IN THE UNITED STATES DISTRICT COURT

FOR THE WESTERN DISTRICT OF WISCONSIN

- -

SANDISK CORP., OPINION and ORDER

Plaintiff, 07-cv-605-bbc

v.

ZOTEK ELECTRONIC CO., LTD.,
ZODATA TECHNOLOGY LTD.,
KINGSTON TECHNOLOGY CO., INC.
and KINGSTON TECHNOLOGY CORP.,

Defendants.

- -

SANDISK CORP., OPINION and ORDER

Plaintiff, 07-cv-607-bbc

v.

ZOTEK ELECTRONIC CO., LTD.,
ZODATA TECHNOLOGY LTD.,
KINGSTON TECHNOLOGY CO., INC.,
KINGSTON TECHNOLOGY CORP.,
IMATION CORP., IMATION ENTERPRISES
CORP., MEMOREX PRODUCTS, INC. and
ADD-ON TECHNOLOGY CO.,

Defendants.

- -

In this patent infringement suit, the parties seek construction of claim terms from 5 different patents, U.S. Patents Nos. 6,757,842 ('842 patent); 6,149,316 ('316 patent); 5,719,808 ('808 patent); 6,426,893 ('893 patent); and 6,763,424 ('424 patent). The parties have not sought construction of any of the terms in U.S. Patent No. 7,137,011 ('011 patent), another patent plaintiff is asserting in this case. (The parties stipulated to the meaning of several terms in that patent, dkt. #543.[1])

The patents in suit all relate to flash memory technology, and in particular, technology related to flash EEPROM (Electrically Erasable Programmable Read Only Memory). As the parties explain in detail, flash memory has limitations related to the fact that the memory cells used to store bits must be "flash erased" rather than one-by-one and the cells wear down after a time. The patented technology involves methods for dealing with these limitations by improving recording and erasure efficiency and decreasing wear. The parties break the challenged patents into three main subgroups: the '842, '316 and '808 patents are called the "Flash EEPROM Patents" and relate to methods for storing "overhead" data (not user data but rather data used to manage the user data and the flash drive) with

[1] This lawsuit involves two cases, 07-cv-605-bbc and 07-cv-607-bbc. The two cases have been consolidated and the relevant documents filed in each case are identical. All references to the docket are for case no. 07-cv-605-bbc.

user data and erasing smaller "sectors" of larger "blocks" or "chips." The '893 patent is an improvement on the earlier patents; it relates to a method for breaking up the overhead data into different groups and storing less variable sorts of data in a different sector from that used for user data and the types of overhead more closely related with the user data (because they are more variable). The '424 patent relates to different methods for performing "partial-block updates" by allowing "rewrites" to occur without erasure and then just keeping track of which of the identical data is the "rewritten" data. This may be accomplished by either storing information about the "timing" of the "pages" of data and ignoring the "older" data with the same "logical address" or by storing newer data after older data and ignoring any page of data having the same "logical address" as data already read.

Claims to be construed:

In the court's order granting in part the parties' motions to construe certain claim terms, thirteen terms were to be construed. Since then, the parties have reached agreement on three of those terms (and other terms not at issue), leaving ten terms in dispute.

A. Terms from the Flash EEPROM patents ('842, '316 and'808 patents):

1. "operating individual blocks of memory cells with non-overlapping portions thereof storing at least user data and overhead information";

2. "linking the address of such unusable blocks with addresses of other blocks that are useable";

3. “an address in a format designating at least one mass memory storage block” and “a mass memory storage block address”;

4. “an array of EEPROM cells”; and

5. “designating a combinations of a plurality of but less than all of said multiple sectors to be erased.”

B. Terms from the ‘893 patent:

1. “individual ones of the redundancy codes being appended to ends of the user data from which they are generated”; and

2. “information of the characteristics of said first group of blocks.”

C. Terms from the ‘424 patent:

1. “recording a relative time of programming the at least one page of new data and the at least one page of superseded data”;

2. “programming individual ones of a first plurality of said given number of pages and a logical page address associated with the original data”;

3. “programming individual ones of a second plurality of a total number of pages less than said given number in a second block with updated data and a logical page address associated with the updated data.”

The parties propose definitions for each term, but as is often the case at this stage of the proceedings, their disputes relate to the scope of the claims, not the ability of a juror to understand the language. The parties’ disputes make it clear that the specific language they propose for each term is simply a vehicle for arguing for a broader or narrower construction of the claim terms at issue. In my experience, attempting to resolve the parties’ disputes by

providing specific definitions to a given claim term is nothing but an invitation to a new round of arguments at a later stage about the meaning of the court's construction, or about the hidden implications of the language adopted. It is counterproductive to resolve claims construction disputes by replacing them with new ones for the parties to dispute about at summary judgment.

Although there may be occasions in which parties seek a construction for the sake of improving jury understanding, this is not one of them. Thus, at this stage I will resolve only the parties' disputes and leave to one side the question of what specific language should be adopted. If concerns arise before trial about jury confusion, the parties can request specific language in limine.

Before turning to the parties' disputes, I note that in several instances, defendants cite the International Trade Commission's construction of the claims in support of its position. However, each time they do so, they fail to point to the Commission's reasoning. The rulings of the Commission are not binding on this court, and because defendants have not identified the Commission's reasoning, they are also unpersuasive.

