(‘906 Patent at 23:34-36 & 24:33-35.)

General legal principles regarding indefiniteness are discussed in Section II., above. As to means-plus-function terms, “[i]f there is no structure in the specification corresponding to the means-plus-function limitation in the claims, the claim will be found invalid as indefinite.” *Biomedino*, 490 F.3d at 950. Further, “the written description must clearly link or associate structure to the claimed function.” *Telcordia*, 612 F.3d at 1376.

The *NVIDIA Markman* considered an argument that “because the structure for incrementing the address for the next-line inquiry is ‘not shown,’ the algorithm cannot be described.” *NVIDIA Markman* at 29. The *NVIDIA Markman* analyzed and rejected that argument:

nVidia is asserting structure linked to a function of actual[ly] implementing a snoop rather than the recited function of initiating a snoop. The patent makes clear that the “next-line” inquiry is “initiated” by the PSNSTR1 signal. ‘906 patent Col. 23:34-36 (“PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer”); Col. 24:32-35 (“As previously described, PSNSTR1 is provided to an input of NAND gate 822 in FIG. 8 and, like LT2, initiates an L1 cache inquiry cycle”). Further, the patent identifies the structure that generates PSNSTR1:

FIG. 9 is a schematic diagram of circuitry in the system controller 116 which produces the PSNSTR1 signal used in FIG. 8. As previously mentioned, PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer.

‘906 patent, Col. 23:32-36; *see also id.*, Col. 23:37-24:35 and Fig. 9. Accordingly, the Court finds that the Pre-Snoop patents do disclose corresponding structure for the recited function and adopts [Plaintiff’s] proposed construction[:] [the logic circuitry of the SYSC/IPC schematically illustrated in Figure 9 of the patent].

*Id.* In the *NVIDIA* case, the disputed term addressed by the above-quoted passage was:

means for *initiating a next-line inquiry*, prior to completion of the transfer of the last data unit before said 1-byte boundary, to determine whether an N+1'th 1-byte line of said secondary memory is cached in a modified state in said first cache memory, said N+1'th 1-byte line being a line of said secondary memory which includes said first data unit beyond said 1-byte boundary

*Id.*, at 27 (emphasis modified). Thus, as the *NVIDIA Markman* noted, the function there at issue

was “initiating” a snoop. In the present case, the disputed term is:

means for, prior to completion of the transfer of the first data unit beyond said 1-byte boundary, *determining whether an N+1'th 1-byte line of said secondary memory is cached in a modified state in said first cache memory*, said N+1'th 1-byte line being the line of said secondary memory which includes said first data unit beyond said 1-byte boundary

The disputed term here thus requires actually “determining” whether the data in the first cache memory has been modified, which amounts to actually implementing a snoop.

The necessary corresponding structure is therefore the structure found in the *NVIDIA Markman* plus whatever additional structure is required for actually determining whether the cached data has been modified. Defendants submit that “whatever structure that is doing the determining must generate an address residing in that N+1'th 1-byte line and drive it onto the host bus.” (Dkt. No. 139, at 26.) Defendants conclude:

As with the “means for sequentially transferring” element above, this claim element appears to be performed by the SYSC/IPC black box, a component whose structure would not be known to a person of ordinary skill. Without a disclosure of the structure that generates the address for the determining means, a person of ordinary skill cannot identify a structure that performs all of the structure necessary to perform the claimed function. Consequently, Claim 26 of the '906 Patent is invalid for indefiniteness.

(*Id.*, at 27.)

On one hand, the *NVIDIA Markman* noted disclosure that circuitry to “increment the secondary memory line address” was “(not shown)”:

The output of NAND gate 910, FTRDTGB, is connected to the D input of a flip-flop 912, which is clocked on LCLKI. Flip-flop 912 thus delays FTRDTGB by one PCICLK to enable other circuitry (not shown) in the system controller 116 to increment the secondary memory line address on HA(31:5) (FIG. 1).

(‘906 Patent at 24:10-15; *see NVIDIA Markman* at 28-29.)

