United States District Court,
N.D. California.

INTERGRAPH HARDWARE TECHNOLOGIES COMPANY,
Plaintiff.
v.
TOSHIBA CORPORATION, et al,
Defendants.

No. C 06-04018 MHP

Aug. 2, 2007.

Background: Assignee of patents related to management of memory and communication among elements of a computer system sued for infringement.

Holdings: The District Court, Marilyn Hall Patel, J., held that:
(1) term "system bus" meant a bus having multiple masters and through which the primary memory communicated with each system element that accessed primary memory;
(2) phrase "system bus monitoring means" was not a means-plus-function term;
(3) term "cache data storage modes" meant strategies for writing data to cache and to primary memory;
(4) the function of term "system tag means" was passively retaining a system tag which indicated one of a plurality of cache data storage modes;
(5) the function of term "primary memory interface means" was coupling data between said primary memory and said cache memory, and selectively transferring data between said primary memory and said cache memory in response to a miss signal; and
(6) function of term "system bus monitoring means" was monitoring reads and/or writes from an input/output (I/O) data processing device over the system bus to an addressed device.

Claims construed.

4,899,275, 4,933,835, 5,091,846. Construed.

Jacob S. Zimmerman, William H. Manning, Brad P. Engdahl, David J. Wallace-Jackson, Robins, Kaplan, Miller & Ciresi, L.L.P., Melissa Wendland, Minneapolis, MN, John P. Bovich, William R. Overend, Reed Smith, LLP, San Francisco, CA, Eric Sean Jackson, Robins, Kaplan, Miller & Ciresi L.L.P., Washington, DC, for Plaintiff.

John J. Feldhaus, Mary M. Calkins, Pavan K. Agarwal, Foley & Lardner, LLP, Washington, DC, Russell J. Barron, Foley & Lardner, Milwaukee, WI, Stephen Michael Lobbin, Foley & Lardner, LLP, Palo Alto, CA, for Defendants.

Claim Construction Memorandum and Order for United States Patent Nos. 4,899,275, 4,933,835**, and** 5,091,846

MARILYN HALL PATEL**, District Judge.**

Plaintiff Intergraph Hardware Technologies Company ("Intergraph") brought this action against defendants Toshiba Corporation, Toshiba America Information Systems, Inc., Toshiba America Medical Systems, Inc., Toshiba America Business Solutions, Inc., Toshiba TEC America Retail Information Systems, Inc. (collectively "Toshiba") and other defendants, asserting claims for infringement of U.S. Patent No. 4,899,275 ("the '275 Patent"), U.S. Patent No. 4,933,835 ("the '835 Patent"), and U.S. Patent No. 5,091,846 ("the '846 Patent"). Now before the court are Intergraph and Toshiba's claim construction briefs, filed pursuant to Patent Local Rule 4-5. Having considered the parties' arguments and submissions, and for the reasons set forth below, the court construes the disputed terms as follows.

BACKGROUND FN1

FN1. For convenience, the court will use column and line numbers from the specifications of the various patents for citations in this order.

Intergraph is the assignee of the three patents at issue in this lawsuit. The "overall invention" covered by the three patents addresses the operation of a computer system, in particular the management of memory and communication among the elements of such a system. The problems addressed by the patents are (1) performance limitations arising from disparities in speeds between a fast microprocessor and a slower main memory, and (2) the need to ensure data consistency between cache memory and main memory. The invention solved these problems by providing for (1) multiple data storage modes on a per-page basis specified by system tags, and (2) a cache controller that monitors the system bus and maintains data consistency.

I. ***Overview of Computer System Components and Operation***

A review of the invention requires a brief overview of the relevant computer technology. The primary components of a typical computer system are the "microprocessor," which includes the CPU, cache memory and cache controller; "I/O (input/output) devices" such as hard drives and external disk drives; the "motherboard," which serves as the primary circuit board holding the electronic components of the computer, the "system bus," containing wires that connect the various components of the system, and the "main memory," which electronically stores data that the CPU is currently processing.

The CPU communicates with the main memory through "read" and "write" commands transmitted over the system bus. A "read" command retrieves data from the main memory without altering the data in the memory. A "write" command sends data to the main memory, altering the data therein. Each of these operations is relatively slow compared to other CPU operations. In order to compensate for this slow speed, computer systems include a "cache memory" located within the microprocessor. The cache memory is smaller than the main memory and communicates with the CPU much faster than the main memory. Due to its smaller size, only a portion of the data in the main memory is stored in the cache memory at any given time. By allowing the CPU to communicate principally with the cache memory rather than the main memory, the overall speed and performance of the computer system increases dramatically.

II. ***Caching Strategies***

In order to maximize the efficiency of cache and main memory usage, various "caching strategies" or "cache data storage modes" have been developed. For the purposes of the invention at issue here, the two main caching strategies are "write-through" and "copyback." In write-through, all CPU writes are sent both to the cache memory and the main memory. In copyback, the CPU writes to the cache memory only. The updated

data in the cache memory is later written to the main memory only if necessary. At the time of invention, each of these strategies was well-known in the art and had specific advantages and disadvantages depending on circumstances.

One of the problems addressed by the inventions was the fact that computer systems were designed with only a single caching strategy. For example, a computer system designed to implement write-through could not perform copyback, and vice versa. Intergraph solved this problem by providing for the selection of multiple cache data storage modes on a "per-page" basis. The term "page" refers to a set of data stored in the main memory. The invention sets the caching strategy for a particular page through the use of control bits known as "system tags," stored in the "cache controller." The cache controller is located in the microprocessor along with the CPU and cache memory. This allows the caching strategy to be determined as each page of data is loaded from the main memory into the cache memory. The use of multiple caching strategies within a single computer system increases the efficiency and flexibility of the system, ensuring that the system uses the most efficient caching strategy for each particular page of data.

III. ***Data Consistency***

A second issue addressed by the invention is the need to maintain data consistency between the cache memory and the main memory. Data inconsistency may arise, for example, where an I/O device such as the hard drive writes to the main memory but not the cache memory. The result is different values in the two memories, which can cause the system to malfunction if the CPU attempts to access the old data in the cache memory. The invention maintains data consistency through the use of a "bus monitor," a component of the cache controller which monitors data transfers on the system bus. When the main memory is accessed by a device other than the CPU via the system bus, the bus monitor detects the data transfer and determines what steps, if any, need to be taken to maintain data consistency. For example, if data at a particular location in the main memory is altered by an I/O device, the bus monitor will invalidate the data at the corresponding location in the cache memory. Typically, the bus monitor invalidates the entire line of data on which the outdated value is stored, rather than invalidating only the value itself.

Similarly, if the CPU alters data at a particular location in the cache memory in copyback mode, a problem may arise if an I/O device attempts to read that data from its corresponding location in the main memory. The bus monitor solves this problem by supplying the updated data from the cache memory to the I/O device, so that the I/O device does not obtain outdated data from the main memory.

LEGAL STANDARD

I. ***Claim Construction***

Under Markman v. Westview Instruments, Inc., 517 U.S. 370, 389-90, 116 S.Ct. 1384, 134 L.Ed.2d 577 (1996), the court construes the scope and meaning of disputed patent claims as a matter of law. The first step of this analysis requires the court to consider the words of the claims. Teleflex, Inc. v. Ficosa N. Am. Corp., 299 F.3d 1313, 1324 (Fed.Cir.2002). According to the Federal Circuit, the court must "indulge a 'heavy presumption' that a claim term carries its ordinary and customary meaning." CCS Fitness, Inc. v. Brunswick Corp., 288 F.3d 1359, 1366 (Fed.Cir.2002). To determine the ordinary meaning of a disputed term, the court may review a variety of sources including the claims themselves, other intrinsic evidence such as the written description and prosecution history, and dictionaries and treatises. Teleflex, 299 F.3d at 1325. The court must conduct this inquiry not from the perspective of a lay observer, but rather "from the standpoint of a person of ordinary skill in the relevant art." Id. (citing Zelinski v. Brunswick Corp., 185 F.3d 1311, 1316 (Fed.Cir.1999)).