A. Terms from the Flash EEPROM Patents ('842, '316 and'808 patents)

1. "operating individual blocks of memory cells with non-overlapping portions thereof storing at least user data and overhead information" ('842 pat., cls. 1, 10)

For this claim term, the parties have two disputes: (1) whether the claimed "operation" can be on "one or more" erase blocks or must be on "multiple erase blocks"; and (2) whether there can be only one user data and one overhead portion in an "individual block."

The first issue comes down to a simple question of interpreting the meaning of "operating individual blocks." Does it allow a reading that only one block may be operated upon, or does it require at least two? As defendants point out, the term says "individual," but that word modifies "blocks," which is plural. Defendants argue that in ordinary speech a person might say she is meeting "individual people" and mean only one person, but I find this unpersuasive. "Individual people" is still plural; at most "individual" suggests the "people" may be met one-by-one. At most, "individual" suggests that the multiple "blocks" may be "operated" upon one-by-one. The claim language resolves this dispute: "individual blocks" must be multiple erase blocks.

The second dispute relates to whether each of the claimed "blocks" may contain only one user data and only one overhead portion. Defendants contend that they must, pointing to the claim language itself. Hrg. trans., dkt. #584, at 62. Both claims 1 and 10 describe

"nonoverlapping portions . . . storing at least user data and overhead information" and later claim "writing data to, or reading data from" both "the user data portion of the . . . block" and "said overhead portion." The claim's reference to "*the* user data portion" and "*said* overhead portion" supports the proposed limitation.

Against this, plaintiff points to claim 16, a claim depending from claim 10. Claim 16 claims "[t]he method of claim 10, wherein the individual blocks include only one user data portion and only one overhead data portion." According to plaintiff, because claim 16 does nothing more than add a limitation to claim 10 that defendants contend is already present in claim 10, claim 16 would be superfluous under defendants' proposed construction. Under the doctrine of "claim differentiation," superfluous claims are to be avoided. ICU Medical, Inc. v. Alaris Medical Systems, Inc., 558 F.3d 1368, 1376 (Fed. Cir. 2009). However, the doctrine of claim differentiation "is not a rigid rule but rather one of several claim construction tools." Id. The doctrine of claim differentiation "cannot alter a definition that is otherwise clear from the claim language, description and prosecution history." O.I. Corp. v. Tekmar, 115 F.3d 1576, 1582 (Fed. Cir. 1997).

In this case, the meaning of the term is "otherwise clear." What the term "nonoverlapping portions . . . storing at least user data and overhead information" leaves unclear is made clear by the surrounding language ("*the* data portion" and "*said* overhead portion"). ACTV, Inc. v. Walt Disney Co., 346 F.3d 1082, 1088 (Fed. Cir. 2003) (court

must consider context of surrounding words of claim when construing term). By using the words "the" to refer to a single user data "portion" and "said" to refer to a single overhead "portion," the claim leaves no doubt that it covers a method involving only one user data portion and one overhead portion.

Although plaintiff contends that language in the specification supports a broader construction, the language it cites does not speak to the number of portions but states only that the "partitioning" between user and overhead portions "need not be rigid" and "the relative size of the various partitioned areas may be logically reassigned. Also the grouping of the various areas is largely for the purpose of discussion and not necessarily physically so." '842 pat., col. 8, lns. 52-57. Although the specification "is always highly relevant to the claim construction analysis," Phillips v. AWH Corp., 415 F.3d 1303, 1323 (Fed. Cir. 2005), in this instance it does not suggest that more than one user data or overhead portion is being claimed, and the claim language itself says otherwise.

In summary, the parties' disputes related to this term are resolved as follows: "operating individual blocks of memory cells with non-overlapping portions thereof storing at least user data and overhead information" must include multiple erase blocks and is limited to a single user data and a single overhead portion.

2. "linking the address of such unusable blocks with addresses of other blocks that are

useable" ('842 pat., cl. 1)

The parties' principal dispute is whether "linking" must be in done in the form of a map or table listing both the unusable and corresponding usable blocks, as defendants contend. A secondary dispute is whether, assuming "linking" does not require a map or table, it is broad enough to include storing a pointer in the unusable block that points to a usable block. (Plaintiff's proposed construction defines linking to include any sort of "substitution," but this term is too vague to be helpful.)

The language itself suggests nothing about the required format for linking addresses, saying only that the claimed method involves "linking the addresses of . . . unusable blocks with addresses of other blocks that are useable." Defendants contend that their construction is "the only construction" consistent with the claim language in light of the order of the following claimed steps: (1) detecting a predefined condition; (2) linking the address of the nonusable block to the address of a usable one; (3) causing the controller to generate an address; and (4) accessing a usable block by referring to the linked address "if the block with the generated address is unusable." '842 pat., cl. 1. As defendants point out, step four anticipates that the "block with the generated address" may be unusable, which means the controller must be able to generate an unusable address after the "linking" occurs.

However, defendants do not explain why the controller would do this only if a map or table is used. Instead, they argue only that the "linking" could not involve plaintiff's

proposed "substitution" because then the controller would not generate the unusable address after the "substitution" occurred. For defendants to be correct, "substitution" must mean "replacement" and they must be arguing that the claimed "replacement" would be carried out in a way that affects what the controller generates. This is not plaintiff's understanding of "substitution," but it is irrelevant. As explained above, the term "substitution" does not precisely capture plaintiff's concern and will not be adopted, so it is unnecessary to consider defendants' attacks on that term.

