

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 March 2008 (20.03.2008)

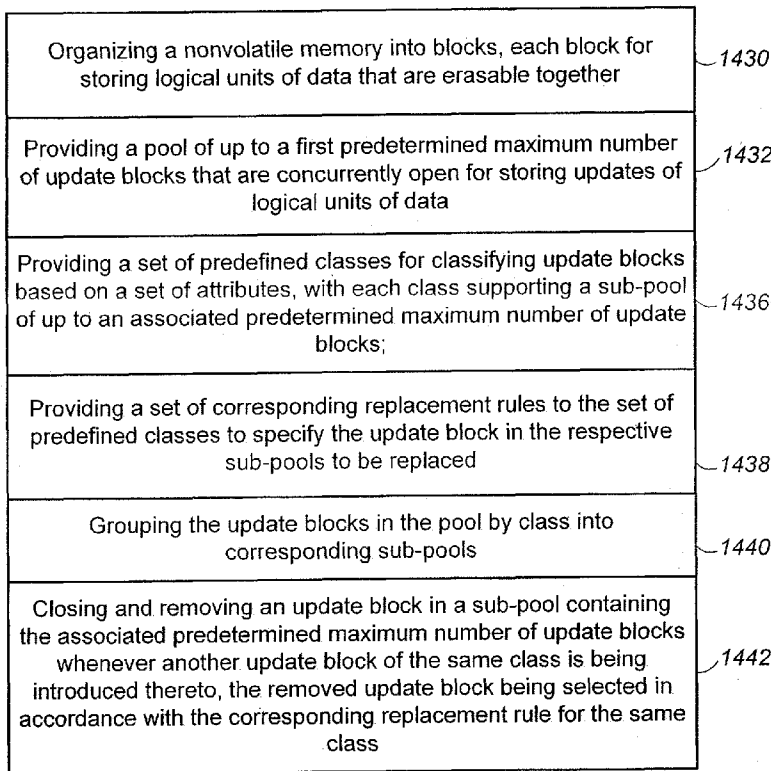
PCT

(10) International Publication Number
WO 2008/033952 A2

- (51) International Patent Classification: Not classified
- (21) International Application Number: PCT/US2007/078311
- (22) International Filing Date: 12 September 2007 (12.09.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
 - 11/532,456 15 September 2006 (15.09.2006) US
 - 11/532,467 15 September 2006 (15.09.2006) US
- (71) Applicant (for all designated States except US): SAN-DISK CORPORATION [US/US]; 601 McCarthy Blvd., Milpitas, California 95035 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): LIN, Jason T. [US/US]; 2670 Iverson Ct., Santa Clara, California 95051 (US).
- (74) Agents: PARSONS, Gerald, P. et al.; Davis Wright Tremaine LLP, 505 Montgomery Street, Suite 800, San Francisco, California 94111 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: NON-VOLATILE MEMORY AND METHOD FOR CLASS-BASED UPDATE BLOCK REPLACEMENT RULES



(57) Abstract: In a nonvolatile memory with block management system, data are written to blocks and are erasable block by block. At any time a pool of blocks are open for storing data concurrently. The number of blocks in the pool is limited. A replacement system allows new blocks to be introduced into the pool without exceeding the limit. In particular, different classes of blocks in the pool each has its own replacement rule, such as closing a least active block before being replaced. In this way, possible inefficiency and premature closure of blocks in the pool can be avoided.

Class-based Replacement Rules for a Set of Update Blocks

WO 2008/033952 A2



Published:

- *without international search report and to be republished upon receipt of that report*

NON-VOLATILE MEMORY AND METHOD FOR CLASS-BASED UPDATE BLOCK REPLACEMENT RULES

FIELD OF THE INVENTION

[0001] This invention relates generally to non-volatile semiconductor memory and specifically to those having a memory block management system with an improved system for managing replacement of a pool of blocks open concurrently for storing data.

