UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

FIFTH GENERATION COMPUTER CORPORATION, Plaintiff, -v- INTERNATIONAL BUSINESS MACHINES CORPORATION, Defendant.	09 CV 2439 (JSR) OPINION AND ORDER

JED S. RAKOFF, U.S.D.J.

In this patent infringement action, plaintiff Fifth Generation Computer Corporation ("Fifth Generation"), the current owner of U.S. Patents 4,860,201 ("the '201 Patent") and 6,000,024 ("the '024 Patent") alleges that defendant International Business Machines Corporation ("IBM") infringed one or both of those patents. Following briefing, the Court held a "Markman" hearing on August 20, 2009 to hear testimony as to the meaning of the disputed terms of the patents here in issue. See Markman v. Westview Instruments, Inc., 517 U.S. 370 (1996). On August 31, 2009, the Court issued a "bottom-line" Order that construed the disputed terms. This Opinion and Order provides the reasons for those constructions.

Familiarity with the parties' submissions and the testimony presented at the Markman hearing is here assumed. Generally speaking, the patents at issue relate to the means by which computers process data and are designed to increase the speed with which

certain computational tasks are completed. In these patents, increases in speed are accomplished through parallel processing, whereby tasks are divided into smaller tasks that are performed simultaneously. Markman Tutorial at 2.[1] The parallel processing is achieved through the use of one type of parallel computer: the "binary tree" computer. Id. In a binary tree computer, each processor (or node) is connected to one or more "child" or "leaf" processors (or nodes) to form communication "trees." Id. at 3.

In particular, the '201 patent (issued on August 22, 1989 to Salvatore J. Stolfo and Daniel P. Miranker) seeks to improve upon previous systems that suffered "propagation delays" (i.e., delays from computing and communicating the data up and down a tree) that resulted from the processors' handling both computing and communicating functions. '201 Patent at 3:29-32. The '201 patent arranges for a separate input/output device to accompany each processing element to handle the broadcast and reporting of information up and down the "binary tree" (a disputed term), thus dedicating the processing elements to other tasks. Id. at 4:64-5:2.

The '024 patent (issued on December 7, 1999 to James L. Maddox) describes a binary tree computer system that attempts to

[1] At the Court's request (and as is this Court's normal practice), the parties, in advance of the Markman hearing, presented the Court, orally and in writing, with a joint "tutorial" on those technical issues relating to these patents as to which the parties were in agreement.

improve on the '201 patent. Here, a series of "bus controllers" (a disputed term) control information flow in the system and are arranged in a binary tree configuration. '024 Patent at 1:48-49. Processors are attached to these bus controllers to form a binary tree of processing elements. Id. at 54-55. In contrast with the '201 patent, in which each processor has its own I/O device, in the '024 patent the bus controllers are themselves arranged in a binary tree configuration. At the extremes, the bus controllers are connected to a set of leaf processing elements -- that is, the bus controllers at the extremes have more than one processing element connected to them. The bus controllers "act collectively as a bucket brigade," '024 Patent at 5:38-39, to transfer information throughout the computer system.

With this background in mind, the Court turns first to the disputed terms of the '201 patent.

1. binary tree

The first term in dispute is "binary tree," which appears in every claim in the '201 Patent, including the four here relevant: Claims 1, 4, 7, and 8. (Claim 1 is an independent claim, and 4, 7, and 8 are dependent.) Claim 1 reads in relevant part: "A parallel processor array comprising . . . means for interconnecting said processing elements in a binary tree in which each processing element except those at the extremities of the binary tree is connected to

one parent processing element and at least first and second child processing elements." '201 Patent at 69:60, 67-69; 70:59-60.

Plaintiff's proposed construction is "a tree where a node has a parent node (except for the root node) and zero, one or two children." Fifth Generation's Rebuttal Markman Brief ("Pl. Reply") at 3. In contrast, defendant reads the term as "an arrangement of nodes where each node has a single parent and two children nodes, except the root node, which has no parent, and the leaf nodes, which have no children." IBM's Opening Markman Brief ("Def. Br.") at 9.

In Phillips v. AWH Corp., 415 F.3d 1303 (Fed Cir. 2005) (en banc), the Federal Circuit explained that "the words of a claim are generally given their ordinary and customary meaning" and that "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 1312-13. Here, plaintiff asserts that its construction reflects the "ordinary meaning" of the term "binary tree," Pl. Reply at 3; see also tr. 08/20/09 at 48, and cites both to the dictionary of the National Institute of Standards and Technology ("NIST"), in which "binary tree" is defined as "a tree with at most two children for each node," and to IBM's own website, on which appears a similar definition. See E-Pass Technologies, Inc. v. 3Com Corp., 343 F.3d 1364, 1367 (Fed. Cir. 2003) ("We resort initially to the relevant dictionary

definitions to determine the ordinary meaning of the [disputed] term . . . .").

However, since the '201 patent was filed in 1986, see '201 Patent at 1, the relevant meaning is the meaning at the time of the invention, see Phillips, 415 F.3d at 1313; PC Connector Solutions LLC v. SmartDisk Corp., 406 F.3d 1359, 1363 (Fed. Cir. 2005) ("A claim cannot have different meanings at different times; its meaning must be interpreted as of its effective filing date."), and there is no evidence before the Court that establishes what the dictionary definition of "binary tree" was in 1986. Indeed, it appears on this record that there was no entry for "binary tree" in either the 1986 version of the NIST dictionary nor in the contemporaneous version of IBM's Dictionary of Computing published in March 1987. See Def. Letter, 09/27/09. Given the huge changes in computer data processing, and accompanying jargon, between 1986 and the present, resort to NIST and IBM dictionaries from the present does not resolve the issue of the term's ordinary meaning.

The Court must therefore look to the usage of the claim term in the context of the particular claim and in the context of the entire patent as submitted in 1986. See Phillips, 415 F.3d at 1313 ("[T]he 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.").