[1] [2] [3] Among the sources of intrinsic evidence, the specification is "the single best guide to the meaning of a disputed term." Vitronics Corp. v. Conceptronic, Inc., 90 F.3d 1576, 1582 (Fed.Cir.1996). By expressly

defining terms in the specification, an inventor may "cho[ose] to be his or her own lexicographer," thereby limiting the meaning of the disputed term to the definition provided in the specification. Johnson Worldwide Assocs., Inc. v. Zebco Corp., 175 F.3d 985, 990 (Fed.Cir.1999). In addition, "[e]ven when guidance is not provided in explicit definitional format, the specification may define claim terms 'by implication' such that the meaning may be 'found in or ascertained by a reading of the patent documents.' " Irdeto Access, Inc. v. Echostar Satellite Corp., 383 F.3d 1295, 1300 (Fed.Cir.2004) (quoting Bell Atl. Network Servs., Inc. v. Covad Commc'ns Group, Inc., 262 F.3d 1258, 1268 (Fed.Cir.2001)). "The specification may assist in resolving ambiguity where the ordinary and accustomed meaning of the words used in the claims lack sufficient clarity to permit the scope of the claim to be ascertained from the words alone." Teleflex, 299 F.3d at 1325. At the same time, the Federal Circuit has cautioned that the written description "should never trump the clear meaning of the claim terms." Comark Comms., Inc. v. Harris Corp., 156 F.3d 1182, 1187 (Fed.Cir.1998) (citations omitted); *see also* Tate Access Floors, Inc. v. Maxcess Techs., Inc., 222 F.3d 958, 966 (Fed.Cir.2000) ("Although claims must be read in light of the specification of which they are part, ... it is improper to read limitations from the written description into a claim ....").

[4] Likewise, the prosecution history may demonstrate that the patentee intended to deviate from a term's ordinary and accustomed meaning. Teleflex, 299 F.3d at 1326. "Arguments and amendments made during the prosecution of a patent application and other aspects of the prosecution history, as well as the specification and other claims, must be examined to determine the meaning of terms in the claims." Southwall Techs., Inc. v. Cardinal IG Co., 54 F.3d 1570, 1576 (Fed.Cir.1995), *cert. denied,* 516 U.S. 987, 116 S.Ct. 515, 133 L.Ed.2d 424 (1995). "In particular, 'the prosecution history (or file wrapper) limits the interpretation of claims so as to exclude any interpretation that may have been disclaimed or disavowed during prosecution in order to obtain claim allowance.' " Teleflex, 299 F.3d at 1326 (quoting Standard Oil Co. v. American Cyanamid Co., 774 F.2d 448, 452 (Fed.Cir.1985)).

[5] [6] [7] Dictionary definitions and other objective reference materials available at the time that the patent was issued may also provide evidence of the ordinary meaning of a claim. Phillips v. AWH Corp., 415 F.3d 1303, 1322 (Fed.Cir.2005) (en banc). A dictionary "has the value of being an unbiased source, accessible to the public in advance of litigation." Phillips, 415 F.3d at 1322 (internal quotation omitted). Thus, district courts "are free to consult such resources at any time in order to better understand the underlying technology and may also rely on dictionary definitions when construing claim terms, so long as the dictionary definition does not contradict any definition found in or ascertained by a reading of the patent documents." Vitronics, 90 F.3d at 1584 n. 6. A court should be cautious, however, not to place too much reliance on dictionaries, as the resulting construction may be too broad. Phillips, 415 F.3d at 1321.

[8] Federal Circuit decisions take a less favorable view of other forms of extrinsic evidence, such as expert testimony and prior art not cited in the specification or the prosecution history, noting that "claims should preferably be interpreted without recourse to extrinsic evidence such as expert testimony, other than perhaps dictionaries or reference books, and that expert testimony should be received only for the purpose of educating the judge." EMI Group N. Am., Inc. v. Intel Corp., 157 F.3d 887, 892 (Fed.Cir.1998), *cert. denied,* 526 U.S. 1112, 119 S.Ct. 1756, 143 L.Ed.2d 788 (1999). Although "extrinsic evidence in general, and expert testimony in particular, may be used ... to help the court come to a proper understanding of the claims[,] it may not be used to vary or contradict the claim language .... Indeed, where the patent documents are unambiguous, expert testimony regarding the meaning of a claim is entitled to no weight." Vitronics, 90 F.3d at 1584.

The Federal Circuit recently revisited the basic approach to claim construction in Phillips. Although Phillips consists largely of an affirmation of ten years of claim construction jurisprudence, it provides at least two pieces of additional guidance. First, the Federal Circuit rejected a line of cases suggesting that claim interpretation must begin with a dictionary definition of the disputed terms. Phillips, 415 F.3d at 1320-21. Second, the Federal Circuit emphasized that claim terms must be interpreted in light of their context,

especially the language used in other claims and the specification. Id. at 1321. Taken as a whole, Phillips appears to signal a small retreat from formalism and bright-line rules in claim construction. As a result, the court will focus primarily on the intrinsic record before it. Cases cited by the parties in support of fixed "rules" of claim construction will accordingly be given somewhat less weight.

II. ***Means-Plus-Function Claims***

[9] Section 112(6) of the Patent Act, 35 U.S.C. section 112 para. 6 ("Paragraph 6" or "Section 112(6)"), provides that an element in a claim "may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof." Claims expressed in this way "shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof." Id. By allowing for a patentee to recite "means-plus-function" claim elements, section 112(6) permits the inventor to describe an element of his or her invention by the result accomplished or the function served rather than describing the item or element to be used. Warner-Jenkinson Co. v. Hilton Davis Chem. Co., 520 U.S. 17, 28, 117 S.Ct. 1040, 137 L.Ed.2d 146 (1997). Although the claim limitation need not recite the structure, material, or acts that comprise the means for carrying out the claimed function, the applicant must describe in the patent specification some structure which performs that function. 35 U.S.C. s. 112 para. 6; Odetics, Inc. v. Storage Tech. Corp., 185 F.3d 1259, 1266-67 (Fed.Cir.1999).

[10] [11] [12] Where an element is expressed as a "means" to perform a particular function,a presumption arises that the claim element should be construed as a means-plus-function claim under section 112(6). Al-Site Corp. v. VSI Int'l, Inc., 174 F.3d 1308, 1318 (Fed.Cir.1999). Conversely, a claim term that does not use the words "means" or "step for" is presumptively not governed by section 112(6). CCS Fitness, 288 F.3d at 1369. This presumption "is a strong one that is not readily overcome." Lighting World, Inc. v. Birchwood Lighting, Inc., 382 F.3d 1354, 1358 (Fed.Cir.2004). However, even in the absence of express means-plus-function language, a claim may be construed to include a means-plus-function element if the proponent of the means-plus-function construction demonstrates that "the claim term fails to 'recite sufficiently definite structure for performing that function.' " CCS Fitness, 288 F.3d at 1369 (quoting Watts v. XL Sys., Inc., 232 F.3d 877, 880 (Fed.Cir.2000)). In making this determination, the court must assess whether the "term, as a name for structure, has a reasonably well understood meaning in the art." Greenberg v. Ethicon Endo-Surgery, Inc., 91 F.3d 1580, 1583 (Fed.Cir.1996).

DISCUSSION

The following chart summarizes the court's construction of the disputed terms. The full analysis supporting each construction is below.