On the other hand, the specification discloses that the SYSC/IPC 116 “drives inquiry cycles”:

Because at least one line of L1 cache 212 supports a write-back protocol, the host processing subsystem 110 also supports inquire cycles, initiated by the external system to determine whether a line of secondary memory is currently being cached in the L1 cache 212 and whether it has been modified in that cache. An external bus master (external to the host processing subsystem 110) (*SYSC/IPC 116* in the system of FIG. 1) *drives inquire cycles to the host processing subsystem 110 prior to an access (read or write) to the secondary memory subsystem 126*, in order to ensure that the secondary memory subsystem 126 contains the latest copy of the data. If the host processing subsystem 110 has the latest copy of the data (i.e., the data is cached modified in the L1 cache 212), then, as soon as permitted by the SYSC 116 and at least for the Pentium processor, the Pentium performs a write-back of the specified data line before the access by the external master is allowed to take place.

(‘906 Patent at 7:64-8:13 (emphasis added).)

In PCI clock cycle 2/3, the PCI master device 138 places the dword address of the first desired transfer onto the AD lines of the PCI-bus 118. It also at this time places a command on the C/BE# lines of PCI-bus 118, and asserts FRAME# to the system controller 116. (See waveforms 414 and 416.) As mentioned, this address ends in `00`, and designates the first quad word in a cache-line-sized block of the secondary memory address space. The *system controller 116 translates this address onto the host bus address lines HA(31:3)* as illustrated in waveform 436 [in Figure 4].

(*Id.* at 12:63-13:5 (emphasis added).)

In order to minimize or eliminate delays at cache line boundaries, as previously described, the *system controller 116 performs a predictive snoop (“pre-snoop”) of the second cache line address of the burst*, prior to completion of the last PCI-bus data transfer from the initial cache line address of the burst. In fact, because the system controller 116 controls the DRAM address on MA(11:0) independently from addresses which the system controller 116 places on the host bus 112 HA(31:5) lines, the pre-snoop takes place simultaneously with at least one data transfer taking place on the PCI-bus 118. The predictive snoop is

“predictive” because it is performed even though the system controller 116 does not yet know whether the PCI device 138 desires to continue the burst beyond the cache line boundary.

In order to accomplish pre-snoop, the system controller 116 detects the first PCI-bus data transfer by sampling IRDY# and TRDY# asserted at the beginning of PCICLK cycle 26/27. *It then increments the cache line address on HA(31:5) at the beginning of PCICLK cycle 28/29, to refer to the next sequential cache line address (line address 20). System controller 116 then, in HCLK cycle 32, asserts EADS# to initiate an inquire cycle of the L1 cache 212 in the host processing subsystem 110.* Two HCLK cycles later, at the beginning of HCLK cycle 35, the system controller 116 samples HITM# negated. Thus, the inquiry cycle for the second cache line has been completed before the last data transfer takes place in the first cache line. Assuming the first transfer does in fact proceed beyond the cache line boundary, the first data transfer (Dword 20) of the second line of data can take place without stopping the burst and without inserting any additional PCI-bus wait states (see arrow 442).

In anticipation of the burst continuing beyond yet another cache line boundary, the *system controller 116 then performs a predictive snoop for the third cache line of the burst, again, while data is still being transferred from secondary memory addresses in the second cache line.* Specifically, at the beginning of PCICLK cycle 58-59, the system controller 116 samples both IRDY# and TRDY# asserted. *It increments the line address to the host processing subsystem 110 in HCLK cycle 60, and asserts EADS# in HCLK cycle 64.* HITM# is again sampled negated at the beginning of HCLK cycle 66, and once again the L1 cache inquiry cycle has been completed before the PCI-bus data transfers have reached the cache line boundary. The process continues until the PCI device 138 terminates the burst, or the inquiry cycle results in HITM# asserted. The latter situation is described below with respect to FIG. 6.

(*Id.* at 14:36-15:16 (emphasis added).)

In about PCICLK cycle 4/5, *the system controller 116 begins driving the second line address, predictively, onto the local bus 112 HA(31:5) address lines.*

(*Id.* at 16:25-27.)