The parties agree that the structure of Claim 1, the means-plus function language of "means for interconnecting," requires that the binary tree contemplated by Claim 1 have two children for each processing element except for those at the extremities. See tr., 08/20/09, at 44. Yet, plaintiff's proposed construction contradicts the language of Claim 1 and flies in the face of this structure. Taking plaintiff's construction on its face, a node with a single parent could have a single child, and that child could, in turn, have another single child, thus creating a straight line from one node to another, and this structure would still be a "binary tree" even though there is nothing binary about it. See tr., 08/20/09, at 46. Indeed, plaintiff's construction would eliminate what plaintiff itself takes to be the preferred embodiment and therefore cannot be the proper construction. See Primos Inc. v. Hunter's Specialities, Inc., 451 F.3d 841, 848 (Fed. Cir. 2006) ("While we are mindful that we cannot import limitations from the preferred embodiments into the claim, we also should not normally interpret a claim term to exclude a preferred embodiment.").

Moreover, plaintiff can point to only one sentence in the patent specification that remotely provides support for its construction: "Since the subtree that is computing the resolve need not be a complete binary tree, the value bits from the children may not arrive in the same clock cycle." '201 Patent at 21:46-49. Plaintiff asserts that this phrase from the patent specification

indicates that the term "binary tree" permits a tree with zero, one or two children. Pl. Reply at 4. Yet this phrase does not define the binary tree qua structure -- it only indicates that some tasks may rely on longer legs of a subtree, and thus the report and resolve functions would not be completed simultaneously. See tr., 08/20/09, at 55.

IBM's construction does not suffer from these difficulties; but plaintiff contends that IBM's construction improperly imports limitations from the preferred embodiment, see Phillips, 415 F.3d at 1323, because the preferred embodiment in no way offers a "clear disavowal" of embodiments of the binary tree computing system having less than two child nodes. Pl. Reply at 4. By way of support, plaintiff underscores the fact that the patent identifies Figure 2 as "the general configuration of a preferred embodiment of the invention comprising a binary tree of 1023 processing elements." '201 Patent at 8:11-13. In actuality, however, defendant's proposed construction does not import a limitation from the specification, but, rather, it relies on the fact that "the specification is the single best guide to the meaning of a disputed term," Phillips, 415 F.3d at 1320 (internal quotation marks omitted). Here, in fact, the specification is more properly read as dictating the correct scope of the claim, and in this respect, the "the inventor's intention, as expressed in the specification, is regarded as dispositive." Id. at 1316. In this regard, defendant has pointed not only to the claim language

itself but numerous instances throughout the specification that confirm that the binary tree contemplated by the patent requires each processing element, except for the root node and the extremities, to have two children. The patent abstract describes the patent as "a plurality of parallel processing elements [. . .] connected in a binary tree configuration, with each processing element except those at the highest and lowest levels being in communication with a single parent processing element as well as the first and second (or left and right) child processing elements." '201 Patent Abstract. The same information is repeated in the section of the patent devoted to the background of the invention, as well as the summary of the invention. Id. at 1:62-65; 4:56-62; see also Modine Mfg. Co. v. U.S. Int'l Trade Comm'n, 75 F.3d 1545, 1551 (Fed. Cir. 1996) ("[W]hen the preferred embodiment is described in the specification as the invention itself, the claims are not necessarily entitled to a scope broader than that embodiment.").

Nor is the holding of E-Pass, upon which plaintiff relies heavily, see tr. 08/20/09 at 25, 29, to the contrary. There, the Federal Circuit overturned the district court's construction of "electronic multi-function card" as "[a] device having the width and outer dimensions of a standard credit card with an embedded electronic circuit." E-Pass, 343 F.3d at 1366. The Federal Circuit found that the district court had taken the preferred embodiment to impose a particular size on the term "card," ignoring language

elsewhere in the patent that suggested that the patent was not based on the standardized size of a credit card. Id. at 1370. Here, in contrast, the binary tree that appears in Figure 2 may be a preferred embodiment, but Figure 2 does not suggest that the requirement of two children except at the extremities is a limitation derived from that preferred embodiment. Indeed, the patent elsewhere states directly: "In a binary tree computer, a large number of processors are connected so that each processor except those at the root and leaves of the tree has a single parent processor and two children processors." '201 Patent at 1:62-65. This statement provides critical support for the idea that the definition given to the term "binary tree" by the patentee differs from plaintiff's proposed construction and conforms more closely with defendant's proposed construction. See Phillips, 415 F.3d at 1316 (emphasizing the role of the inventor as "lexicographer").

Based on the foregoing, the Court construes "binary tree" to mean an arrangement of nodes where each node has a single parent and two children nodes, except the root node, which has no parent, and the leaf nodes, which have no children.

2. subtree

The term "subtree" likewise appears in every claim. Claim 1 explains that the input/output provides a "means for broadcasting received from a parent processing element to said child processing elements, such that common information is distributed to each

processing element of the binary tree or a subtree thereof . . . ." '201 Patent at 70:62-66.

Plaintiff construes "subtree" to mean "a partitionable portion of a tree that is less than the entire tree created by arbitrarily denoting a node to be the root node in a tree," Fifth Generation's Opening Markman Brief on Claim Construction ("Pl. Br.") at 13, asserting that this construction is consistent with ordinary meaning "in the art of computing."[2] Id. Defendant construes the term as "a subset of the binary tree such that each node has a single parent and two children nodes, except the root node, which has no parent, and the leaf nodes, which have no children." Def. Br. at 10.

The parties essentially agree that a subtree is a subdivision of the main binary tree and that a subtree is itself a binary tree with the full functionality of the main binary tree of which it is a subpart. Pl. Br. at 13; Def. Br. at 10-11; Pl. Reply at 7. U.S. Patent No. 4,843,540 ("the '540 Patent") is incorporated by reference into the '201 Patent, and the '540 Patent establishes that "one of the characteristics of the binary tree" is that "it includes sub-sets which are also binary trees." '540 Patent at 3:49-51.