Term	**Construction**
"system bus"	a bus having multiple masters and through which the primary memory communicates with each system element that accesses primary memory
"system bus means"	Not subject to Section 112(6); construed as "system bus"
"access indicating means"	the conductors of the system bus which transmit the Active Cycle (AC) and MS<4> signals
"first data processing element [coupled to the system bus for processing data from the main memory]"	Not construed.
"cache data storage modes"	strategies for writing data to cache and to primary memory
"system tag means"	function: storing (in this case, passively retaining) a system tag which indicates one of a plurality of cache data storage modes. structure: a storage location in the cache memory management means

"translation means"	structure: array 352/354 of TLB 350 of Fig. 9
"primary memory interface means"	function: (a) coupling data between said primary memory and said cache memory; (b) and selectively transferring data between said primary memory and said cache memory in response to a miss signal
	structure: (a) SIR 260 and SOR 250 of '846 Patent Fig. 8, also labeled as 643 and 644 of Fig. 21; (b) microengine 650
"primary memory means" in Claims 1 & 4-14	Not subject to Section 112(6), construed as "a memory for storing data"
"primary memory means" in Claim 15	Subject to Section 112(6). function: coupling data from the selected storage location to the system bus when the second cache inquiry signal is provided and the first ready signal is absent
	structure: the portion of the Drivers/Receivers and Interface circuitry depicted in box 140 of Fig. 1 that performs the protocols set forth in '835 13:15-16, 41-43, 46-48
"data consistency means"	structure: system bus control 840 and microengine 650 programmed to execute a set of instructions to perform an internal cycle to determine whether to purge the cache of its data and supply data from an addressed location or to supply its data instead of supplying data from an addressed location.
"system bus monitoring means"	function: "monitoring reads and/or writes from an I/O data processing device over the system bus to an addressed device"

I. ***System Bus***

The term "system bus" appears in independent Claim 1 of the '846 Patent. The parties agree that the term "system bus" had an ordinary and customary meaning at the time of invention, but disagree as to the meaning. Intergraph's original proposed construction of "system bus" is "one or more conductors in a computer along each of which information is transmitted from any of a plurality of sources to any of a plurality of destinations and to which the first data processing element, the cache memory management means, and the primary memory, as well as potentially numerous other system elements including multiple I/O processors, can be coupled." Toshiba's proposed construction is "a bus having multiple masters and through which the primary memory communicates with each system element that accesses primary memory."

As an initial matter, Toshiba claims that Intergraph's language is unclear and unworkable, and should be rejected on that basis. Toshiba claims that Intergraph's construction contains numerous terms that would themselves require interpretation, such as "source," "destination," "potentially numerous" and "can be coupled," creating indefiniteness problems and other uncertainties. In particular, the conditional phrase "can be coupled" adds an unacceptable degree of uncertainty and indefiniteness to the construction of the patent. In response to this argument, Intergraph now states that it is willing to accept a modified version of Toshiba's construction: "a bus having multiple masters and through which the primary memory communicates with the cache memory." This is Toshiba's proposed construction replacing the phrase "each system element that accesses primary memory" with "the cache memory."

[13] The main dispute, therefore, is whether the term "system bus" must be construed to require that every component communicates with the main memory via the system bus. Under Toshiba's proposed construction, all components that communicate with the main memory must be able to do so via the system bus, though other communication channels may be available. Under Intergraph's proposed construction, only the cache memory must be able to communicate with the main memory via the system bus. The other components need not be able to use the system bus to communicate with the main memory.

In support of its proposed construction, Intergraph raises only two arguments, one based on claim language superfluity and the other based on the prosecution history. First, Intergraph points to Claim 1 of the '846 patent, which includes the element "system bus monitoring means for monitoring I/O requests over the system bus." According to Intergraph, because Toshiba's construction would require all I/O requests to go over the system bus, the phrase "over the system bus" in Claim 1 would be superfluous under Toshiba's construction. In response, Toshiba claims that the phrase "over the system bus" specifies the location where the monitoring takes place, not where the communication occurs.

Additionally, Intergraph argues that the prosecution history does not indicate that Intergraph adopted Toshiba's interpretation of the system bus. In particular, the Examiner initially rejected claims 16-40 of the '835 patent application, claims that recited a "system bus," as obvious based in part on U.S. Patent No. 4,701,844. Pl's Exh. M at 5. The '844 Patent discloses a system bus involved in communications between the CPU and the "Main Memory," but also discloses other paths for communications with the Main Memory. Pl's Exh. F. This, Intergraph asserts, demonstrates that Intergraph's understanding of the term "system bus" during prosecution did not require that all communications with the primary memory take place over the system bus. As Toshiba points out, however, nothing in the prosecution history shows that Intergraph ever advanced the interpretation set forth in the '844 Patent during prosecution. Furthermore, at the time the PTO rejected the claim based on the '844 Patent, the PTO also cited JP Application No. 58-58666, which teaches a system bus as shown in the Intergraph application. Def's Exh. M at 5. Taken together, these two prior art references show not only that Intergraph had no reason to adopt one reference over the other, but also that both types of system bus were present and understood in the prior art. In light of this ambiguity, the court looks to the patent specification for clarification.

Turning to the specification, Toshiba claims that the '846 Patent and the ' 835 Patent each require the system bus 141 to be the central system artery that provides the pathway for the system elements to access the main memory. ' 846 Patent at 4:22-25, 4:40-42, 15:17-20; '835 Patent at 1:45-62, 2:1-4. Intergraph's expert, Dr. Wolfe, confirmed this conclusion, stating that the specification has no technical disclosure of alternatives for the system bus. Def's Exh. 2, Wolfe Dep. at 40:21-41:10, 53:23-57:2, 58:3-11, 64:17-20. The diagram denoted Figure 1, which is common to the two patents, clearly shows that all communications between any component and the main memory pass through the system bus.

Additionally, the specifications of the two patents differentiate between different buses, including the I/O bus, the cache/processor instruction bus, and the cache/processor data bus. The claims of the '846 Patent also identify different types of buses, in particular the "processor bus" as opposed to the "system bus." '846 Patent, Claim 17. Toshiba claims that Intergraph's proposed construction would eliminate this distinction between the system bus and other buses which is explicitly drawn in the claims and specification. Dr. Wolfe was unable to identify any language in Intergraph's proposed construction of "system bus" that would not also apply to the processor buses, and acknowledged that the proposed construction does not differentiate between buses based on purpose. Def's Exh. 2, Wolfe Dep. at 101:16-105:17, 111:24-112:18, 122:23-123:2.

Toshiba also points to the context of the claim language in support of its interpretation. Claim 1 of the '846 Patent recites "a system bus coupled to the primary memory," "a first data processing element *coupled to the system bus* for processing data from the primary memory," and "cache memory management means, *coupled to ... said system bus*" including "a cache memory for storing data from said primary memory" (emphasis added). Toshiba asserts that this language comports with the system bus being the sole pathway to the main memory in light of the written description. Additionally, Claims 1 and 9 each recite "data consistency means" that include a system bus monitoring means to monitor I/O requests "over the system bus." Toshiba claims that because the only monitored I/O requests affecting this data consistency are directed to the main memory, they must traverse the system bus.

Toshiba additionally claims that Intergraph's statement of the invention confirmsToshiba's construction of "system bus." The stated aim of the invention is to ensure data consistency between the cache and the main memory, and this is achieved by a cache controller that monitors the system bus and maintains data consistency. Toshiba claims that monitoring "system bus" transactions would not ensure data consistency if data could flow into and out of the main memory through channels other than the system bus. Def's Exh. 1 para. 64. As discussed above, the data consistency function of the invention is important where components other than the cache memory are communicating with the main memory. For example, if an I/O device alters data in the main memory, the system bus monitoring means detects the change and reacts appropriately to ensure data consistency between the cache memory and the main memory. Because the system bus monitoring means monitors the system bus only, if I/O devices could alter data in the main memory through channels other than the system bus, the system bus monitoring means would fail to detect the data modification, and the data consistency operation would not be triggered. Under such circumstances, one of the key aims of the invention-maintaining data consistency between the cache memory and the main memory-would be thwarted.