Defendants add that the specification supports the map or table requirement because the examples of linking that it provides involve a map or table. '842 pat., col. 8, lns. 49-51 ("The addresses of the defective cell and the backup cell are stored as defect pointers in the defect map 409."); id., col. 11, lns. 48-52 ("When the number [of bad memory cells] in a sector exceeds a predetermined value, the controller marks that sector as defective and maps it to another sector. The defect pointer for the linked sectors may be stored in a sector defect map."). However, as the Court of Appeals for the Federal Circuit has "repeatedly warned," the scope of a claim should not be limited simply because a specific embodiment shows the requested limitation. Phillips, 415 F.3d at 1323. Even if an embodiment is the only one disclosed, it may serve to limit a claim only if it is clear that the patentee intended to limit the scope of the claims to the disclosed embodiment. Id.; see also On Demand Machine Corp. v. Ingram Industries, Inc., 442 F.3d 1331, 1339-40 (Fed. Cir. 2006)

(limitation warranted because specification used the term "customer" repeatedly in specialized context); Nystrom v. TREX Co., Inc., 424 F.3d 1136, 1144-45 (Fed. Cir. 2005) (limitation warranted because written description and prosecution history used term "board" consistently to refer to wood decking materials cut from log). Defendants' cited evidence falls short of showing an intent to limit "linking" to using a map or table.

In addition, the limitation defendants seek to import into claim 1 is already found in claim 7 of the '824 patent, which depends from claim 1. Claim 7 discloses a form of "linking" that

> includes maintaining a list within the card that links such unusable blocks with addresses of corresponding ones of the other blocks that are useable, and wherein accessing a usable block includes referring to the list to translate the address of the unusable block into an address of a usable block.

Defendants point out that dependent claim 7 involves more than just using a map or table, but they do not deny that using a map or table is one of the features disclosed in this claim. The fact that the "map or table" requirement is spelled out explicitly in a dependent claim is strong evidence that the patentee did not intend the requirement to be a part of claim 1. If the requirement were already present in the independent claim, there would be no reason to repeat the limitation in the dependent claim. In short, defendants fail to support their contention that the claimed "linking" must be by map or table.

The fact that defendants' proposed limitation is unsupported does not mean that

plaintiff is right when it says that the usable and unusable addresses can be "linked" by storing a single pointer in the unusable block that points to a usable block. A natural reading of "linking" does not suggest such an indirect relationship. The claim does not call simply for linking a defective block to a good block, but rather for linking the *addresses* of the defective and good blocks. According to plaintiff, a "linking" occurs when a defect pointer containing the address of a good block is placed somewhere in a defective block. Once the defect pointer is in place, if the controller generates an address in the defective block, it will find the defect pointer in that block and move on to the good block referred to by the defect pointer. In this setting, the blocks are linked, but the addresses are linked only in the sense that the good block address stored in the defect pointer is stored in the defective block (located in an address in the defective block). Plaintiff's proposed construction would stretch the claim language beyond what is acceptable. The claim language says explicitly that the "addresses" of the blocks must be linked but plaintiff's proposed construction does not include that requirement.

In conclusion, "linking the address of such unusable blocks with addresses of other blocks that are useable" need not take the form of a map or table, but must be something more than simply storing a pointer in the unusable block that points to a usable block.

3. "an address in a format designating at least one mass memory storage block" and "a mass memory storage block address" ('316 pat., cls. 67 and 79; '842 pat., cls. 1, 10 and 61)

For this term, the parties have one major and one minor dispute. The minor dispute is whether the claimed address must "reference a block of user data" as plaintiff contends. In a sentence, defendants assert that "[t]here is simply no requirement in claim language" to support this limitation. Dkt. #536 at 20. Plaintiff responds in a footnote that elsewhere the claim language requires at least "user data" and "overhead data" and asserts that overhead data can be excluded because it is not "involved" in the method. Dkt. #559 at 16 n.5. The parties' cursory treatment of this matter suggests it is not a matter of importance to them, and may not even be related to a question of infringement or invalidity. At any rate, I am not persuaded that the limitation is required and decline to accept plaintiff's request to add it.

The major dispute is whether the mass memory storage block address must specify a "physical" address for the mass memory system or whether they could also specify a "logical" address. By "physical" address, the parties are referring to a description of a physical location within the mass memory system. A "logical" address describes the data without regard to its physical location; it is simply a number assigned to data to distinguish that data from other data stored in mass memory. In either instance, the system maintains

a way for "finding" the data (which is presumably why these physical and logical descriptions are both considered "addresses").

As plaintiff points out, the type of addressing used by a mass memory system (such as a hard disk) matters little to the functioning of a flash memory device. Regardless what sort of address the mass memory gives it, the flash memory device will have to "remap" the memory into its unique storage system. A mass memory physical address cannot become the physical address of the data stored in the flash memory. To take an example, data stored in a magnetic memory disk may be given a "physical address" according to the Cylinder, Head and Sector in which the data is located but a flash memory device does not have cylinders or heads. Therefore, it must convert the previous physical address into an address that reflects *its* own physical organization.

For plaintiff, this point weighs in favor of including both physical and logical addresses within the claim. Plaintiff adds that its construction is supported by both the "intrinsic record" of the '316 and '842 patents and the "understandings" of persons of ordinary skill in the art at the time of the inventions. In particular, plaintiff asserts that, by the time the patents were filed, it was "well understood" that disk drives could use either physical or logical addressing. Moreover, a prior art patent, U.S. Pat. No. 4,924,331 (the "Robinson patent"), stated that it was already "well known in the disk drive field to convert logical block addresses into cylinder head and sector addresses as most seek commands are

initially input as logical block addresses."