To support its construction, plaintiff cites as extrinsic evidence IBM's Terminology webpage, which defines "subtree" as "a

[2] Plaintiff originally omitted the term "arbitrarily" in the construction it proposed in a pre-briefing conference with defendant and subsequently added it in its briefs for the Markman hearing. IBM's Rebuttal Markman Brief ("Def. Reply") at 4 n.4.

tree structure created by arbitrarily denoting a node to be the root node in a tree." Pl. Br. at 13; Pl. Reply at 6. The '201 patent, according to plaintiff, is consistent with this definition, which describes a binary tree as capable of being "partitioned into any number of subtrees, which maintain the functionality of the tree." Pl. Br. at 13. Yet this extrinsic evidence is insufficient to support this construction, especially in light of the intrinsic evidence, as well as this Court's determination of the proper construction of "binary tree." See Vitronics Corp. v. Conceptronic, 90 F.3d 1576, 1583 (Fed. Cir. 1996) ("In most situations, an analysis of the intrinsic evidence alone will resolve any ambiguity in a disputed claim term. In such circumstances, it is improper to rely on extrinsic evidence."). Not only does the definition upon which plaintiff relies not reflect the meaning of the patent at the time of filing in 1986, see Phillips, 415 F.3d at 1313, but it impermissibly permits any node, including a node with zero children, such as one of the extremities, to be designated a root node, despite the fact that plaintiff concedes that the subtree must also be a binary tree. See Nystrom v. Trex, Co., 424 F.3d 1136, 1145 (Fed. Cir. 2005) ("In the absence of something in the written description and/or prosecution history to provide explicit or implicit notice to the public -- i.e., those of ordinary skill in the art -- that the inventor intended a disputed term to cover more than the ordinary and customary meaning revealed by the context of the intrinsic record, it is improper to read the term to encompass a broader definition simply because it may

be found in a dictionary, treatise, or other extrinsic source."). Moreover, since the Court has rejected plaintiff's construction of binary tree, which would have permitted a parent node with zero children, so too must it reject plaintiff's similar construction of "subtree."

Plaintiff objects to defendant's construction on the ground that it requires that the "root node" have no parent node. Pl. Reply at 7. The '201 and '540 Patents both establish clearly that subtrees, as partitions, remain connected to higher level nodes in a larger tree. See '540 Patent at 15:61-64. Yet plaintiff concedes that defendant's construction is accurate if it is understood that a subtree's lack of a parent node does not mean that the subtree is not connected to a higher level node, but rather that the higher level node is itself not considered part of the subtree. Pl. Reply at 8 n. 2; see also tr., 08/20/09, at 55. The Court agrees that this is the proper understanding of that construction.

Nor does defendant improperly rely on the functionality of the subtree to dictate its structure, as plaintiff suggests. See Pl. Reply at 7. As it did in the context of a binary tree, plaintiff cites to a single statement in the '201 Patent for the idea that a subtree need not have the same structure: "Since the subtree that is computing the resolve need not be a complete binary tree, the value bits from the children may not arrive in the same clock cycle." '201 Patent at 21:46-49. Yet, as noted previously, this establishes only that the report and resolve functions would not happen simultaneously

if a longer leg of the subtree were handling a particular task. See tr., 08/20/09, at 55. It does not, contrary to plaintiff's position, permit a subtree to have a different structure than a binary tree, namely, a parent with one or zero children. Indeed, plaintiff concedes that the subtree must have the same structure as the larger tree. See Pl. Reply at 8 n.3.

The Court thus confirms its construction that "subtree" means a subset of the binary tree such that each node has a single parent and two children nodes, except the root node, which has no parent, and the leaf nodes, which have no children. The Court's construction should be understood to mean that the parent node nevertheless remains connected to a higher level node in the main binary tree of which the subtree is a part because a subtree is always part of a larger binary tree.

3. without direct control of the processors of the processing elements

The third disputed term (actually, a phrase) appears in independent Claim 1 and dependent claims 4, 7, and 8. Claim 1 of the '201 Patent reads in relevant part:

> A parallel processor array comprising:
> said input/output means comprising:
> means for broadcasting information received from a parent processing element to said child processing elements, such that common information is distributed to each processing element of the binary tree or a subtree thereof without direct control of the processors of the processing elements; and
> means for determining a priority among respective values of information received from said child processing elements and information received from the processor with which said

input/output means is associated without direct control of the processors of the processing elements.

Patent '201 at 69:60; 70:61-69; 71:1-4.

Plaintiff construes the disputed claim term to mean that "the main processor at a node is not burdened with controlling all aspects of an operation."[3] Pl. Br. at 9. Defendant, in contrast, construes it to mean that the "[broadcasting]/[determining a priority] function is performed independently by the I/O device without receiving instructions from its associated processor." Def. Br. at 11.

The parties thus chiefly dispute the amount of control meant by the words "direct control," and both assert that their respective constructions reflect the ordinary, plain meaning of the words. See Pl. Reply at 9; Def. Br. at 11-12. Resolution of this issue requires focusing on the relationship of the I/O device to its associated processing element. In prior art, the processing elements also handled communication of data to other processing elements, and therefore "propagation times of the query and the result through the binary tree introduce significant delays in overall throughput comparable to those of a serial computer." '201 Patent at 3:29-33. The patent specification establishes that the patent's novel contribution, in part, was that the I/O device "communicates data and queries from the root processing element to all other N processing

[3] Plaintiff originally construed the term to mean that "the main processor at a node is not interrupted such that very little computational overhead is required for controlling the operation," but proposed the alternative construction in its briefs for the Markman hearing. See Def. Reply at 7 n.6.

elements in the array" so as to "minimize propagation delays in the binary tree computer." '201 Patent at 5:10-16. Communication from the root processor to the processing elements is the broadcast function. The specification also provides that "the I/O Circuit of the present invention also provides a high speed function resolve and report function that determines a most favorable value of a set of values stored in the processing elements and reports this value." '201 Patent at 5:56-59. This resolve and report function is the determination of a priority from a data set. In short, the patent specification establishes that the purpose of the I/O device was to increase processing speed by allowing it to handle broadcast and resolve functions, leaving the processing element to perform other tasks.