In response, Intergraph argues that its data consistency mechanism can operate in a computer containing a separate pathway other than the system bus for I/O to access primary memory. According to Intergraph, in such a configuration data consistency could be maintained where an I/O request is sent over the system bus and monitored by the system bus monitoring means. This argument appears wholly irrational. Intergraph appears to confirm that data consistency is attained only where I/O requests are sent over the system bus, which supports Toshiba's argument. In other words, if the goal of the invention is data consistency and data consistency is only maintained for communications passing through the system bus, then the only reasonable interpretation is that all communications must pass through the system bus.

Toshiba further claims that Intergraph's construction would render the patents invalid as non-enabling. In particular, the fact that the specification teaches only a system whereby all main memory communications to go over the system bus would render invalid a claim that covers main memory communications that do not involve the system bus. While this type of invalidity argument is inappropriate at the claim construction stage, the court is mindful of the tenet that claims should be construed to avoid invalidity.

Finally, Toshiba claims that Intergraph knows that its construction is wrong, pointing to previous statements by Intergraph's CEO and an Intergraph expert in a prior litigation indicating that all communication involving the main memory passed through the system bus. Defs' Exhs. 5 & 8. Intergraph responds that these references are irrelevant, and that they express opinions not relied upon by Intergraph. Intergraph further claims that the CEO opinions cited by Toshiba referred to a configuration in which I/O data was not required to go over the system bus to access memory, and that the expert report did not address whether all communication with the main memory took place over the system bus. Because it is not clear that the specific issue at hand here-whether all communications with the primary memory in the claimed system must pass over the system bus-were at issue in the contexts of these previous statements, the court affords this evidence little weight.

Focusing, as the court should, on the intrinsic evidence available, Toshiba's position that the system bus must carry all communications with the primary memoryis correct. The cited references support both readings, but the patent itself uniformly treats the system bus as a component carrying all main memory communications. Intergraph does not seem to dispute the treatment of the system bus in the patent itself, or point to anything in the specification suggesting that communications with the main memory can take place via paths not involving the system bus. Furthermore, Intergraph's superfluity argument is unconvincing. Even if the court were to reject Toshiba's argument regarding the meaning of the phrase at issue, the inclusion of this phrase could easily be interpreted as adding clarity to the claim rather than specifying a particular limitation. In any case, superfluity is not so rigid a doctrine as to require the court to disregard the unambiguous teachings of the specification itself. Finally, Toshiba is correct that Intergraph's broad

construction would confuse the system bus with other, distinct buses disclosed by the patents.

In sum the, court adopts Toshiba's construction of "system bus": "a bus having multiple masters and through which the primary memory communicates with each system element that accesses primary memory."

II. ***System Bus Means***

[14] The term "system bus means" appears in Claim 1 of the '835 Patent. The clause at issue is "system bus means, coupled to the primary memory, for communicating data with the primary memory means." The parties disagree as to whether the term is a means-plus-function element.

[15] The use of the word "means" raises a strong presumption that Section 112(6) applies. Sage Prods., Inc. v. Devon Indus., Inc., 126 F.3d 1420, 1427 (Fed.Cir.1997). However, "where a claim recites a function, but then goes on to elaborate sufficient structure, material, or acts within the claim itself to perform entirely the recited function, the claim is not in means-plus-function format." Id. at 1427.

During prosecution, Intergraph canceled certain claims that recited "system bus" and replaced them with claims that recited "system bus means," Pl's Exh. N at 2, lending further support to the contention that "system bus means" is not coterminous with "system bus." In response, Toshiba cites a patent attorney's opinion from an earlier litigation, in which the attorney opined that Section 112(6) did not apply to the term "system bus means." Defs' Exh. 11 at 8 & 15.

In light of the above discussion and construction of the term "system bus," the term "system bus monitoring means" recites sufficient structure to perform the function of "communicating data with the primary memory means." Accordingly, "system bus means" is not a means-plus function term.

III. ***Access Indicating Means***

The term "access indicating means" appears in Claim 2 of the '835 Patent. The parties agree that this term is a means-plus-function term, and that the function is "for indicating when the primary memory means is being accessed." The parties disagree as to the corresponding structure. Intergraph's proposed structure is "the conductors of the system bus which transmit the ActiveCycle (AC) and MS<4> signals." Toshiba's structure is "system bus lines MS and AC that indicate the timing and duration of each primary memory access." The dispute therefore centers on the corresponding MS bits. Intergraph claims that only the MS<4> bit corresponds to the function, while Toshiba identifies all MS bits.

The '835 Patent identifies five MS bits on the system bus. '835 Patent at 10:35-69. Intergraph claims that Toshiba's expert, Dr. Colwell, admitted that only MS<4> is necessary to indicate whether or not a request is directed to primary memory means. Pl's Exh. E, Colwell Dep. at 260:14-20. Toshiba argues that Intergraph's proposed construction, which only identifies the signals by name, is not sufficiently specific to identify them based on how the signals work. This, according to Toshiba, provides no rational way of determining whether an accused device has the same or equivalent corresponding structure. In particular, Toshiba claims that its construction clarifies the corresponding structure by stating the meaning of the signal lines MS and AC as indicating the timing and duration of each primary memory access. Intergraph rejects this approach, disputing that the AC indicates "when" a system bus cycle is active and instead arguing that the AC indicates "that" a bus cycle is active. *See* '835 Patent at 11:15-17. Intergraph's argument in this regard is contrary to the claim language, however, which states that the access indicating means is "for indicating *when* the primary memory means is being accessed." '835 Patent, Claim 2.

As used in the specification, only the MS<4> bit indicates when the main memory is being accessed. The value of MS<4> is 0 if the main memory is being accessed, and 1 if memory-mapped I/O space or boot

loader space is being accessed. Id. at 10:42-55. Toshiba's generalized formulation of the structure is helpful only if read in conjunction with an additional functional element: "indicat[ing] the timing and duration of each primary memory access." Having separately defined the corresponding function, it would be improper to define the corresponding structure in terms of an additional function. Because Intergraph's proposed structure clearly identifies a structural element that performs the corresponding function, the court adopts Intergraph's structure: "the conductors of the system bus which transmit the ActiveCycle (AC) and MS<4> signals."

IV. ***First Data Processing Element***

Claim 1 of the '846 Patent and Claim 1 of the '835 Patent each recite a "first data processing element, coupled to the system bus [means], for processing data from primary memory." Intergraph claims that this does not require construction, but proposes the construction: "a device that is coupled to the system bus and that processes data stored in primary memory." Toshiba's proposed construction is "a data processing element that processes data communicated from the primary memory through the system bus." Additionally, Toshiba claims that the term to be construed is the broader term "first data processing element coupled to the system bus for processing data from the main memory." In any case, the dispute once again boils down to whether all communication from the primary memory must be conducted over the system bus. Because this issue has been resolved via the construction of "system bus," no construction of "first data processing element" or "first data processing element coupled to the system bus for processing data from the main memory" is necessary.

V. ***Cache Data Storage Modes***

[16] The term "cache data storage modes" appears in independent claims 1 and 9 of the '846 Patent. The parties agree that the cache data storage modes are "strategies for writing data to cache and to primary memory," though Toshiba proposes adding "including at least copyback and write-through cacheable strategies" as a second clause. Intergraph claims that Toshiba's approach improperly imports limitations from the preferred embodiment into the claims. Additionally, Intergraph raises a claim differentiation argument, citing dependent claim 4 which recites specific data storage modes of "copyback" and "write-through." Toshiba rejects this argument, stating that claim differentiation is "a guide, not a rigid rule," Curtiss-Wright Flow Control Corp. v. Velan, Inc., 438 F.3d 1374, 1381 (Fed.Cir.2006), and that any differentiation arguments are trumped by the intrinsic record limiting the claims to copyback and write-through. Furthermore, Claim 4 contains additional limitations beyond specifying copyback and write-through.