Defendants acknowledge that logical addressing in mass memory storage was known at the time the patents at issue were filed but attempt to distinguish the Robinson patent (which also suggests that logical addressing was common). For the sake of argument, I assume defendants are correct in saying that logical addressing would not have been common at the time of the invention. Even so, plaintiff's point is valid: it would make sense for a flash memory device to receive both logical and physical addresses.

That is not enough to say it does. The starting point in construing any term is the language of the claim itself, which in this instance is "an address *in a format* designating at least one mass memory storage block." (The parties agree that this language is interchangeable with the term "mass memory storage block address.") According to defendants, plaintiff's construction would make the phrase "in a format" superfluous because format falls out of the picture under that construction.

Defendants may be right that logical block addressing is not regularly considered a mass memory storage "format" even if it can be used in that setting. Nonetheless, I am not persuaded that the phrase is as empty as defendants say it is under plaintiff's reading. Leaving out the phrase and claiming "an address designating at least one mass memory storage block" tends to suggest that the address is actively "designating" the block while the phrase "in a format" helps emphasize that the address need be only of a sort designed to

work in a mass memory storage block. Moreover, nothing about the phrase or the purpose of the patents at issue suggests an intent to limit the patent to converting only one type of mass memory storage addressing. It would make no sense or produce any additional benefit to include such a limitation. Defendants do not suggest that the limitation was necessary to distinguish prior art. At most, the phrase "in a format" is inartful; it is not intended to impose a requirement that the address of the mass memory storage block be physical.

Thus, I conclude that "an address in a format designating at least one mass memory storage block" and "a mass memory storage block address" need not "refer to a block of user data" and may specify either a physical or a logical address for the mass memory system.

4. "an array of EEPROM cells" ('808 pat., cl. 16)

The parties disagree about whether the term "an array of EEPROM cells" requires the claimed memory storage elements to be "contiguous" or contain "dedicated row and column decoders." Defendants seek both limitations on the ground that the sole embodiment shows storage elements that are arguably "contiguous" (assuming that term is sufficiently clear to have meaning) and contains decoders labeled "column decoder" and "row decode." (The embodiment is from a separate application incorporated into the '808 patent by reference, Figure 4 from U.S. Patent Application No. 07/337,579, which became U.S. Patent No. 5,172,338).

Even if I assume that the sole embodiment includes contiguous elements and dedicated decoders, this does not necessarily mean that an "array of EEPROM cells" must always do so. As explained above, the mere fact that a sole embodiment includes certain features does not mean that the claimed invention is so limited. Something in the specifications must indicate that the embodiment was intended to demonstrate the contours of the invention, not just provide one example of the invention. Phillips, 415 F.3d at 1323. Defendants do not point to any intrinsic evidence suggesting the inventor intended to limit "an array of EEPROM cells" to the embodiment it incorporated by reference.

Defendants point out that the term is not described in the specification, which is true, but that does not mean the embodiment must be the sole source for interpreting the term. Indeed, the starting point is the "ordinary and customary meaning" of the terms themselves, as understood by a person of ordinary skill in the art at the time of the invention. Id. at 1312-14. The cited embodiment may suggest that, at the time of the invention, an "array" in the relevant field included dedicated row and column decoders and contiguous memory storage elements. However, this does not suggest that either of these features was a requirement of an array. Indeed, as plaintiff's expert avers, at the time of the invention (the late 1980s), "it was well-known in the semiconductor memory industry to divide a memory array into physically distinct units called 'sub-arrays' or 'subarrays.'" Taylor Decl., dkt. #540, ¶ 12. This matters because a memory system using a sub-array might not have "contiguous"

memory storage elements throughout the array and may have dedicated decoders at the sub-array level as opposed to the array level, as demonstrated in a later patent of plaintiff's using sub-arrays:

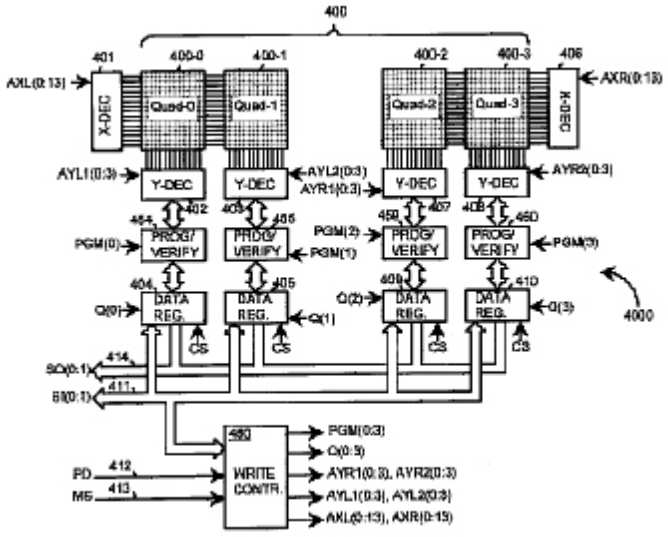

Dkt. #60-5, U.S. Pat. No. 5,890,192, Fig. 3. (Although each separate group of memory storage elements is labeled as a "quad" in Fig. 3, elsewhere the patent describes the group of storage elements as a "plurality of flash EEPROM cells 400 organized into four sub-arrays or quadrants, 400-0 to 400-3." Id., col. 3, lns. 57-58.)