The dispute over the instant term (or phrase) turns on whether the processor controls, even indirectly, the I/O device during the broadcast and resolve functions. Plaintiff's construction attempts to capture the implication of "no direct control" by stating that the "main processor" is "not burdened with controlling all aspects of an operation," yet the term "main processor" appears nowhere in the patent nor does plaintiff's construction clearly indicate how much of a burden is shifted to the I/O device. Worse still, this construction would permit plaintiff to claim that the processor directly controls some aspects of the broadcast and resolve functions, in contradiction of the claim language "without direct control."

Although plaintiff's construction is not persuasive, plaintiff argues that defendant's construction erroneously requires that the processor have "absolutely no control" over the "broadcast" or "priority" functions and that intrinsic evidence supports the claim that the processor has some control over the state of the I/O device.[4] Pl. Reply at 9. For example, the processor's low level routines send instructions to the I/O device, thus controlling the resolve function by having the bus drivers transmit data back up the binary tree. Pl. Reply at 10 (citing '201 Patent at 28:32-37). Thus, while the I/O device may execute the broadcast and priority functions independently of the processor, according to plaintiff, the processor nevertheless exercises some control because it must send instructions to have the I/O device initiate those functions. See tr., 08/20/09, at 69.

Although there is no doubt that the processor communicates with the I/O device, see tr. 08/20/09 at 99, the disputed claim language here focuses on the execution of these particular functions, i.e., how "common information is distributed to each processing element" and how "determining a priority among respective values of

[4] Plaintiff goes so far as to claim that the processing element "operates to 'greatly assist the I/O procedure.'" Pl. Reply at 9, citing '201 Patent at 9:15-16. This citation takes the phrase out of context, however, as the full sentence indicates that the processor "is provided with four parallel, 8-bit ports which simplify interconnection between processors and greatly assist the I/O procedure." '201 Patent at 9:13-16. Thus, it is not quite the case that the processor itself "greatly assists" the I/O device.

information" is handled by the I/O devices. '201 Patent at 70:64-65, 68-69. As the patent specification suggests and the patent prosecution confirms, these functions are performed independently from the associated processor. The '201 Patent applicants originally omitted the phrase "without direct control," and the United States Patent and Trademark Office ("USPTO") rejected the claims as obvious in light of prior art, specifically U.S. Patent No. 4,583,164 ("Tolle"). See '201 Patent Prosecution History, Def. Br., Ex. F, at FBC 829-30. The applicants subsequently sought to distinguish their contribution from prior art by emphasizing that while "broadcasting" was directly controlled by the processors, here "once the broadcast operation is initiated . . . it proceeds independently of the processors . . . . [T]he 'means for determining priority' in each processing element is not under the control of the processor therein." Id. at FGC 793-94. After a second rejection, the applicants modified the claim language to include "without direct control." Id. at FGC 843-44. The applicants asserted that their system handled broadcasting faster than prior systems because it did not have to wait for the execution of processor instructions:

> Because 'broadcasting' in Tolle is directly controlled, by the processors of the cells of the binary tree through the execution of a storage management algorithm rather than through cooperation among input/output means within each cell . . ., the time required to move data through the binary tree in Tolle is on the order of the number of cells in the binary tree multiplied by the execution time of the instructions required for moving information from a parent cell to its left and right child cells.

Id. at 852.

Hence, by shifting the execution of the functions to the I/O devices, and away from the processors, the propagation delays were reduced. The prosecution history thus dictates that the execution of the broadcast and resolve functions are performed "independently" by the I/O devices. See Gillespie v. Dywidag Systems Intern., USA, 501 F.3d 1285, 1291 (Fed. Cir. 2007) ("The patentee is held to what he declares during the prosecution of his patent.").

It follows that "without direct control of the processors of the processing elements" means that the [broadcasting] /[determining a priority] function is performed independently by the I/O device without receiving instructions from its associated processor.

> 4. each in a time on the order of the logarithm of the number of processing elements . . .

Claim 1 of the '201 Patent states that the broadcasting and resolve functions are performed

> each in a time on the order of the logarithm of the number of processing elements in said binary tree or subtree multiplied by the time for the broadcasting of information from a parent processing element to child processing elements connected thereto, and the time required to determine priority amoung values of information received from the processor of a processing element and the child processing elements connected thereto, respectively.

'201 Patent at 71:12-21.

Plaintiff construes the phrase here in issue as "each operation is performed in time based on the base 2 logarithm (log2) of the total number of nodes in a tree or subtree multiplied by the time required to broadcast from a parent to a child node and the time required to determine the result data between a parent and its child

nodes." Pl. Br. at 16. Defendant's construction is "the broadcasting operation is performed in one clock cycle, multiplied by the base 2 logarithm of the number of processing elements in the binary tree; the priority is determined in two clock cycles, multiplied by the base 2 logarithm of the number of processing elements in the binary tree." Def. Br. at 14. According to defendant, "broadcasting" takes one clock cycle, and priority determination (or "resolve") takes 2 cycles (one for comparison and one to propagate the result to the next level). Thus, according to defendant, each operation is performed in the time that it takes to broadcast or resolve at each level multiplied by the number of levels in the tree. The time at each level is one clock cycle for broadcasting, and two for the priority determination. Def. Reply at 14.

Plaintiff's construction is flawed to the extent it suggests that "each operation" occurs in a time that is based on both the broadcast and priority determinations functions, omitting the words "each" and "respectively." See Def. Br. at 15. The parties are nevertheless essentially in agreement except for defendant's inclusion of a specific number of clock cycles, which plaintiff asserts contradicts the specific claim language of "on the order of." Pl. Reply at 12. The patent specification indicates, with respect to the broadcast function, that "data may be broadcast to all the [processing elements] of the array with a delay of only one cycle for each level of the array . . . . Since a processor instruction cycle

typically requires over twenty clock cycles, there is sufficient time in one processor instruction cycle to broadcast data to all the [processing elements]." '201 Patent at 19:35-37, 42-45.