On balance, these disclosures of a “system controller” and an “integrated peripherals controller” are sufficient corresponding structure to avoid indefiniteness. Disclosure of specific circuitry is not required. *See Intel*, 319 F.3d at 1365-67 (finding that “core logic” was sufficient corresponding structure despite absence of disclosure of circuitry); *see also Tech. Licensing*, 545

F.3d at 1338-39. Defendants have not met their burden to prove indefiniteness by clear and convincing evidence. *See Halliburton*, 514 F.3d at 1249-50. Nonetheless, the corresponding structures must include the SYSC/IPC 116, including Plaintiff’s “82C557 (SYSC) and 82C558 (IPC)” chips, as well as Plaintiff’s “82C556 data buffer controller (DBC)” chip. (*See* ‘906 Patent at 9:31-39; *cf. Odetics*, 185 F.3d at 1268.)

The Court therefore finds that the **“means for . . . determining . . .”** term is not indefinite and that the corresponding structure is **“a system controller and an integrated peripherals controller (SYSC/IPC) 116, including Plaintiff’s 82C556 (DBC), 82C557 (SYSC), and 82C558 (IPC) and including the circuitry shown in Figs. 8 & 9 that generates the PSNSTR1 and EADS# signals, and equivalents thereof.”**

**G. “initiating one and only one snoop access of said cache memory, said snoop accesses each specifying the respective N+1’th l-byte line” (‘291 Patent, Claims 73 & 88)**

Plaintiff’s Proposed Construction	Defendants’ Proposed Construction
“initiating one and only one next line inquiry”	“performing exactly one snoop of cache memory, said snoop specifying the N+1th L-byte line”

(Dkt. No. 135, at 22; Dkt. No. 139, at 18.)

(1) The Parties’ Positions

Plaintiff submits that the *Apple Markman* rejected a proposal similar to Defendants’ and that the *Apple Markman* instead adopted the construction that Plaintiff proposes here. (Dkt. No. 135, at 22-23.) Plaintiff also submits that in *AMD*, the parties agreed upon the construction that Plaintiff here proposes. (*Id.*, at 22.) Plaintiff also argues that “under [Defendants’] proposed construction, even though the N+1th line is being snooped exactly once, as contemplated by the specification and the claim, it is possible that a jury might be confused into believing that the claimed method is not being practiced because an additional and extraneous snoop meant that

more than ‘exactly one snoop’ was occurring.” (*Id.*, at 23.) Plaintiff further argues that “[t]he prosecution history confirms the clear intent of this language, as [Plaintiff] was explicit in stating that the ‘once and only once’ limitation had been added to differentiate the claimed invention from an alleged VIA prior art chipset that asserted multiple predictive snoop accesses for each cache line, thereby choking the CPU with unnecessary inquiry cycles.” (*Id.*)

Defendants respond that although this disputed term was the subject of an appeal in *Apple*, the parties settled before resolution of the appeal. (Dkt. No. 139, at 18.) Defendants therefore urge that the construction of this term should be revisited. (*Id.*) Defendants argue that “[a] construction that requires ‘exactly one snoop’ properly captures th[e] explicit claim language” “because the claim language states that the cache memory can only be snooped once during the transfer of a particular cache line.” (*Id.*, at 19.)

Plaintiff replies that “[i]f anything, the [*Apple*] settlement indicates Apple’s deep concern that the district court’s rulings would be affirmed.” (Dkt. No. 143, at 7.) Plaintiff notes that there can be multiple bus masters, any of which might access memory and thereby trigger a snoop. (*Id.*, at 8.) Plaintiff submits that the specification provides “no reason . . . for filtering out snoops that are not redundant and which are extraneous to the burst transfer, but may be critical to other functions being performed by other masters in the computer at the same time that a burst transfer is proceeding.” (*Id.*)

## (2) Analysis

Claims 73 and 88 of the ‘291 Patent recite (emphasis added):