Plaintiff's expert has more than his say-so to back this up; he identifies three patents from the time of the invention that included arrays subdivided into sub-arrays: U.S. Pat. No. 4,694,433, Abstract ("[a] memory structure for very large memory arrays on a chip . . . where the memory array is divided into a number of subarrays,); U.S. Pat. No. 4,758,993, Abstract

("having an array of memory cells which are divided in several sub-arrays."); U.S. Pat. No. 4,807,191, col. 1, lns. 13- 15 ("[i]ntegrated circuit memories are frequently characterized by being divided into sub-arrays, or blocks, of memory cells."). Defendants point out that none of the patents involve flash EEPROM memory, but do not explain why the term "array" would mean something different in the flash memory field from what it would mean in other semiconductor memory devices. Even if defendants are correct that at the time of the invention "sub-arrays" had not been contemplated in the flash memory context, that would not mean an "array" must be limited in the ways that defendants propose. Only if "arrays" in the field of flash memory *excluded* sub-arrays would there be a ground for imposing the limitations defendants request. Defendants have no evidence of this, so the limitations will not be imposed. "An array of EEPROM cells" need not be "contiguous" or contain "dedicated row and column decoders."

5. "designating a combinations [sic] of a plurality of but less than all of said multiple sectors to be erased" ('808 pat., cl. 16)

For this claim term, the parties disagree about whether the claimed "designation" requires setting a tag in a dedicated register for each sector. Defendants say it must, relying in part on their view that the claim requires the ability to designate "any combination" of sectors. The parties disagree on whether this limitation applies.

As to the question whether the device must be capable of setting "any combination" of sectors for erase, defendants have the better argument. The claim language calls for only "a combinatio[n]" of more than one but less than all sectors. However, at times language used in the specification may require a particular limitation to be read into the claim language, such as when the specification includes "repeated and definitive remarks" that a particular limitation applies to the claims. Computer Docking Station Corp. v. Dell, Inc., 519 F.3d 1366, 1374 (Fed. Cir. 2008). Moreover, describing a limitation as part of "the present invention" or "the invention" is strong evidence that the claims should be so limited. Trading Technologies International, Inc. v. eSpeed, Inc., 595 F.3d 1340, 1353-54 (Fed. Cir. 2010); see also Honeywell International, Inc. v. ITT Industries, Inc., 452 F.3d 1312, 1318 (Fed. Cir. 2006) (four references to a fuel filter as "this invention" or "the present invention" warranted limiting the invention to a fuel filter). This case includes precisely such limiting language in the specification. For example, the specification states that "[t]he invention allows any combination of sectors among the chips to be selected and then erased simultaneously," '808 pat., col. 1, lns. 65-66, and "allows any combination of sectors selected for erase to be deselected," id., col. 2, lns. 3-4. The specification later repeats this sentiment:

> In the present invention, the Flash Eeprom memory is divided into sectors where all cells within each sector are erasable together. Each sector can be addressed separately and selected for erase. One important feature is the ability to select any combination of sectors for erase together. This will allow for a

> much faster system erase than by doing each one independently as in prior art.

Id., col. 4, lns. 51-55. Plaintiff responds that the specification describes the "present invention" in permissive terms ("the present invention allows"), so it should not be considered a limitation. This is a misreading of the specification. The cited language does not suggest that the present invention *may* include the "any combination" feature, it says it *does*. What makes the language appear permissive is the fact that the feature itself is a description of what the system is *capable* of doing (selecting any combination), not what it is actively doing (selecting a given combination). Because the specification repeatedly describes "the present invention" as one capable of selecting any combination of sectors for erase together, this limitation must be part of the claim.

It is a separate question whether the claimed designation requires setting a tag in a dedicated register for each sector, as defendants contend. According to defendants, this limitation is warranted both because there is no other way to insure erasure of "any combination" of sectors and because the only embodiment disclosed in the patent uses designated registers. As to the first point, defendants fail to establish that tagging designated registers is the only possible method for selecting any combination of sectors. More important, there is no need to decide this question. If defendants are correct, the question of infringement of their products should rise or fall according to whether they are capable

of designating "any combination" of sectors.

As for the argument that the only embodiment disclosed uses registers, defendants fail to tie the embodiment to any intent to limit the claims to the contours of the embodiment. Indeed, the specification refers to the embodiment, Fig. 3A, as merely "illustrat[ing]" the claimed feature. Defendants add that the dedicated erase register was a "novel inventive concept" identified by an inventor of the '808 patent, but this is irrelevant; the inventors received a patent describing the invention in broader terms.

Moreover, as plaintiff points out, requiring designation to include setting tag bits in a dedicated register would make some of the language in dependent claim 18 superfluous. Claim 18, which depends (indirectly) from claim 16, adds "setting a tag bit for individual ones of the sectors to be erased"; if claim 16 already included this limitation, it would make no sense to repeat this limitation in claim 18.

In conclusion, "designating a combinations [sic] of a plurality of but less than all of said multiple sectors to be erased" requires the ability to designate "any combination" of sectors for erase but does not require setting a tag in a dedicated register for each sector to be erased.

B. Terms from the '893 Patent

1. "individual ones of the redundancy codes being appended to ends of the user data from which they are generated" ('893 pat., cls. 13 and 58)

For this term, the parties' dispute relates to what is required for the claimed "redundancy code" (error correcting code) to be "appended to the ends of" the user data. Defendants contend that the claimed redundancy code must be attached immediately adjacent to the user data and plaintiff says the only limitation is that "no other user data" may intervene between the end of the user data and the redundancy code generated from that data.