As the patent specification indicates, a single processor instruction cycle typically requires more than twenty clock cycles, '201 Patent at 19:42-43, but the patent was intended to ensure that the I/O device "communicate data and queries from the root processing element to all other . . . processing elements in one processor instruction cycle," id. at 5:11-14. According to the prosecution history, the applicants' original claim language specified that each of the functions would take "less than an average processor instruction cycle." '201 Patent Prosecution History, Def. Br., Ex. F, at FGC 408. The USPTO rejected that language because of the inherent ambiguity of the word "average," id. at FGC 829, and the applicants further amended the claim language to specify "each in a time." Thus, the time in which each function can be performed is calculated by multiplying the number of levels in the binary tree by the time required for the broadcast or priority determination at each level.

With respect to the resolve function, the patent specification states: "[T]he entire resolve operation can be performed in the time it takes a number offered by a [processing element] to be clocked through the comparator plus the propagation time through the tree of one clock cycle per level. If the number offered is only a byte and the processor array only has ten levels, .

. . the entire resolve operation can be completed in less than twenty-clock cycles which is less than the average processor instruction cycle." '201 Patent at 21:67-22:7. Plaintiff argues that this is the preferred embodiment using the example of only a byte, with the implication that if the number were larger than a byte, the time would exceed one clock cycle. See tr., 08/20/09, at 74, 99. Yet, while the preferred embodiment here specifies a ten-level binary tree, see '201 Patent at 19:39-40, it in no way indicates that the speed at each level would vary based on the size of the number offered for processing. To the contrary, in the same section of the specification, the patent clearly states: "Advantageously, there is a delay of only a single clock cycle in performing the comparison at each level in the binary tree . . . ." '201 Patent at 21:63-65. Hence, the comparison requires only one clock cycle. The report and resolve function, or determining a priority, thus takes one cycle for the comparison and one clock cycle per level to propagate the date back up the binary tree, for a total of two clock cycles, multiplied by the number of levels in the tree.

At the Markman hearing, counsel for defendants conceded that the specification indicated that the functions were capable of being performed in that time, not that they must be performed in that time. See tr., 09/21/09, at 105-06. Accordingly, the Court's construction of the disputed term is that "the broadcasting operation is capable of being performed in one clock cycle, multiplied by the base 2 logarithm of the number of processing elements in the binary tree;

the priority is capable of being determined in two clock cycles, multiplied by the base 2 logarithm of the number of processing elements in the binary tree.

5-6. single instruction multiple data mode (SIMD mode)/ multiple instruction multiple data mode (MIMD mode)

These terms appear in dependent claim 8, which reads "the apparatus of claim 7 wherein each sub-tree is operated in a single instruction multiple data mode and the plurality of sub-trees are operated in a multiple instruction multiple data mode." '201 Patent at 72:29-32. The parties agree that SIMD mode and MIMD mode are two modes by which the binary tree computer system divides tasks and handles data and instructions. Pl. Reply at 14; Def. Br. at 16-17. The patent specification states that "in SIMD mode, each [processing element] is first loaded with its own data and then a single stream of instructions is broadcast to all [processing elements]. . . . In MIMD mode, each [processing element] is first broadcast its local program and data and then each [processing element] is logically disconnected from its neighbor [processing elements] and executes independently." '201 Patent at 10:41-43, 48-51.

Defendant's constructions track almost verbatim the definitions of these terms provided in the specification. See Def. Br. at 15. By contrast, plaintiff construes SIMD mode to mean that "[e]ach processing element has its own data but executes the same instruction as other processing elements," and construes MIMD mode to mean that "[e]ach processing element has its own program and data and

executes independently of its neighbor processing elements." Pl. Reply at 14. According to plaintiff, in SIMD mode, processors execute the same instructions on different sets of data, whereas in MIMD mode, processors execute different instructions on different sets of data. Id. The crux of plaintiff's objection to defendant's construction is that defendant has imported into the claim from the specification how each mode is initiated, thus imposing a limitation from the preferred embodiment. In short, plaintiff objects that the manner in which these modes are implemented are distinct and separate from the modes themselves. See tr., 08/20/09, at 75.

The Court disagrees. The terms are undefined in the claim, and thus the specification, upon which defendant predominantly relies, provides the appropriate definition of those terms. See Phillips, 415 F.3d at 1321 (emphasizing that the specification "acts as a dictionary when it expressly defines terms used in the claims . . . ."); Irdeto Access Inc., v. Echostar Satellite Corp., 383 F.3d 1295, 1300 (Fed. Cir. 2004) ("Even when guidance is not provided in explicit definitional format, the specification may define claim terms by implication . . . .").

The '540 Patent, which is incorporated by reference, see Pl. Br. at 7, provides additional support for defendant's construction. It confirms that the SIMD mode is operative where "data processing elements contain multiple subsets of the data set F and each operate on the same instruction." '540 Patent at 6:43-45. Moreover, it further specifies that "the data processing elements remain logically

connected to the root data processing element to receive inquires [sic] (as simply instructions or unknown data) and to report the results of the data processing operation." '540 Patent at 46-50. This at least confirms that in SIMD mode, in contrast with MIMD mode, the processing elements are provided a data set and then remain connected to the root processor so as to receive the instructions that all of the processing elements will execute. Similarly, for MIMD mode, the patent indicates that data processing elements can be provided with different instruction sets to be used on their own data: "The logical disconnection state of a data processing element from its parent during the data processing is called multiple instruction multiple data (MIMD) mode." Id. at 6:59-62. Hence, the '540 Patent corroborates that in MIMD mode, the processing elements are disconnected from other processing elements while executing their independent set of instructions.

Accordingly, the Court confirms the construction of the two terms set forth in the August 31 Order. SIMD mode is "where each processing element is first loaded with its own data and then a single stream of instructions is broadcast to every processing element in the binary tree." MIMD mode is "where each processing element has broadcast to it its local program and data and then each processing element is logically disconnected from its neighbor processing element and executes independently."

The Court turns now to the disputed terms of the '024 Patent.