73. A method for transferring a plurality of data units to a bus master from a respective plurality of memory locations at sequential memory location addresses in an address space of a secondary memory, for use with a host processing unit and a cache memory which caches memory locations of said secondary memory for said host processing unit, said cache memory having a line size of 1 bytes, and

each data unit having a size equal to the largest size that can be transferred to said bus master in parallel, comprising the steps of:

sequentially transferring data units to said bus master from said secondary memory according to a PCI-bus burst transaction, beginning at a starting memory location address in said secondary memory address space and continuing beyond at least first and second 1-byte boundaries of said secondary memory address space, each 1-byte line of said transaction requiring at least 8 data unit transfers to said bus master; and

during the transfer of the data units for each entire N'th 1-byte line in said step of transferring, *initiating one and only one snoop access of said cache memory, said snoop accesses each specifying the respective N+1'th 1-byte line* and being initiated early enough such that they can be sampled by said host processing unit prior to completion of the transfer to said bus master of the last data unit in the respective N'th 1-byte line,

wherein said step of transferring comprises the step of transferring to said bus master three sequential data units including the last data unit before said first 1-byte boundary and the first data unit beyond said first 1-byte line, all at a constant rate,

and wherein said step of transferring further comprises the step of transferring to said bus master three sequential data units including the last data unit before said second 1-byte boundary and the first data unit beyond said second 1-byte line, all at a constant rate.

\* \* \*

88. Controller apparatus for a computer system which includes a secondary memory having an address space, a bus master, a host processing unit and a cache memory which caches memory locations of said secondary memory for said host processing unit, said cache memory having a line size of 1 bytes, and each data unit having a size equal to the largest size that can be transferred to said bus master in parallel, said controller apparatus comprising circuitry which in a mode of operation, in response to a PCI-bus burst read transaction initiated by said bus master,

sequentially transfers data units to said bus master from said secondary memory according to said PCI-bus burst transaction, beginning at a starting memory location address in said secondary memory address space and continuing beyond at least first, second and third 1-byte boundaries of said secondary memory address space, each full 1-byte line of said transaction requiring at least 8 data unit transfers to said bus master, a plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries being transferred to said bus master at a constant rate, said constant rate being dependent upon the frequency of a PCI-bus clock provided to said bus master; and

during the transfer of the data units for each entire N'th 1-byte line according to said transaction, *initiates one and only one snoop access of said cache memory, said snoop access specifying the respective N+1'th 1-byte line* and being initiated early enough such that it can be sampled by said host processing

unit prior to completion of the transfer to said bus master of the last data unit in the respective N'th 1-byte line, said snoop accesses being sampled by said host processing unit in accordance with a host clock signal having a frequency that is at least twice said PCI-bus clock frequency.

The specification discloses:

According to the Pentium databooks, every data transfer to or from the memory address space which is cached by the L1 cache should be preceded by an inquire cycle.

(‘906 Patent at 5:20-30.) The patentees explained during prosecution that the claim limitation “*one and only one snoop access*” “is important because . . . a chipset that asserts multiple snoop accesses *for each cache line transferred* can tend to choke the processor and reduce the performance of the overall system.” (Dkt. No. 135, Ex. 10, 12/10/2001 Ghosh Decl., at p. 16 (emphasis added).)

On balance, Defendants’ proposal—which would exclude overlapping snoops rather than merely requiring that a next line is snooped only once—is at odds with the context of the claims, the specification, and the prosecution history. Defendants’ proposal of the word “exactly” is therefore rejected, and the Court adopts its prior construction in the *Apple* case. *Apple Markman* at 4.

The Court therefore hereby construes “**initiating one and only one snoop access of said cache memory, said snoop accesses each specifying the respective N+1'th 1-byte line**” to mean “**initiating one and only one next line inquiry.**”

**H. '291 Patent, Claims 88 and 89**

<b>'291 Patent, Claim 88</b>	
<b>Plaintiff's Proposed Construction</b>	<b>Defendants' Proposed Construction</b>
<p>[Plaintiff] denies that Claim 88 is in means plus function form, because, among other things, the claim specifically recites a structure (a “controller apparatus for a computer”)</p> <p>In the event that Claim 88 is construed pursuant to 35 U.S.C. §112, paragraph 6, [Plaintiff] contends that the disclosed structure which meets limitations 88.2 [(“sequentially transfers”)] and 88.3</p>	<p>Indefinite<sup>6</sup></p> <p>In the parties’ P.R. 4-5(d) Joint Claim Construction Chart, Defendants propose:</p> <p>This term is indefinite because the specification does not disclose sufficient structure for performing the claimed function. To the extent this term is not indefinite the structure associated with it should include at least the following structures:</p> <p>the partially disclosed circuitry inside the System Controller &amp; Integrated Peripherals Controller 116</p>