The plain meaning of the language controls here. As plaintiff points out, the word "appended" could mean "attached to" without requiring immediate adjacency. Think of a set of exhibits "appended to" a brief; not all of the exhibits are immediately adjacent, but any of them could be said to be "appended to" the brief. The problem is that the claim language is more specific: the phrase is "appended *to the ends of*." The added phrase "to the ends of" emphasizes exactly where the code must be attached.

Plaintiff contends that the phrase "appended to the ends to" has been used loosely to indicate only that data comes *after* other data. In particular, U.S. Pat. No. 5,551,020 (the Flax patent) discloses information "appended to the ends" of a string and contrasts it with information "appended to the beginning" of the string. There are two problems with this argument. First, plaintiff fails to point to anything in the Flax patent suggesting that the

attachment "to the ends" and "to the beginning" was not attached immediately to the ends. Second, even if plaintiff could show the term was used loosely in that patent, the circumstances were different; there the patent contrasted data "appended to the ends" with data "appended to the beginning," suggesting that the phrase "to the ends" served the purpose of emphasizing at which location the data was stored. In this case, that added emphasis is not needed because there is no contrast of the data with data stored elsewhere.

Plaintiff's next argument is that the claim cannot require immediate adjacency because that would interfere with another feature claimed in the invention related to "fill bytes," which are used to fill in defects in the block. Figure 4 of the '893 patent shows an example of how a relationship between user data (labeled "DATA" in the figure) and redundancy code (labeled ECC in the figure) would appear in a "perfect world."

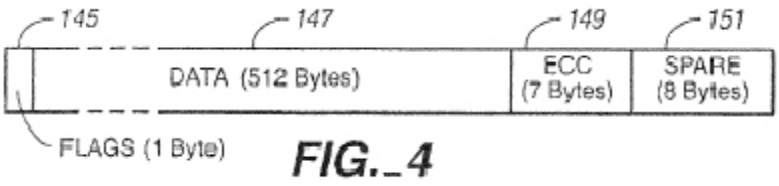

FIG._4

The "SPARE" bytes represent the space that the data may shift if fill bytes are inserted at places in the data "to avoid bad columns." '893 pat., col. 13, lns. 26-29. According to plaintiff, fill bytes could be inserted anywhere, including at the very beginning

of the redundancy code. If so, the redundancy code would no longer be immediately adjacent to the user data.

Plaintiff's theory has two problems. First, plaintiff does not explain why the "fill byte" feature of the invention would have to work alongside the feature requiring the redundancy code to be "appended to the ends" of the user data. The only claim plaintiff identifies that includes the "fill byte" feature is claim 60, which depends from claim 52. The claim at issue here is claim 58, which also depends from claim 52. Plaintiff does not explain why these separate dependent claims must be construed in a way that makes them both work together and no principled basis for doing so suggests itself.

Second, even if the present claim language should be construed in a way to accommodate the "fill byte" feature, it would not warrant leaving the construction as wide open as plaintiff proposes (to allow anything but user data to intervene). At most, such an accommodation would allow the redundancy data to be attached either directly to the user data or to fill bytes that are attached directly to the user data. Plaintiff has not indicated it would be satisfied with this accommodation, so adopting it seems unjustified. I conclude that "individual ones of the redundancy codes being appended to ends of the user data from which they are generated" requires the redundancy code to be attached immediately adjacent to the user data.

2. "information of the characteristics of said first group of blocks" and "characteristics of the memory cell blocks" ('893 pat., cls. 1 and 65)

Both claims at issue, claims 1 and 65 of the '893 patent, require separating data involving "characteristics of" the blocks of user data and storing it separately. The parties disagree about which sorts of addresses are considered "characteristics of" the user data blocks, and therefore which addresses must be stored separately from user data. According to plaintiff, the only addresses that are "characteristics of" the user data blocks are "good block addresses," a physical address of a substitute block. Defendants contend that the claim language covers any "block" address, including logical addresses.

The starting point for determining the scope of the claim language is, of course the language itself. Claim 1 splits "a first group of . . . blocks for storing user data" from "a second group of . . . blocks for storing information of the characteristics of said first group of blocks" and explains what is stored in each group of blocks:

> storing, in individual ones of the first group of said blocks, user data plus characteristics of the user data being written therein but not including characteristics of said first group of blocks, and
>
> storing, in individual ones of the second group of said blocks, a plurality of records of characteristics of individual ones of the first group of blocks but without storing either user data or characteristics of the user data into the second group of blocks.

'893 pat., cl. 1. Thus, claim 1 recognizes three distinct types of data: (1) user data; (2) characteristics of the user data; and (3) characteristics of the blocks storing the first two

types of data. (Claim 65 simply requires that "the sectors of user data stored in the memory cell blocks . . . not include characteristics of the memory cell blocks in which they are stored.")

Logical block addressing does not fit well in any of the three categories. It is not user data. As defendants point out, a logical block address "remains the same irrespective of the user data" and does not otherwise describe the user data. Thus, it would be a stretch to call a logical address a "characteristic" of the user data any more than a physical address would be.

At the same time, it would be a stretch to call a logical address a "characteristic" of the user data block as well. As defendants acknowledge, "[l]ogical addresses are by definition generic and specify nothing on their own." Defs.' Resp. Br., dkt. #555, at 24. A logical block address does not describe the user data block.