The disputed terms appear in independent Claim 1 and dependent Claim 7. Claim 1 reads in relevant part:

> A binary tree computer system for connection to and control by a host computer, comprising:
> N bus controllers connected in a binary tree configuration in which each bus controller, except those at the extremes of the tree, are connected to left and right child bus controllers, where N is an integer greater than 2, one of said bus controllers being a root bus controller for connecting said binary tree connected bus controllers to said host computer.

'024 Patent at 7:11-19.

1. binary tree computer system

Plaintiff construes the phrase "binary tree computer system" to mean "a computing system (or partitionable portion of a computing system) with nodes connected in a binary tree configuration." Pl. Br. at 19. Defendant's competing construction is "a computer system of nodes connected in a binary tree configuration." Def. Br. at 18.

The parties therefore principally dispute whether a partitionable portion of a binary tree computer system is itself a binary tree computer system. Because subtrees have the full functionality of a main tree, plaintiff argues that its construction captures the notion that a "binary tree computer system" can be a "portion of a computing system," asserting that "[i]t is "well known that a binary tree computer can be divided into a 'subtree' that functions as the full computer system." Pl. Br. at 20. Defendant, for its part, does not dispute that trees are partitionable into subtrees but rather maintains that a subtree would not itself constitute a "binary tree computer system." Def. Br. at 19; Def.

Reply at 16-17. This reflects the plain meaning of the phrase, since the patent specifies that a binary tree computer system includes, inter alia, connection a host processor, see ACTV, Inc. v. Walt Disney Co., 346 F.3d 1082, 1088 (Fed. Cir. 2003) ("[T]he context of the surrounding words of the claim also must be considered in determining the ordinary and customary meaning of those terms."), whereas a subtree alone, as the Court has construed the term, would not require such a connection to the host computer.

Although plaintiff cites to the '201 and '540 Stolfo Patents as intrinsic evidence that subtrees are themselves "binary tree computer systems," the '024 Patent incorporates those earlier patents by reference but also deliberately distinguishes itself from them. '024 Patent at 1:23-28. The '024 patent itself does not reference "subtrees" or "partitionability" and as such the incorporation by reference provides no support for the idea that the patent applicant intended for "binary tree computer system" to include a partitionable portion thereof. Cf. Modine Mfg. Co., 75 F.3d at 1553 (arguing that incorporation by reference "does not convert the invention of the incorporated patent into the invention of the host patent" and finding that the single reference to the incorporated patent did not otherwise alter the "presentation of the invention" in the specification). Plaintiff's extrapolation from the fact that a subtree is itself a binary tree to asserting that a subtree of a binary tree computer system is itself a binary tree computer system exceeds the bounds of claim construction, as it reads into the claim

language a phrase that is altogether absent from the specification. In SunRace Roots Enterprise Co, Ltd. v. Sun Victory Trading Co., Inc., 336 F.3d 1298 (Fed. Cir. 2003), for example, the Federal Circuit overturned the district court's construction of "shift actuator," which held that the term required the use of a cam, even though no explicit definition of the term appeared in the specification. Rather, the district court imposed the limitation from the preferred embodiment, despite the fact that the invention could conceivably be embodied in a structure that lacked a cam. Id. at 1302. Here, the terms "subtree" or "partitionability" are not even present in the specification, and thus cannot even be plausibly read back into the claim language, which clearly does not include those terms. Cf. Interactive Gift Express, Inc. v. Compuserve Inc., 256 F.3d 1323, 1331 (Fed. Cir. 2001) ("If the claim language is clear on its face, then our consideration of the rest of the intrinsic evidence is restricted to determining if a deviation from the clear language of the claims is specified.").

Nor is plaintiff's assertion that "individual processor circuit cards" are themselves "binary tree computer systems" (thus justifying plaintiff's argument for partitionability) particularly persuasive. See Pl. Br. at 20-21. The specification itself indicates that "the binary tree computer system can . . . be constructed from a plurality of sub-units," '024 Patent at 5:53-55, and refers to specific processor cards -- a three processing element unit ("3-PEU") or a four processing element expansion unit

("4-PEXU"), id. at 5:57-58 -- from which a binary tree may be built, which plaintiff marshals as evidence that subtrees are computer systems. But nowhere in the patent are these cards referred to as "subtrees," and thus plaintiff's reliance on the 3-PEU or 4-PEXU cards as "subtrees" is misplaced: the patent's language clearly indicates that the system is constructed from the cards, not that each card is itself a binary tree computer system.

Accordingly, the Court hereby reaffirms its construction that "binary tree computer system" means a computer system of nodes connected in a binary tree configuration.

2. binary tree configuration

The disputed term appears in Claim 1, which indicates that "N bus controllers [are] connected in a binary tree configuration in which each bus controller, except those at the extremes of the tree, are connected to left and right child bus controllers . . . ." '024 Patent at 7:11-16.

The parties essentially argue for constructions that track their respective constructions of "binary tree" from the '201 Patent. Plaintiff's construction is "a tree arrangement where a node has a parent (except for a root node) and zero, one or two children nodes," Pl. Reply at 15, and defendant puts forward "an arrangement of nodes where each node has a single parent and two children nodes, except the root node, which has no parent, and the leaf nodes, which have no children," Def. Br. at 19. Here, the claim language clearly indicates that the binary tree configuration, as contemplated by the

patent, has bus controllers with two children bus controllers except at the extremities. See Phillips, 415 F.3d at 1314 ("[T]he claims themselves provide substantial guidance as to the meaning of particular claim terms."); Catalina Mktg. Int'l v. Coolsavings.com, Inc., 289 F.3d 801, 807 (Fed. Cir. 2002) ("Claim language defines claim scope."). This construction is confirmed by the specification. See Interactive Gift Express, 256 F.3d at 1331 (noting that if claim language is clear, specification is read only to determine if deviation from claim language is specified); see also Phillips, 415 F.3d at 1316 ("[T]he specification necessarily informs the proper construction of the claims."). The specification outlines the "basic structure of the binary tree parallel computer system of the invention" and specifies that "each node . . . is also connected downstream to its own [processing element] and either to two child nodes . . . , or in the case of the nodes at the extremes of the tree, to right and left leaf [processing elements]. '024 Patent at 2:50-56.