<sup>6</sup> In the parties’ pre-briefing claim chart, Defendants proposed (Dkt. No. 130, Ex. A, at 11-12):

To the extent this term is not indefinite the structure associated with it should include at least the following structures:

the partially disclosed circuitry inside the System Controller & Integrated Peripherals Controller 116 in Fig. 1 that moves data between the H-Bus and the PCI Bus performs the function of sequentially transferring data units to said bus master from said secondary memory according to said PCI-bus burst transaction, beginning at a starting memory location address in said secondary memory address space and continuing beyond at least first, second and third 1-byte boundaries of said secondary memory address space, each full 1-byte line of said transaction requiring at least 8 data unit transfers to said bus master, a plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries being transferred to said bus master at a constant rate.

The circuitry shown in Figs. 8 & 9 that generates the EADS# signal performs the part of the function of initiating one and only one snoop access of said cache memory, said snoop access specifying the respective N1’th [*sic*] 1-byte line and being initiated early enough such that it can be sampled by said host processing unit prior to completion of the transfer to said bus master of the last data unit in the respective N’th 1-byte line. The EADS# signal causes the function of determining if the secondary memory is cached in a modified state in the cache memory to be performed by an external cache controller whose circuitry [the patentee] did not disclose.

<p>[[“initiates”]] is the system controller and the integrated peripherals controller and known equivalents thereof</p>	<p>(SYSC/IPC) in Fig. 1 that moves data between the secondary memory and the PCI Bus which performs the function of sequentially transferring data units to said bus master from said secondary memory according to said PCI-bus burst transaction, beginning at a starting memory location address in said secondary memory address space and continuing beyond at least first, second and third 1-byte boundaries of said secondary memory address space, each full 1-byte line of said transaction requiring at least 8 data unit transfers to said bus master, a plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries being transferred to said bus master at a constant rate, the SYSC/IPC further comprising the following integrated OPTi, Inc. integrated circuit chips 82C557 (SYSC) and 82C558 (IPC) as described in OPTi, Inc., “Viper-M 82C556M/82C557M/82C558M, Data Book, Version 1.0” (April 1995), incorporated by reference herein, and further including an OPTi, Inc. 82C556 data buffer controller (DBC), also described in the above-incorporated data book, which includes some buffers not shown in FIG. 1, and equivalents thereof. ‘906 Patent, col. 9, ll. 30- 38; ‘291 Patent, col. 9, ll. 34-42.</p> <p>The circuitry shown in Figs. 8 &amp; 9 that generates the EADS# signal performs the part of the function of initiating one and only one snoop access of said cache memory, said snoop access specifying the respective N1’th [<i>sic</i>] 1-byte line and being initiated early enough such that it can be sampled by said host processing unit prior to completion of the transfer to said bus master of the last data unit in the respective N’th 1-byte line. The EADS# signal causes the function of determining if the secondary memory is cached in a modified state in the cache memory to be performed by an external cache controller whose circuitry [the patentee] did not disclose.<sup>7</sup></p>
---	--

<sup>7</sup> Plaintiff objects that Defendants have “abandoned and waived all arguments not addressed in its Markman briefing; namely, the alternative structures for Claim 26 of the ‘906 Patent and Claims 88 and 89 of the ‘291 Patent that [Defendants] attempt[] to disclose in this chart.” (Dkt. No. 142, Ex. A, at 5 n.2.) At the December 18, 2012 hearing, however, Plaintiff was agreeable to including the 82C556 (DBC), 82C557 (SYSC), and 82C558 (IPC) as part of the corresponding structure.