The description and examples provided in the specification suggest that "characteristics" of the user data block relate to its *physical* properties. The specification explains that the relevant data includes "[o]verhead data of the condition, characteristics, status, and the like, of the individual blocks," including "an indication of how many times the block has been programmed and erased, voltage levels to be used for programming and/or erasing the block, whether the block is defective or not, and, if so, an address of a substitute good block, and the like." Id., col. 2, lns. 56-63. Later, the specification distinguishes this

type of overhead data from the type that may be stored with user data, explaining that "the overhead information that is stored in a block along with a sector of data is limited to information about the data itself and does not include physical overhead information about the block or its operation," id., col. 13, lns. 61-64, while "information about the physical block is stored in another block," id., col. 14, lns. 6-7. Additional examples of "this type of information about the physical block" include "experience cycles, numbers of pulse or voltages required to program or erase the block of cells, defects within the block, and like information," "flags including an indication that the user data block is a good one," voltage for programming or erasing the user data block and data indicating the wear of the user data block, id., col. 14, lns. 1-4, 35-47, Fig. 9.

One example in the specification is silent about whether the "characteristic" involves a physical characteristic of the block. Figure 10 illustrates "an overhead record for a user data block that has exceeded its useful lifetime." Id., col. 14, lns. 66-67. The figure shows a block of data containing "FLAGS," and "SPARE UNIT, BLOCK ADDRESS." The specification explains that the "SPARE UNIT, BLOCK ADDRESS" portion of Figure 10 specifies "the spare block's address." The specification says nothing about whether the address must be physical.

Although logical block addressing does not fit well in any of the three possible groups of data, it *must* fit somewhere, which means the question is where it fits best. I conclude that

the best fit is with the group including "characteristics of" the blocks. Although the specification emphasizes that the data tends to relate to physical characteristics of the block, it refers to "addresses" generally in its example and does not suggest any reason for excluding logical addresses, which are at least indirectly about the physical block (each logical address is ultimately tied to a physical location when the controller "translate[s the] logical block address to a physical one." Id., col. 18, ln. 30.

Moreover, treating logical block addresses like physical ones makes sense in light of the stated purpose of the invention, which includes avoiding "frequent rewriting of the overhead data, each time the user data is rewritten into the block" and "reduc[ing] the amount of time necessary to access and read the block overhead data." Id., col. 3, lns. 3-7. Because a logical address is tied to a piece of user data in much the way a physical address would be and enjoys the same sort of stability, it would make sense to store logical addresses outside the user data block. I conclude that both logical and physical addresses are "characteristics of" user data blocks.

C. Terms from the '424 patent

1. "recording a relative time of programming the at least one page of new data and the at least one page of superceded data" ('424 pat., cls. 1 and 3)

For this term, the parties disagree about whether the time of programming a page must be recorded on that page and allow a comparison of the page to pages written both before and after the page (rather than just before it). In addition (and related to their other disputes), the parties disagree about whether the claim is broad enough to include a "block" recording method described in the specification:

> A second specific implementation of the inventive technique can also be described with respect to Fig. 8. In this example, the time stamp is used only to determine the relative age of the data stored in blocks, while the most recent pages among those that carry the same LBN [logical block number] and page number are determined by their relative physical locations. The time stamp 43 then does not need to be stored as part of each page. Rather, a single time stamp can be recorded for each block, either as part of the block or elsewhere within the non-volatile memory, and is updated each time a page of data is written into the block. Data is then read from pages in an order of descending physical address, starting from the last page of the most recently updated data pages having the same LBN.

'424 pat., col. 9, lns. 40-53.

As to this last dispute, plaintiff contends that this sort of block recording is included in the claim. If plaintiff is correct, then neither of the two disputed limitations could be imposed. If block recording is part of the claim, then the time of programming need not be

recorded on the page and there is no need to allow a comparison of the page to pages after the recorded page.

However, the claim language does not support plaintiff's reading. The block recording method does not involve recording a time for each page, but instead records only one time for a given block. The relative newness of each *page* is determined by other means (by relative physical location). The claim requires "recording a relative time of programming" the *pages*. The claim says "relative," so there is an argument that a time need not be recorded for each page; even so, the "relative time" that must be recorded must compare the newness of pages, not blocks.

Plaintiff points out that excluding the block recording method from claim 1 would conflict with claim 8, which depends from claim 1. Claim 8 adds the requirement that the data programmed in claim 1 must be programmed "into the first available unused pages within [a given block] in a predefined order." According to plaintiff, because the block recording method is the one that uses the "predefined order" of pages to determine newness, that method must be a part of claim 1 from which claim 8 is located. Plaintiff is mistaken. The "predefined order" limitation in claim 8 would be more useful if the block recording method were a part of claim 1, which suggests that claim 1 was supposed to include that method. However, narrower language was used in claim 1 than would allow for the block recording method. The fact that claim 8 will not be as useful without a broad reading of

claim 1 cannot overcome the plain language of the claim itself. Therefore, I conclude that "recording a relative time of programming the at least one page of new data and the at least one page of superceded data" cannot be performed by the block recording method described in the specification.

The next question is whether the claimed "recording" requires programming the relative time in the page being recorded and requires allowing a comparison of relative times of recording of pages both before and after the recorded page. Defendants contend that both of these limitations are required because the embodiments in the specification include such an in-page limitation (which would necessarily allow comparisons with pages both before and after), col. 8, lns. 35-50, and because the invention will not work properly without such page-by-page stamping. As to the first point, as explained repeatedly above, the presence of an embodiment showing one limitation or another is not enough by itself to warrant importing that limitation into the claims. Defendants identify nothing else to suggest an intent to limit the claims to the features of the embodiments. As to the second point, the law prohibits construing claims more narrowly just to fix an unworkable patent. Chef America, Inc. v. Lamb-Weston, Inc., 358 F.3d 1371, 1374 (Fed. Cir. 2004) ("[C]ourts may not redraft claims, whether to make them operable or to sustain their validity.") Defendants have failed to explain why their specific limitations should be imported into the claims, so they will not be. In summary, "recording a relative time of programming the at least one page of new data

and the at least one page of superseded data" is not satisfied by performing the block recording method described in the specification, but it does not require recording the relative time of programming on the page being recorded or allowing a comparison of the page to pages written after the page.