Considering as well the Court's construction of "binary tree" from the '201 Patent, which patent is here incorporated by reference, see '024 Patent at 1:24-26, the evidence overwhelmingly supports the Court's construction of "binary tree configuration" to mean "an arrangement of nodes where each node has a single parent and two children nodes, except the root node, which has no parent, and the leaf nodes, which have no children."

3. host computer

As Claim 1 indicates, the binary tree computer system is "for connection to and control by a host computer." Patent '024 at 7:11-12. Plaintiff construes the term "host computer" to be "a computer connected to a network that provides access to that network." Pl. Reply at 19. Defendant, in turn, construes it to be "a computer that is connected to and controls the binary tree of bus controllers." Def. Br. at 19.

Plaintiff's construction relies primarily on Figures 1, 7, and 8 in the patent specification, asserting that the host computer pictured therein is "connected to a network of bus controllers," thus justifying the importation of the term "network" into the definition of host computer. Pl. Reply at 19. As in its construction of "binary tree configuration," however, plaintiff's effort to insert terms that do not appear elsewhere in the patent is unavailing. Cf. Catalina Mktg., 289 F.3d at 307. Rather, the Figures support defendant's construction, since they establish that the "network" is better described as a "binary tree of bus controllers." See Phillips, 415 F.3d at 1316 ("The construction that stays true to the claim language and most naturally aligns with the patents description of the invention will be, in the end, the correct construction."). The specification further adds that the "host computer [] generates instructions . . . to control the operation of the system." '024 Patent at 3:35-38.

Plaintiff argues that this construction, by referring to the instructions to the bus controllers, imports a limitation from dependent Claim 3 into Claim 1 and thus runs afoul of the doctrine of claim differentiation. Pl. Reply at 20. Claim 3 reads "the binary tree computer system of claim 1 wherein each of said bus controllers further includes means for interpreting instructions received from the host computer." Under the doctrine of claim differentiation, limitations from a dependent claim are normally not read into the independent claim. See Phillips, 415 F.3d at 1315 ("[T]he presence of dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim."). Here, however, Claim 3 is adding further detail about the bus controllers, not imposing a particular limitation on the host computer. Moreover, independent Claim 7 confirms that the host computer sends instructions to the bus controllers by defining the patent invention as "a binary tree computer system for connection to and control by a host computer" comprised of bus controllers in a binary tree configuration in which "each of said bus controllers includ[es] means for interpreting instructions received from the host computer." '024 Patent 8:11-12, 29-30. Thus, there is no doubt that Claim 7 does not treat the bus controllers as a dependent limitation, as plaintiff asserts. See Regents of the Univ. of Cal. v. DakoCytomation Cal., Inc., 517 F.3d 1364 (Fed. Cir. 2008) ("While it is true that dependent claims can

aid in interpreting the scope of claims from which they depend, they are only an aid to interpretation and are not conclusive. Indeed, the presumption created by the doctrine of claim differentiation is not a hard and fast rule and will be overcome by a contrary construction dictated by the written description or prosecution history.") (internal quotations marks and citations omitted).

Nor does the defendant's construction, which emphasizes the connection between the binary tree and the host computer, rely improperly on the preamble. A preamble is generally construed as a limitation "if it recites essential structure or steps, or if it is 'necessary to give life, meaning, and vitality to the claim.'" Catalina Mktg. Int'l, Inc., 289 F.3d at 808(quoting Pitney Bowes, Inc. v. Hewlett-Packard Co., 182 F.3d 1298, 1305 (Fed. Cir. 1999). In Catalina Marketing, the Federal Circuit clearly found that the preamble of a patented invention, which referred to terminals located at "predesignated sites such as consumer stores," was not a limitation on the claim, since the patented invention did not depend on the location for the patent's significance or novelty and was not relied upon to distinguish prior art. Id. at 810. But that court also noted that where the preamble "is essential to understand limitations or terms in the claim body" or "when reciting additional structure or steps underscored as important by the specification," it may limit claim scope. Id. at 808. Here, the term "host computer" appears in both the preamble and the body of Claims 1, 3, and 7, and

therefore reliance on the preamble does not unduly limit the scope of the claim language. See id. ("[D]ependence on a particular disputed preamble phrase for antecedent basis may limit claim scope because it indicates a reliance on both the preamble and claim body to define the claimed invention."); Bell Commc'ns Research v. Vitalink Commc'ns Corp., 55 F.3d 615, 620 (Fed. Cir. 1995) ("[W]hen the claim drafter chooses to use both the preamble and the body to define the subject matter of the claimed invention, the invention so defined.").

In light of the unambiguous claim language and support from the specification, the Court construes "host computer" to be a computer that is connected to and controls the binary tree of bus controllers.

4. bus controllers

This disputed term appears in all the claims except dependent Claims 6 and 10. Claims 1 and 7 establish that the bus controllers are connected in a binary tree configuration, and each bus controller includes "a buffered interface connecting said processing element to said bus controller for transmitting instructions and data between the bus controller and the connected processing element" (Claim 1), '024 Patent 7:29-32, and "means for interpreting instructions received from the host computer and for executing such instructions addressed to it and for passing instructions to bus controllers down the tree. . . ." (Claim 7), id. at 8:29-33.

Plaintiff asserts that the term is not generally known in the art of computing but that "buffers" are, and that, when such meaning is combined with the "means of writing information," the construction of "bus controller" is "clear." Pl. Reply at 23. Despite that assertion of clarity, however, the construction plaintiff then offers is, on its face, circular and confusing, as it uses the term "bus controller" in the definition of the same: "A buffered interface connecting the processing element to the bus controller that enables writing information into the memory of the processing element without involving the microprocessor of said connected processing element." Pl. Br. at 21.[5] Even though it borrows phraseology directly from the claims, plaintiff's construction conflates the buffered interface, which is part of the bus controller, with the bus controller itself.