Toshiba further argues that the invention disclosed in the '846 Patent specifically stated that it was superior to the prior art because it implements "all three caching strategies, cacheable non cacheable, write-through, and copyback, to be utilized in the same machine, and which allows for selection of any one of the three strategies under software control using software control bits." '846 Patent at 1:62-67, 3:1-4. Toshiba claims that no cacheable modes are described or suggested in the patent other than write-through and copyback.

The parties also raise substantial arguments regarding two pieces of prior art cited by the PTO during the prosecution of the '846 Patent. In rejecting the claims, the PTO relied on the fact that U.S. Patent No. 4,433,374 ("the Hanson reference") disclosed "three modes of written data into the cache memory": "non-write through mode [normal write command]," "write through mode [Store-Thru command]" and "non-cacheable mode [Dispersed Write]." Pl's Exh. O at 4 (brackets in original). Intergraph claims that the PTO's rejection was based on the *number* of modes disclosed in the Hanson reference rather than the types of modes disclosed, because the Hanson reference did not disclose a copyback mode. Intergraph attempted to overcome this rejection by amending its claims to include limitations related to data consistency rather than asserting the failure of the Hanson reference to disclose a copyback mode. Pl's Exh. P at 11-12.

In response, Toshiba cites statements from Dr. Wolfe claiming that, in the "normal write" mode disclosed in the Hanson reference and cited by the PTO, "data was held in a disk-drive cache until the disk drive was not busy, and was then written to the disk." Wolfe Dec., Exh. A para. 45. Despite the fact that Dr. Wolfe's very next sentence is "Hanson does not disclose a copyback mode," Toshiba claims that Dr. Wolfe's explanation of the "normal write" is consistent with the copyback mode recited in Claim 4 of the '846 Patent. Aside from drawing an obtuse parallel between Dr. Wolfe's explanation and Claim 4's recitation of immediately storing data only in the cache memory, Toshiba provides no explanation as to why Dr. Wolfe allegedly contradicted himself in his report. Accordingly, Toshiba's argument in this regard is unconvincing.

Toshiba additionally relies on a second piece of prior art asserted by the PTO. In attempting to distinguish from U.S. Patent No. 4,669,043 ("the Kaplinsky reference"), Intergraph stated that its claimed invention accomplished combining write-through and copyback caching strategies whereas the prior art disclosed only systems using one strategy or the other. Defs' Exh. Z at 11, 13-14. Intergraph distinguished the Kaplinsky reference based on the inability to add additional modes to the reference disclosed therein, not the absence of a particular mode. Id. at 13. As with the Hanson reference, the salient point regarding the Kaplinsky reference was numerosity rather than type. A system's ability to select a caching strategy as new requirements arise, rather than committing to a particular strategy in advance and therefore being unable to adapt, increasesthe flexibility and functionality of the system. Id. at 11.

Intergraph further claims that the PTO's citation of Hanson, which did disclose a plurality of modes, was in response to Intergraph's distinction of Kaplinsky. The Examiner stated that "when Hanson et al's teachings of different cache data storage modes is incorporated into Kaplinsky's system, it would have been obvious to one of ordinary skilled [sic] in the art that the different cache data storage modes can be controlled by the system attributes stored in the descriptor table." Pl's Exh. O at 4. This further supports the contention that the ability to select from among different available strategies was the salient issue in this exchange.

While the specification cites specific examples known at the time to demonstrate the invention's multiple cache storage modes, the claims themselves do not limit the invention to the delineated modes. The intrinsic record supports Intergraph's contention that the PTO was chiefly concerned with the number, rather than type, of modes. Furthermore, the claim language itself refers to "a plurality of cache data storage modes." '846 Patent at Claims 1 & 9. The term "cache data storage modes" is broad, and to limit the modes to those specifically identified in the specification would be to improperly import claim limitations where none exist. Accordingly, the court adopts Intergraph's construction of "cache data storage modes": "strategies for writing data to cache and to primary memory."

VI. *System Tag Means*

The term "system tag means" appears in independent claims 1 and 9 of the ' 846 Patent. The parties agree that Section 112(6) applies to this term, but disagree as to both the function and the structure.

A. *Function*

[17] Intergraph's proposed function is "storing (in this case, passively retaining) a system tag which indicates one of a plurality of cache data storage modes." Toshiba's proposed function is "storing a system tag that indicates which one of a plurality of selectable cache data storage modes is effected." The dispute therefore centers on Toshiba's use of the terms "selectable" and "will be effected." Intergraph claims that Toshiba's construction improperly "adopt[s] a function different from that explicitly recited in the claim." JVW Enters., Inc. v. Interact Accessories, Inc., 424 F.3d 1324, 1331 (Fed.Cir.2005) (internal quotations omitted). In particular, Intergraph argues that the term "selectable" is wrongly imported from the preferred embodiment, and that "will be effected" conflates the term with the separate function of cache mode

effecting means.

Toshiba acknowledges that the use of the term "will be effected" is based on the additional recitation in Claims 1 and 9 of "cache mode effecting means for effecting the cache data storage mode indicated by said system tag." Toshiba offers no justification for adding this extraneous limitation to the stand-alone "system tag means" term, and the court declines to do so.

As for the use of the term "selectable," Toshiba likewise all but admits that this is an importation of a limitation from the specification, which is unsupported by the claims. Accordingly, the term "selectable" need not be included in the construction of "system tag means." The court adopts Intergraph's proposed function: "storing (in this case, passively retaining) a system tag which indicates one of a plurality of cache data storage modes."

B. *Structure*

[18] Intergraph's proposed structure for "system tag means" is "memory 350 of TLB 270 that stores a system tag ST field of control bits." Toshiba's proposed structure is "a storage location in a memory." Intergraph claims that Toshiba's structure is improperly broad because there are many memories described in the specification, including I/O memory 154, cache memory 320, and ROM 700. PL's Exh. E, Colwell Dep. at 74:3-6, 74:16-18, 80:9-81:8. Intergraph claims that none of these structures performs the claimed function, but all would satisfy Toshiba's definition. Specifically, I/O memory 154 holds disk controller programs, '846 Patent at 19:65-20:1, cache memory 320 holds data, *id*. at Fig. 10B, and ROM 700 holds instructions for microengine 650. *Id*. at 35:44-68. Intergraph claims that none of these structures store system tags and therefore cannot serve as corresponding structures.

In support of its own proposed structure, Intergraph points to a statement of Dr. Wolfe that the TLB is linked to the storage of system tags that indicate cache data storage. Wolfe Report para. 53. Additionally, Figure 11B of the ' 846 Patent depicts a line of the TLB 350, including the ST field. The ST field holds bits that, according to Intergraph, indicate cache data storage modes. Pl's Exh. E, Colwell Dep. at 87:4-13. Dr. Colwell likewise testified that the TLB stores the relevant system tags. Id. at 84:8-19. Finally, Intergraph claims that the ST field of the TLB is the only structure that stores system tags and that is part of the cache memory management system. The patent states that "[t]he various access modes can be intermixed on a page by page basis within the translation logic of the cache-memory management system". '846 Patent Abstract. Dr. Colwell testified that the memory of this translation logic is TLB 350, and that the ST field is stored within memory 350 of the TLB. Pl's Exh. E, Colwell Dep. at 84:13-24, 85:22-25, 86:19-87:23. This, according to Intergraph, confirms that only the ST field within memory 350 of the TLB is linked to the claimed function.