BACKGROUND OF THE INVENTION

[0002] Solid-state memory capable of nonvolatile storage of charge, particularly in the form of EEPROM and flash EEPROM packaged as a small form factor card, has recently become the storage of choice in a variety of mobile and handheld devices, notably information appliances and consumer electronics products. Unlike RAM (random access memory) that is also solid-state memory, flash memory is non-volatile, and retaining its stored data even after power is turned off. Also, unlike ROM (read only memory), flash memory is rewritable similar to a disk storage device. In spite of the higher cost, flash memory is increasingly being used in mass storage applications. Conventional mass storage, based on rotating magnetic medium such as hard drives and floppy disks, is unsuitable for the mobile and handheld environment. This is because disk drives tend to be bulky, are prone to mechanical failure and have high latency and high power requirements. These undesirable attributes make disk-based storage impractical in most mobile and portable applications. On the other hand, flash memory, both embedded and in the form of a removable card is ideally suited in the mobile and handheld environment because of its small size, low power consumption, high speed and high reliability features.

[0003] Flash EEPROM is similar to EEPROM (electrically erasable and programmable read-only memory) in that it is a non-volatile memory that can be erased and have new data written or “programmed” into their memory cells. Both utilize a floating (unconnected) conductive gate, in a field effect transistor structure,

positioned over a channel region in a semiconductor substrate, between source and drain regions. A control gate is then provided over the floating gate. The threshold voltage characteristic of the transistor is controlled by the amount of charge that is retained on the floating gate. That is, for a given level of charge on the floating gate, there is a corresponding voltage (threshold) that must be applied to the control gate before the transistor is turned "on" to permit conduction between its source and drain regions. In particular, flash memory such as Flash EEPROM allows entire blocks of memory cells to be erased at the same time.

[0004] The floating gate can hold a range of charges and therefore can be programmed to any threshold voltage level within a threshold voltage window. The size of the threshold voltage window is delimited by the minimum and maximum threshold levels of the device, which in turn correspond to the range of the charges that can be programmed onto the floating gate. The threshold window generally depends on the memory device's characteristics, operating conditions and history. Each distinct, resolvable threshold voltage level range within the window may, in principle, be used to designate a definite memory state of the cell. When the threshold voltage is partitioned into two distinct regions, each memory cell will be able to store one bit of data. Similarly, when the threshold voltage window is partitioned into more than two distinct regions, each memory cell will be able to store more than one bit of data.

[0005] The transistor serving as a memory cell is typically programmed to a "programmed" state by one of two mechanisms. In "hot electron injection," a high voltage applied to the drain accelerates electrons across the substrate channel region. At the same time a high voltage applied to the control gate pulls the hot electrons through a thin gate dielectric onto the floating gate. In "tunneling injection," a high voltage is applied to the control gate relative to the substrate. In this way, electrons are pulled from the substrate to the intervening floating gate. While the term "program" has been used historically to describe writing to a memory by injecting electrons to an initially erased charge storage unit of the memory cell so as to alter the memory state, it has now been used interchangeable with more common terms such as "write" or "record."

[0006] The memory device may be erased by a number of mechanisms. For EEPROM, a memory cell is electrically erasable, by applying a high voltage to the substrate relative to the control gate so as to induce electrons in the floating gate to tunnel through a thin oxide to the substrate channel region (i.e., Fowler-Nordheim tunneling.) Typically, the EEPROM is erasable byte by byte. For flash EEPROM, the memory is electrically erasable either all at once or one or more minimum erasable blocks at a time, where a minimum erasable block may consist of one or more sectors and each sector may store 512 bytes or more of data.

[0007] The memory device typically comprises one or more memory chips that may be mounted on a card. Each memory chip comprises an array of memory cells supported by peripheral circuits such as decoders and erase, write and read circuits. The more sophisticated memory devices also come with a controller that performs intelligent and higher level memory operations and interfacing.