At the Markman hearing, plaintiff offered the unpersuasive argument that "is" and "includes" were synonymous, and therefore its construction meant that a bus controller "includes" a buffered interface. See tr., 08/20/09, at 143-45. The Federal Circuit has held that "including" is a broad term that refers to essential elements but permits additional elements to be added, Lucent Techs., Inc. v. Gateway, Inc., 525 F.3d 1200, 1214 (Fed. Cir. 2008); but

[5] Prior to submitting its opening brief, plaintiff had indicated to defendant that its construction was "a buffer interface connecting the processing element to the bus controller and the means for writing information into the memory of the connected processing element without control of the main processor of the connected processing element." See Def. Reply at 20 n.13.

plaintiff can provide no case that supports the premise that "is" can be read more broadly to mean "includes." Moreover, even assuming arguendo plaintiff's tortured linguistic alchemy permitted "is" to mean "includes," plaintiff's construction would still suffer from circularity by relying for its definition on the very term it seeks to define.

In contrast, defendant defines the plural "bus controllers" as "controllers that transfer instructions and data from the host computer to the connected processing elements, and data from the connected processing elements to the host computer over a bus." Defendant's construction also fairly tracks Claim 1's language, as well as Claim 7's language which says that the bus controllers "includ[e] means for interpreting instructions received from the host computer . . . and for passing instructions to bus controllers down the tree." Def. Br. at 22. The patent specification provides support for this construction, which reads: "The BCxs [bus controllers] act as buffered repeaters that transfer Function Calls and data from the Host Computer to the selected PE(s), and data with its Fault Message from the selected PE to the Host Computer." '024 Patent at 2:64-3:1. In short, the specification tells us that the bus controllers, by acting collectively as a "bucket brigade," transfer data and instructions from the host computer to the processing elements and from processing elements to other processing elements. Id. at 2:38-40. So too does the abstract provide that the

bus controllers "include[] a buffered interface . . . for transmitting instructions and data between the bus controller and the processing element. . . ." '024 Patent Abstract.

Accordingly, the Court construes "bus controllers" to be controllers that transfer instructions and data from the host computer to the connected processing elements and data from the connected processing elements to the host computer over a bus.

5. root bus controller

Claim 1, as noted above, indicates that when the computer system of the patent comprises more than 2 bus controllers, one of the bus controllers is a "root bus controller" that connects the remaining bus controllers, which are in a binary tree configuration, to the host computer. Plaintiff construes this disputed term to mean "any bus controller that is the highest level bus controller in the tree or subtree," Pl. Br. at 21, thus incorporating its view that "binary tree computer system" includes a partitionable portion such as a "subtree." Pl. Reply at 21.[6] Defendant asserts that the term is best defined as "one bus controller at the highest order position of the binary tree computer system that connects the binary tree to the host computer and which has no parent bus controller." Def. Br. at 23; Def. Reply at 23. Both parties thus essentially agree that

[6] Prior to submitting its opening brief, plaintiff had indicated to defendant that its construction was "a bus controller that is designated the highest level bus controller in the tree of subtree." See Def. Reply at 23 n.14.

the root bus controller is the "highest order" bus controller in the system, and therefore has no parent bus controller. Claims 1 and 7 also clearly indicate that the root bus controller as the link between the binary tree of bus controllers and the host computer. Hence, the principal dispute is whether a bus controller of a portion of the system could be designated the root bus controller.

Since the Court has already construed "binary tree computer system" to not include a partitionable portion thereof, such as a subtree, plaintiff's arguments emphasizing that a binary tree can be divided into subtrees is unavailing, see Pl. Reply at 22-23, as there is no support in the patent for the idea that the parent node of a subtree could be designated a "root bus controller." That the root bus controller is located at the highest order of the binary tree computer system itself is unequivocally established by both Claims 1 and 7, which specify that there be "one" root bus controller. '024 Patent at 7:17; 8:18. The specification also repeatedly refers to "one" bus controller that connects the binary tree to the host computer. See id. at 1:52-53; 2:41-42; and 2:51-53.

Plaintiff cites Professor Stolfo's testimony as extrinsic evidence that "every node can be a root." Pl. Reply at 24, Ex. L, at 190. Although the court may consider expert testimony in claim construction, see Phillips, 415 F.3d at 1318, here plaintiff seeks to take Professor Stolfo's assertion that every node can be a root out of its original context of a discussion of a binary tree

configuration and extrapolate from it support for the notion that any bus controller can also be a root bus controller. This contradicts the portion of the written record that the root bus controller links the binary tree to the host computer. '024 Patent at 7:18-19. Arbitrarily designating a particular bus controller to be a "root bus controller" would violate this requirement, and thus the Stolfo testimony, as applied in this context, must be discounted. See Phillips, 415 F.3d at 1318 ("[A] court should discount any expert testimony 'that is clearly at odds with the claim construction mandated by the claims themselves'. . . .") (citation omitted).

Accordingly, the Court hereby reaffirms its construction of "root bus controller" as the bus controller at the highest order position of the binary tree computer system that connects the binary tree to the host computer and which has no parent bus controller.

Additionally, the parties have agreed upon the construction of a number of terms from the '201 and '024 patents. See Def. Br, Ex. A, at 1-4; see also tr., 08/20/09, at 155. Several of them are "means-plus-function" terms, which must be interpreted by the Court in accordance with 35 U.S.C. § 112, ¶ 6. The Court hereby adopts the parties' jointly agreed upon constructions of the "means-plus-function" terms.

Finally, by stipulation dated September 15, 2009, the parties agree that, upon issuance of the instant Opinion, the Court may enter final judgment in favor of defendant and dismiss defendant's

counterclaims without prejudice as moot, thus reserving for appeal only the claim constructions here set forth. Accordingly, the Clerk of the Court is directed to close all open docket entries and to enter final judgment in favor of defendant, dismissing the complaint with prejudice and dismissing defendant's counterclaims without prejudice.

SO ORDERED.

JED S. RAKOFF, U.S.D.J.

Dated: New York, New York
January 5, 2010