In response, Toshiba identifies a number of structures other than TLB 270 that store system tags: page tables in main memory, '846 Patent Abstract, 2:46-52, 25:61-62, system bus input register 260 when the tags are transferred into the cache MMU from main memory, the DAT of cache MMU, id. at 31:39-53, 31:67-32:3, and the hardwired TLB, id. at 26:26-27. Toshiba claims that Intergraph's proposition that the memory must be stored in TLB 270 lacks a technical explanation because the actual structure that performs the storing function is the memory rather than the surrounding TLB 270 structure. Regarding these additional structures identified by Toshiba, Intergraph claims that the page tables are not part of the cache memory management means as required by Claim 1. Additionally, Intergraph claims that there is nothing to suggest a link between the additional structures and the performance of the recited function. Intergraph further claims that the hardwired TLB and the DAT do not have storage locations. '846 Patent at Fig. 8B, 31:67-32:2.

While Toshiba's construction is unacceptably broad, the intrinsic record does not support limiting the corresponding structure to the single component identified by Intergraph. Because many of Intergraph's

arguments relate to the requirement that the system tag means be located in the cache memory management means, as is required by Claim 1, the court adopts a modification of Toshiba's proposed structure: "a storage location in the cache memory management means."

VII. ***Translation Means***

The term "translation means" appears in independent claim 9 of the '846 Patent. The parties agree that the term is covered by Section 112(6) and that the function is "providing an associative map of virtual to real addresses corresponding to data stored in said primary memory." Once again, the dispute over the corresponding function is a matter of general or specific. Intergraph proposes "array 352/354 of TLB 350 of Fig. 9," while Toshiba proposes "an array of a TLB." Dr. Colwell acknowledged that "an array of TLB" refers to array 352 or 354 of TLB 350. Pl's Exh. E, Colwell Dep. at 258:7-16. Toshiba's only argument in support of its construction is that the inclusion of the numbers would add confusion. It is difficult to see how increased specificity adds confusion. Accordingly, the court adopts Intergraph's proposed structure: "array 352/354 of TLB 350 of Fig. 9."

VIII. ***Primary Memory Interface Means***

The term "primary memory interface means" appears in independent Claim 9 of the '846 Patent. The parties agree that the term is governed by Section 112(6) but disagree as to the function and corresponding structure.

A. ***Function***

[19] Intergraph's proposed function is "coupling data between said primary memory and said cache memory." Toshiba's proposed function is "(a) coupling data between said primary memory and said cache memory; (b) and selectively transferring data between said primary memory and said cache memory in response to a miss signal." The dispute centers on whether the "wherein" clause in the claim imposes an additional function to the term. Claim 9 contains the clause "primary memory interface means for coupling data between said primary memory and said cache memory," followed several clauses later by the clause "wherein said primary memory interface means selectively transfers data between said primary memory and said cache memory in response to a miss signal." Accordingly, the court must decide whether the latter clause constitutes an additional function which must be performed by the primary memory interface means.

The parties dispute the legal effect of the "wherein" clause. Intergraph asserts that the word "wherein" should be treated the same as "whereby." Intergraph's sole support for this contention is a citation to a section of the Manual of Patent Examining Procedure ("MPEP") dealing with the terms "adapted to," "adapted for," "wherein" and "whereby." Pl's Exh. G, MPEP s. 2111.04. Significantly, "adapted to" and "adapted for" are treated as a single type of clause, while "whereby" and "wherein" are treated as separate types of clauses. *Id*. The MPEP gives these clauses as "examples of claim language ... that may raise a question as to the limiting effect of the language in a claim." *Id*. The MPEP further states that "[t]he determination of whether each of these clauses is a limitation in a claim depends on the specific facts of the case." *Id*.

Intergraph cites no case where a "wherein" clause was held not to impose a limitation. Instead, Intergraph cites two cases dealing with "whereby" clauses: Texas Instruments, Inc. v. U.S. Int'l Trade Comm'n, 988 F.2d 1165, 1172 (Fed.Cir.1993) (holding that "[a] 'whereby' clause that merely states the result of the limitations in the claim adds nothing to the patentability or substance of the claim"); Lockheed Martin Corp. v. Space Sys./Loral, Inc., 324 F.3d 1308, 1319 (Fed.Cir.2003) (holding that "[t]he function is properly identified as the language after the 'means for' clause and before the 'whereby' clause"). The court in Griffin v. Bertina, 285 F.3d 1029, 1034 (Fed.Cir.2002), held that "wherein" and "whereby" were two different parts of speech and confirmed that the inquiry as to whether these words create limitations is fact-dependent.

There, the court held that a "wherein" clause imposed a limitation because it "expresse[d] the inventive discovery" rather than "merely stat[ing] the inherent result of performing the manipulative steps." Id.

[20] Accordingly, the court cannot consider the word "wherein" in isolation, but must determine whether the "wherein" clause in Claim 9 expresses an inventive component or merely the result of the delineated limitations. Intergraph characterizes Claim 9 as reciting seven separate structural elements, including the primary memory interface means, each with its own clause. The two "wherein" clauses following these seven component clauses, according to Intergraph, simply express how the components are intended to operate together in a computer system. In response, Toshiba argues that treating the "wherein" clause in Claim 9 as non-limiting would force the court to ignore several dependent claims that consist of nothing more than "wherein" clauses. Toshiba's argument misses the mark. As with the specific "wherein" clause at issue in Claim 9, each "wherein" clause elsewhere in the patents must be individually considered to determine whether the clause simply recites a result or sets forth an inventive limitation.

However, the court cannot accept Intergraph's proposed function, because the "wherein" clause at issue does not merely recite a result. Rather, it clearly recites an additional function of the primary memory interface means. Notably, in each of the cases cited by Intergraph, the "whereby" clauses stated pure results rather than actions directly involving the structures or limitations at issue. In Texas Instruments, 988 F.2d at 1169, the result of the claimed steps was that "the fluid will not directly engage the device and electrical connection means at high velocity, and the conductors will be secured against appreciable displacement by the fluid." Likewise, in Lockheed Martin, 324 F.3d at 1315, the term at issue was "means for rotating said wheel in accordance with a predetermined rate schedule which varies sinusoidally over the orbit at the orbital frequency of the satellite," and the immediately ensuing "whereby" clause set forth the necessary result of rotating the wheel as described, i.e. offsetting the attitude of a satellite in a particular manner. In Griffin, 285 F.3d at 1031, the clause at issue was

assaying for the presence of a point mutation in the nucleotides of codon 506 within EXON 10 of the human Factor V gene, wherein said point mutation correlates to a decrease in the degree of inactivation of human Factor V and/or human Factor Va by activated protein C, wherein the presence of said point mutation in said test nucleic acid indicates an increased risk for thrombosis or a genetic defect causing thrombosis.

Because the correlation was the inventive discovery, the wherein clauses were necessary limitations.

Here, the "wherein" clause does not recite a necessary result of arranging the delineated components of the claimed apparatus. Rather, it recites a specific function of the primary memory means. Selectively transferring data between the primary memory and cache memory in response to a miss signal is not a necessary result of coupling data between the primary memory and the cache memory. It is a separate function performed by the primary memory interface means which must be read as part of the limitation.

Accordingly, the court adopts Toshiba's proposed function: "(a) coupling data between said primary memory and said cache memory; (b) and selectively transferring data between said primary memory and said cache memory in response to a miss signal."

B. ***Structure***

Intergraph's proposed structure for this term is "SIR 260 and SOR 250 of ' 846 Patent Fig. 8, also labeled as 643 and 644 of Fig. 21." Toshiba claims that no corresponding structure is disclosed in the specification tied to performing the entire recited function, and that the claim is therefore indefinite under Section 112. The dispute over the structure tracks the dispute over the function. Intergraph asserts that its proposed structure satisfies its proposed function, while Toshiba claims that no structure performs the entire function. Toshiba appears to acknowledge that the proposed structure performs the data coupling function. Accordingly, the

issue is whether Intergraph's proposed structure performs the selective data transfer function.