<b>‘291 Patent, Claim 89</b>	
<b>Plaintiff’s Proposed Construction</b>	<b>Defendants’ Proposed Construction</b>
Same as for Claim 88, above.	<p>Indefinite<sup>8</sup></p> <p>In the parties’ P.R. 4-5(d) Joint Claim Construction Chart, Defendants propose:</p> <p>This term is indefinite because the specification does not disclose sufficient structure for performing the claimed function.</p> <p>To the extent this term is not indefinite the structure associated with it should include at least the following structures: the partially disclosed circuitry inside the System Controller &amp; Integrated Peripherals Controller (SYSC/IPC) 116 in Fig. 1 that moves data between the secondary memory and the PCI Bus which performs the function of reading data from said secondary memory at a constant rate for said plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries, the SYSC/IPC further comprising the following integrated OPTi, Inc. integrated circuit chips 82C557 (SYSC) and 82C558 (IPC) as described in OPTi, Inc., “Viper-M 82C556M/82C557M/82C558M, Data Book, Version 1.0” (April 1995), incorporated by reference herein, and further including an OPTi, Inc. 82C556 data buffer controller (DBC), also described in the above-incorporated data book, which includes some buffers not shown in FIG. 1, and equivalents thereof. ‘906 Patent, col. 9, ll. 30-38; ‘291 Patent, col. 9, ll. 34-42</p>

---

<sup>8</sup> In the parties’ pre-briefing claim chart, Defendants proposed (Dkt. No. 130, Ex. A, at 14):

To the extent this term is not indefinite the structure associated with it should include at least the following structures: the partially disclosed circuitry inside the System Controller & Integrated Peripherals Controller 116 in Fig. 1 that moves data between the H-Bus and the PCI Bus performs the function of reading data from said secondary memory at a constant rate for said plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries

(Dkt. No. 135, at 24-26; Dkt. No. 139, at 27 & 29; Dkt. No. 142, Ex. A, at 14-16 & 19-20.)

Claim 89 depends from Claim 88.

(1) The Parties' Positions

Plaintiff submits that the *AMD Markman* agreed that these are not means plus function claims and that the parties agreed in *Apple* that Claim 88 is not a means-plus-function claim (Claim 89 was not at issue in *Apple*). (Dkt. No. 135, at 26-27.) Plaintiff argues that the “controller” recited in Claim 88 is sufficient structure and that “the mere fact that they [(Claims 88 and 89)] also claim functional attributes of the called out controllers does not make claims 88 and 89 susceptible to construction under Section 112 ¶6.” (*Id.*, at 28.) Plaintiff further notes that Claim 88 provides specific structure for the controller as a “controller apparatus for a computer system which includes a secondary memory having an address space, a bus master, a host processing unit and a cache memory which caches memory location of said secondary memory for said host processing unit, said cache memory having a line size of L bytes. . .” (*Id.*, at 27.)

Alternatively, if the Court finds that Claims 88 and 89 are in means-plus-function form, Plaintiff submits that the corresponding structure is “the SYSC/IPC, . . . described by the Patent as a ‘core logic chipset.’” (*Id.*, at 29 (quoting ‘906 Patent at 7:1-3; citing ‘906 Patent at Abstract, 6:7-27, 6:61-15:16, 20:55-29:59 & Figs. 1 & 8-12).)

Defendants respond that despite the absence of the word “means,” Claim 88 is subject to 35 U.S.C. § 112, ¶ 6. (Dkt. No. 139, at 27.) Defendants argue that unlike cases where “circuit” has been found to be sufficient structure to avoid 35 U.S.C. § 112, ¶ 6, “the multiple functions the circuitry performs are complex and constitute the entirety of the claimed invention.” (*Id.*) Defendants submit that “determining the structure of the circuitry that is actually performing the claimed functions requires examining the specification.” (*Id.*, at 28.)

Defendants also argue that “[a]llowing [Plaintiff] to claim functionally here [(without application of 35 U.S.C. § 112, ¶ 6)] would be particularly egregious because the specification of the ‘291 Patent (which has the same disclosure as the ‘906 Patent) does not disclose sufficient structure to perform the claimed function.” (*Id.*) Defendants explain that “the function of sequentially transferring data units requires copying data from the host bus to the PCI bus, and the circuitry for actually accomplishing that transfer is not disclosed as explained above in conjunction with the ‘sequentially transferring’ element of Claim 26 of the ‘906 Patent.” (*Id.*) Defendants further explain that “the ‘291 Patent does not disclose any structure for generating an address residing within the N+1’th L-byte line, a necessary prerequisite for being able to initiate a snoop access of that line.” (*Id.*, at 29.)