2. "Logical page address" ('424 pat., cls. 20 and 24)

The parties disagree about whether the claimed "logical page address" is limited to a "logical block number plus logical page offset." (The parties seek construction of two different claim terms referring to the "logical page address": (a) "programming individual ones of a first plurality of said given number of pages and a logical page address associated with the original data"; and (b) "programming individual ones of a second plurality of a total number of pages less than said given number in a second block with updated data and a logical page address associated with the updated data." There are no other disputes related to the first term and only one other dispute related to the second, which is discussed in the section below.)

Defendants contend that the embodiments specify only a "logical block number plus logical page offset" and therefore that the claims should be so limited. Plaintiff's attempt to show that the embodiments involved other types of "logical page addresses" is unpersuasive, but ultimately this does not matter. Once again, embodiments alone do not create

limitations. Defendants identify no evidence that the claims were intended to be limited to the embodiments listed. The term "logical page address" is not limited to a "logical block number plus logical page offset."

3. "programming individual ones of a second plurality of a total number of pages less than said given number in a second block with updated data and a logical page address associated with the updated data" ('424 pat., cls. 20 and 24)

Aside from their disagreements about a "logical page address," the other issue the parties have relates to whether the claimed "programming" must occur "without marking the original (now superseded) pages of data in the first block with an invalid data flag." Plt.'s Resp. Br., dkt. #559, at 40. This comes down to whether there was prosecution disclaimer. According to plaintiff, it disclaimed marking with data flags during prosecution so the claim must perform the step without so marking. In their briefs, defendants argue that the cited disavowal does not rise to the level of disclaimer and cannot limit the meaning of the term. However, at the claims construction hearing, they changed course, acknowledging twice that plaintiffs did disclaim the "flagging" feature that they are now seeking as a limitation. Hrg. trans. at 103 ("that's right, they did disclaim it"); id. at 117-18.

Because defendants have agreed that disclaimer occurred, they have no ground for arguing that the limitation should be imposed. Therefore, I conclude that "programming individual ones of a second plurality of a total number of pages less than said given number

in a second block with updated data and a logical page address associated with the updated data" must occur "without marking the original (now superseded) pages of data in the first block with an invalid data flag."

ORDER

IT IS ORDERED that the following determinations are made about the following claim terms:

1. As used in claims 1 and 10 of U.S. Pat. No. 6,757,842, "operating individual blocks of memory cells with non-overlapping portions thereof storing at least user data and overhead information"
    a. must include multiple erase blocks; and
    b. is limited to a single user data and a single overhead portion.
2. As used in claim 1 of U.S. Pat. No. 6,757,842, "linking the address of such unusable blocks with addresses of other blocks that are useable"
    a. need not take the form of a map or table listing both the unusable and corresponding usable blocks; and
    b. is not satisfied by simply storing a pointer in the unusable block that points to a usable block.

3. As used in claims 67 and 79 of U.S. Pat. No. 6,149,316 and claims 1, 10 and 61 of U.S. Patent No. 6,757,842, "an address in a format designating at least one mass memory storage block" and "a mass memory storage block address"

   a. need not refer to a block of user data; and

   b. may specify either a physical or a logical address for the mass memory system.

4. As used in claim 16 of U.S. Pat. No. 5,719,808, "an array of EEPROM cells"

   a. is not required to be "contiguous"; and

   b. is not required to contain "dedicated row and column decoders."

5. As used in claim 16 of U.S. Pat. No. 5,719,808, "designating a combinations [sic] of a plurality of but less than all of said multiple sectors to be erased"

   a. requires the ability to designate "any combination" of sectors for erase; and

   b. does not require setting a tag in a dedicated register for each sector.

6. As used in claims 13 and 58 of U.S. Pat. No. 6,426,893, "individual ones of the redundancy codes being appended to ends of the user data from which they are generated" requires the redundancy code to be attached immediately adjacent to the user data.

7. As used in claims 1 and 3 of U.S. Pat. No. 6,426,893, "information of the characteristics of said first group of blocks" and "characteristics of the memory cell blocks" may include logical and physical block addresses.

8. As used in claims 1 and 3 of U.S. Pat. No. 6,763,424, "recording a relative time of programming the at least one page of new data and the at least one page of superseded data"

   a. is not satisfied by performing the block recording method described at column 9, lines 40-53 of the specification;

   b. does not require recording the relative time of programming on the page being recorded; and

   c. does not require allowing a comparison of the page to pages written after the page.

9. As used in claims 20 and 24 of U.S. Pat. No. 6,763,424,

   a. "logical page address" is not limited to a "logical block number plus logical page offset"; and

   b. "programming individual ones of a second plurality of a total number of pages less than said given number in a second block with updated data and a logical page address associated with the updated data" must

occur "without marking the original (now superseded) pages of data in the first block with an invalid data flag."

Entered this 22d day of September, 2010.

BY THE COURT:
/s/
BARBARA B. CRABB
District Judge