Intergraph asserts that Dr. Colwell acknowledged that the proposed structure performs this function. In response to the question of whether he was "aware of any structure described in the specification that takes action in response to a miss signal," Dr. Colwell identified "boxes 120 and 130" of Figure 1 of the '846 Patent. Pl's Exh. E, Colwell Dep. at 270:2-13. Toshiba claims that Dr. Colwell's was making the general point that having a structure which responds to a miss is necessary for a cache to be present. In particular, Dr. Colwell stated that "[w]hen you take a miss, there's a bunch of things you have to do; otherwise, you don't have a cache." Id. at 270:5-7.

Intergraph identifies an additional structure for the first time in its reply brief, apparently having abandoned its assertion that the selective data transfer function is satisfied by boxes 120 and 130. Instead, Intergraph now claims that microengine 650 is "clearly linked to the function." Intergraph argues that microengine 650 sends an address to memory on a read request, that an artisan would have understood that data is returned from memory on a read request, and that an artisan would have made the link between a cache miss and a miss signal. The specification states that microengine 650 transmits data between the main memory and the cache-MMU memories in response to a "cache miss." '846 Patent at 6:60-8; *see also* Wolfe Dec., Exh. A para.para. 85-86. This corresponds to a selective data transfer between the primary memory and the cache memory in response to a miss signal.

IX. ***Primary Memory Means***

The parties agree that Section 112(6) applies to the term "primary memory means" as it is used in Claims 15 and 16 FN2 of the ' 835 Patent. Toshiba claims that Section 112(6) does not apply to "primary memory means" as it is used in Claims 1 and 4. Claim 15 reads as follows:

FN2. Although Claim 16 is mentioned in Intergraph's opening brief, the parties raise no specific claim construction arguments directed toward this claim. Accordingly, the court will not reach a separate construction of "primary memory means" as the term is used in Claim 16.

15. The computer system according to claim 14 wherein the primary memory means couples data from the selected storage location to the system bus when the second cache inquiry signal is provided and the first ready signal is absent.
In Claim 1, the first component of the claimed "computer system" is "primary memory means for storing data at storage locations therein." Claim 4 is a dependent claim "wherein the primary memory means further comprises access control means for selectively allowing and inhibiting access to the primary memory means when the storage location to be accessed in the primary memory means may correspondto data stored in the cache memory means."
"[W]here a claim recites a function, but then goes on to elaborate sufficient structure, material, or acts within the claim itself to perform entirely the recited function, the claim is not in means-plus-function format." Sage Prods., 126 F.3d at 1427-28. Toshiba claims that, while the added functions in Claim 15 bring "primary memory means" within Section 112(6), the recitation of "primary memory means" in Claim 1 recites sufficient structure to perform the recited function, i.e. "memory." *See* Katz v. AT & T Corp., 63 F.Supp.2d 583, 640 (E.D.Pa.1999).

Intergraph raises a number of arguments against treating the term differently according to which claim it appears in. First, there is the general principle that claims should be construed in a manner that "will preserve the patent's internal coherence." Markman, 517 U.S. at 390, 116 S.Ct. 1384. Second, there is the presumption that the use of the word "means" is intended to invoke Section 112(6). Al-Site Corp., 174 F.3d at 1318. Finally, Intergraph argues that failing to apply Section 112(6) to the term "primary memory means"

in independent Claim 1 while applying the provision to the term in other claims which antecedently define the same term would result in the latter claims being construed more broadly than the independent claim. Under Section 112(4), "[a] claim in dependent form shall be construed to incorporate by reference all the limitations of the claim to which it refers." However, Section 112(6) provides that a means-plus-function claim "shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof." Accordingly, if "primary memory means" were treated as a structure in Claim 1 but as a means-plus-function term in the dependent claims, the dependent claims would cover equivalents not covered by the independent claim. Intergraph's approach avoids this problem by uniformly treating "primary memory means" as a means-plus-function term, with the dependent claims adding additional limitations to the independent claim.

Intergraph's reading of Section 112(4) is overly restrictive. The language in the statute refers to claims, not claim terms. The additional limitations recited in a particular dependent claim necessarily renders the entire claim less broad than the independent claim, even if individual claim terms may be interpreted more broadly in the dependent claim. Furthermore, in light of the flexible approach to means-plus-function interpretation, construing each term in the context of its specific claim to determine the appropriate function and structure does not harm the overall consistency and coherency of the patent. Means-plus-function terms may be used in different ways in different claims.

As used in Claim 1, therefore, the term "primary memory means" clearly recites sufficient structure-a memory which stores data at storage locations within the memory-to exclude the term from Section 112(6) in Claim 1. *See* Katz, 63 F.Supp.2d at 640. As used in other claims, however, "primary memory means" implicates Section 112(6) because the claims do not recite sufficient structure to perform the stated functions.

For Claim 15, for which Toshiba agrees that Section 112(6) applies, Toshiba's proposed function is "coupling data from the selected storage location to the system bus when the second cache inquiry signal is provided and the first ready signal is absent," and its proposed structure is "drivers/receivers and interface circuitry, and a memory for storing data." Intergraph apparently agrees with this asserted function, and argues that the corresponding structure is "the portion of the Drivers/Receiversand Interface circuitry depicted in box 140 of Fig. 1 that performs the protocols set forth in '835 13:15-16, 41-43, 46-48." Again, Toshiba's generalized formulation lacks the necessary specificity, and the court therefore adopts Intergraph's proposed structure.

As for Claims 4-8 and 9-14, Toshiba argues that Section 112(6) does not apply, and offers varying constructions depending on where the term appears in the patent. For Claims 4-8, Toshiba construes "primary memory" as "a memory for storing data, and including 'access control means' as recited." For claims 9-14, Toshiba construes the term as "a memory for storing data." The issue with respect to these claims is whether the construction for the purposes of Claims 4-8 should include "and including 'access control means' as recited." Intergraph argues that this is a separately recited limitation, and that it would be redundant to incorporate this limitation into the primary memory means. Toshiba asserts that its construction follows the plain claim language. The court agrees with Intergraph that the inclusion of the "access control means" language in Claims 4-8 is a recitation of an additional function of the primary memory means recited in the other claims. Accordingly, the court adopts Intergraph's construction.

X. ***Data Consistency Means***

The term "data consistency means" appears in Claim 1 of the '846 Patent. The parties agree that the term is covered by Section 112(6), and agree on the function. The parties further agree that the structure includes the system bus control 840 and microengine 650, that microengine 650 is a general purpose microcontroller, and that the structure of microengine 650 includes an algorithm. The parties disagree as to what algorithm is

linked to the function. Intergraph's proposed structure is: "system bus control 840 and microengine 650 programmed to execute a set of instructions to perform an internal cycle to determine whether to purge the cache of its data and supply data from an addressed location or to supply its data instead of supplying data from an addressed location." Toshiba's proposed structure is:

includes (a) system bus control 840; (b) microengine 650 programmed to execute a set of instructions to perform an internal cycle (I) for writes to the primary memory where the data is in the cache, writing to the cache for writes of less than a full cache line, and marking the cache line invalid for writes of a full cache line, and (ii) for reads, if the data has been modified in the cache, supply data, as specifically found at column 40, lines 30-37; and (c) system bus monitoring means, as construed infra.

The dispute centers on the issue of whether the algorithm contains limitations on write sizes. Intergraph's algorithm contains no write size limitation.