As to Claim 89, Defendants argue that “[w]ithout a disclosure of the data buffers and their associated control logic, a person of ordinary skill cannot identify sufficient structure to perform the data reading function claimed.” (*Id.*, at 30.) Defendants conclude that “in addition to being indefinite because it depends from a claim that is indefinite, Claim 89 is indefinite because the ‘291 Patent fails to disclose additional structure to perform the function added by this claim.” (*Id.*)

Defendants also submit the Declaration of Joe McAlexander in support of their indefiniteness arguments. (*Id.*, Ex. F, 11/27/2012 McAlexander Decl.)

Plaintiff replies to Defendants’ arguments on these claims together with Plaintiff’s reply to the means-plus-function terms, as discussed in subsection IV.E.(1), above.

## (2) Analysis

Claims 88 and 89 of the ‘291 Patent recite:

88. *Controller apparatus for a computer system which includes a secondary memory having an address space, a bus master, a host processing unit and a*

*cache memory* which caches memory locations of said secondary memory for said host processing unit, said cache memory having a line size of 1 bytes, and each data unit having a size equal to the largest size that can be transferred to said bus master in parallel, *said controller apparatus comprising circuitry* which in a mode of operation, in response to a PCI-bus burst read transaction initiated by said bus master,

sequentially transfers data units to said bus master from said secondary memory according to said PCI-bus burst transaction, beginning at a starting memory location address in said secondary memory address space and continuing beyond at least first, second and third 1-byte boundaries of said secondary memory address space, each full 1-byte line of said transaction requiring at least 8 data unit transfers to said bus master, a plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries being transferred to said bus master at a constant rate, said constant rate being dependent upon the frequency of a PCI-bus clock provided to said bus master; and

during the transfer of the data units for each entire N'th 1-byte line according to said transaction, initiates one and only one snoop access of said cache memory, said snoop access specifying the respective N+1'th 1-byte line and being initiated early enough such that it can be sampled by said host processing unit prior to completion of the transfer to said bus master of the last data unit in the respective N'th 1-byte line, said snoop accesses being sampled by said host processing unit in accordance with a host clock signal having a frequency that is at least twice said PCI-bus clock frequency.

89. Apparatus according to claim 88, wherein said circuitry further reads data from said secondary memory at a constant rate for said plurality of sequential data units bracketing at least said first, second and third 1-byte boundaries.

Claims 88 and 89 do not use the word “means.”

A claim limitation that actually uses the word “means” will invoke 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.

*CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1369 (Fed. Cir. 2002) (citations omitted).

The presumption that a limitation lacking the term “means” is not subject to section 112 ¶ 6 can be overcome if it is demonstrated that the claim term fails to recite sufficiently definite structure or else recites function without reciting sufficient structure for performing that function.

\* \* \*

The task of determining whether the limitation in question should be regarded as a means-plus-function limitation, like all claim construction issues, is a question of

law for the court, even though it is a question on which evidence from experts may be relevant.

*Lighting World, Inc. v. Birchwood Lighting, Inc.*, 382 F.3d 1354, 1358 (Fed. Cir. 2004) (citations and internal quotation marks omitted).

On balance, the recitals regarding the “controller” and “host processing unit” in Claim 88 are sufficient to avoid application of 35 U.S.C. § 112, ¶ 6 to either Claim 88 or Claim 89. *See Telcordia*, 612 F.3d at 1376-77 (holding that “controller” was sufficient disclosure because “[t]he record shows that an ordinary artisan would have recognized the controller as an electronic device with a known structure”). In other words, Defendants have failed to overcome the presumption that in the absence of the word “means,” 35 U.S.C. § 112, ¶ 6 does not apply. This is the same conclusion that the Court reached in the *AMD* case. *AMD Markman* at 8-11.

In sum, Claims 88 and 89 do *not* contain means-plus-function limitations and are *not* invalid as indefinite. The parties present no other dispute regarding Claims 88 and 89, so the Court does not further construe those claims.

## V. 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.

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 21st day of December, 2012.**

  
\_\_\_\_\_  
RODNEY GILSTRAP  
UNITED STATES DISTRICT JUDGE