Intergraph asserts that the specification discloses multiple alternative algorithms for maintaining data consistency, and that the court should adopt as the corresponding structure only the algorithm or algorithms necessary to perform the function. *See* Creo Prods., Inc. v. Presstek, Inc., 305 F.3d 1337, 1344-46 (Fed.Cir.2002). Intergraph specifically claims that no size limitation is necessary to maintain data consistency, and that Toshiba's algorithm would effectively import limitations from the preferred embodiment into the function (as opposed to the structure). Dr. Colwell apparently acknowledged that Toshiba's algorithm imports limitations from the preferred embodiment, stating that in developing the proposed construction he only considered the embodiment set forth in the specification. Pl's Exh. E, Colwell Dep. at 235:25-237:17. Finally, Intergraph asserts that Toshiba's algorithm is not disclosed in the specification. The Federal Circuit has held that a "computer-implemented means-plus-function claim is limited to a computer programmed to perform the algorithm disclosed in the specification." Creo Prods., 305 F.3d at 1345.

Toshiba defends its additional size limitations by claiming that only one algorithm set forth in the specification is directly linked to microengine 650. Toshiba's algorithm is a slightly modified version of the algorithm set forth at Column 40, lines 29-37 of the '846 Patent. In contrast, Intergraph's proposed algorithm appears at Column 13, lines 3-17, in a section which does not mention microengine 650.

Although it is settled that microengine 650 is part of the corresponding structure, Toshiba's approach is far too rigid. The fact that microengine 650 is not mentioned by name in the portion of the specification discussing Intergraph's proposed algorithm does not indicate that microengine 650 could not perform the described algorithm. What matters is what the claims describe, not the relationships set forth in the specification. Toshiba has pointed to nothing in the claim language that would justify importing its proposed size limitations. Although the specification uses single-words and four-word lines as examples, the technology at the time of invention was not limited to these sizes. The court therefore adopts Intergraph's algorithm.

XI. ***System Bus Monitoring Means***

[21] The term "system bus monitoring bus means" appears in Claim 1 of the '846 Patent as part of the data consistency means. The parties agree that Section 112(6) applies, and disagree as to the function.FN3 Intergraph's proposed function is "monitoring (i.e., watching) reads and writes from, or on behalf of, I/O devices over the system bus." Toshiba's proposed function is "detecting reads and/or writes as issued by an I/O data processing device over the system bus to an addressed device." The controversy centers on three terms in Toshiba's construction: "as issued," "detecting" (rather than "monitoring"), and "to an addressed device," and on Intergraph's use of the term "from, or on behalf of, I/O devices."

FN3. The parties appear to agree on the corresponding structure, though Toshiba argues that the structure is too indefinite. This indefiniteness argument is inappropriate at the claim construction stage.

The disputes over the terms "as issued" and "from, or on behalf of, I/O devices" amount to a disagreement as to whether I/O requests are sent from I/O data processing devices only, as Toshiba claims, or may be sent directly from I/O devices through the system bus, as Intergraph claims. Under Intergraph's construction, I/O requests are sent either directly from I/O devices, or are sent from I/O devices and then processed by the I/O processing device. The I/O processor as described in the specification is able to perform protocol conversion, meaning the request as issued by the I/O device is converted by the I/O processing device before it is transmitted over the system bus. '846 Patent at 19:1-15, Wolfe Dec., Exh. A para. 65. The parties agree that, if the request is ultimately sent by the I/O processing device, it is monitored "as issued." Intergraph does not appear to argue that a request is modified if it is sent directly by an I/O device without being converted by an I/O processing device. Accordingly, the real dispute is whether the claim should encompass requests sent by an I/O processing device only, or requests sent by either an I/O processing device or an I/O device.

The claim language is ambiguous in this regard, stating that the system bus monitoring means is "coupled to the system bus, for monitoring I/O requests over the system bus." '846 Patent at 41:31-33. The claim therefore does not specify the origin of the I/O requests. In describing one embodiment, the relevant portion of the specification reads as follows:

[A]s illustrated in FIG. 1, other system's elements can be coupled to the system bus 141, such as an I/O processing unit, IOP 150, which couples the system bus 141 to the I/O bus 151. [...] I/O bus 151 can couple to the secondary storage or other peripheral devices, such as hard disks, floppy disks, printers etc. Multiple IOPs can be coupled to the system bus 141 and thereby can communicate with the main memory 140.

'846 Patent at 4:33-42. The specification therefore makes it clear that the I/O processing unit, not the secondary storage or peripheral I/O devices, communicates with the main memory via the system bus. Intergraph nonetheless argues that a person of ordinary skill in the art would have known that I/O devices could issue I/O requests. Intergraph claims that Dr. Colwell acknowledged that I/O devices could issue I/O requests, though the portion of Dr. Colwell's deposition indicates only that he "vaguely" remembered the portion of the specification discussed above. Defs' Exh. 3, Colwell Dep. at 247:8-12. Dr. Wolfe testified that although nothing in the patent specifically states that the I/O processor 150 initiates the read/write operation "on behalf of" an I/O device, he believed that a person of ordinary skill would interpret the specification as such in light of the fact that the I/O processor itself would have no need for data except when processing requests initiated by I/O devices. Defs' Exh. 2, Wolfe Dep. at 33:9-35:3. The passage that Dr. Wolfe was discussing, however, confirms that the I/O read/write operations are "initiated by an I/O processor, IOP 150." '846 Patent at 16:9-10. Intergraph has therefore pointed to nothing in the patent suggesting that I/O operations are initiated by I/O devices rather than the I/O processor. Furthermore, the I/O processor is the component which communicates with the system memory. Accordingly, Toshiba's construction is correct to the extent that it requires the I/O requests to be issued by an I/O processing device. However, Toshiba has provided no justification for the inclusion of the extraneous term "as issued." However, the reads and writes must still be "issued," and therefore the construction should include the word "issued," but not the word "as." Accordingly, the proper construction does not include that term.

Additionally, Intergraph claims that the substitution of the word "detecting" for "monitoring" is problematic in that it "adopt[s] a function different from that explicitly recited in the claim." JVW Enters., 424 F.3d at 1331 (internal quotations omitted). Dr. Colwell acknowledged that "monitoring" is different from "detecting," Colwell Dep. at 246:23-247:7, indicating that the substitution of the word "detecting" in place of the recited word "monitoring" alters the meaning of the claim. This is in contrast to Intergraph's own use

of the word "watching," which Intergraph claims is consistent with the specification and is merely a clarification of the word "monitoring." Toshiba's only proffered support for its construction is Intergraph's acknowledgment that part of the corresponding structure for "system bus monitoring means" includes a component labeled as a "detector." This argument is unconvincing. Nor does the court see any reason to adopt Toshiba's proposed synonym "watching." The court will instead use the term "monitoring," which is sufficiently clear.

Finally, Intergraph claims that the addition of the term "to the addressed device" improperly imports a limitation from the preferred embodiment which is not required by the claims. Intergraph acknowledges that, in the embodiment, I/O requests are directed to addressed devices such as main memory, but asserts that this is not a necessary limitation. In response, Toshiba claims that if the I/O request were not to an addressed device, nothing would determine the location to read data from or write data to. Toshiba points to the broader context of the claim language at issue, asserting that the system bus monitoring means must be considered in the context of the data consistency means. Both parties' proposed interpretation of "data consistency means" includes references to addressed components. In the context of the claim, an I/O request only makes sense if there is an addressed location. Accordingly, Toshiba is correct in this regard.

In sum, the function of "system bus monitoring means" is "monitoring reads and/or writes issued from an I/O data processing device over the system bus to an addressed device."

CONCLUSION

For the foregoing reasons, the court construes the disputed claims in the manner described above.

IT IS SO ORDERED.

N.D.Cal.,2007.
Intergraph Hardware Technologies Co. v. Toshiba Corp.