[0008] There are many commercially successful non-volatile solid-state memory devices being used today. These memory devices may be flash EEPROM or may employ other types of nonvolatile memory cells. Examples of flash memory and systems and methods of manufacturing them are given in United States patents nos. 5,070,032, 5,095,344, 5,315,541, 5,343,063, and 5,661,053, 5,313,421 and 6,222,762. In particular, flash memory devices with NAND string structures are described in United States patent nos. 5,570,315, 5,903,495, 6,046,935. Also nonvolatile memory devices are also manufactured from memory cells with a dielectric layer for storing charge. Instead of the conductive floating gate elements described earlier, a dielectric layer is used. Such memory devices utilizing dielectric storage element have been described by Eitan et al., "NROM: A Novel Localized Trapping, 2-Bit Nonvolatile Memory Cell," IEEE Electron Device Letters, vol. 21, no. 11, November 2000, pp. 543-545. An ONO dielectric layer extends across the channel between source and drain diffusions. The charge for one data bit is localized in the dielectric layer adjacent to the drain, and the charge for the other data bit is localized in the dielectric layer adjacent to the source. For example, United States patents nos. 5,768,192 and 6,011,725 disclose a nonvolatile memory cell having a trapping dielectric sandwiched between two silicon dioxide layers. Multi-state data storage is implemented by separately reading the binary states of the spatially separated charge storage regions

within the dielectric.

[0009] In order to improve read and program performance, multiple charge storage elements or memory transistors in an array are read or programmed in parallel. Thus, a “page” of memory elements are read or programmed together. In existing memory architectures, a row typically contains several interleaved pages or it may constitute one page. All memory elements of a page will be read or programmed together.

[0010] In flash memory systems, erase operation may take as much as an order of magnitude longer than read and program operations. Thus, it is desirable to have the erase block of substantial size. In this way, the erase time is amortized over a large aggregate of memory cells.

[0011] The nature of flash memory predicates that data must be written to an erased memory location. If data of a certain logical address from a host is to be updated, one way is rewrite the update data in the same physical memory location. That is, the logical to physical address mapping is unchanged. However, this will mean the entire erase block contain that physical location will have to be first erased and then rewritten with the updated data. This method of update is inefficient, as it requires an entire erase block to be erased and rewritten, especially if the data to be updated only occupies a small portion of the erase block. It will also result in a higher frequency of erase recycling of the memory block, which is undesirable in view of the limited endurance of this type of memory device.

[0012] Another problem with managing flash memory system has to do with system control and directory data. The data is produced and accessed during the course of various memory operations. Thus, its efficient handling and ready access will directly impact performance. It would be desirable to maintain this type of data in flash memory because flash memory is meant for storage and is nonvolatile. However, with an intervening file management system between the controller and the flash memory, the data can not be accessed as directly. Also, system control and directory data tends to be active and fragmented, which is not conducive to storing in a system with large size block erase. Conventionally, this type of data is set up in the controller RAM, thereby allowing direct access by the controller. After the memory

device is powered up, a process of initialization enables the flash memory to be scanned in order to compile the necessary system control and directory information to be placed in the controller RAM. This process takes time and requires controller RAM capacity, all the more so with ever increasing flash memory capacity.

[0013] US 6,567,307 discloses a method of dealing with sector updates among large erase block including recording the update data in multiple erase blocks acting as scratch pad and eventually consolidating the valid sectors among the various blocks and rewriting the sectors after rearranging them in logically sequential order. In this way, a block needs not be erased and rewritten at every slightest update.

[0014] W0 03/027828 and W0 00/49488 both disclose a memory system dealing with updates among large erase block including partitioning the logical sector addresses in zones. A small zone of logical address range is reserved for active system control data separate from another zone for user data. In this way, manipulation of the system control data in its own zone will not interact with the associated user data in another zone. Updates are at the logical sector level and a write pointer points to the corresponding physical sectors in a block to be written. The mapping information is buffered in RAM and eventually stored in a sector allocation table in the main memory. The latest version of a logical sector will obsolete all previous versions among existing blocks, which become partially obsolete. Garbage collection is performed to keep partially obsolete blocks to an acceptable number.

[0015] Prior art systems tend to have the update data distributed over many blocks or the update data may render many existing blocks partially obsolete. The result often is a large amount of garbage collection necessary for the partially obsolete blocks, which is inefficient and causes premature aging of the memory. Also, there is no systematic and efficient way of dealing with sequential update as compared to non-sequential update.

[0016] Therefore there is a general need for high capacity and high performance non-volatile memory. In particular, there is a need to have a high capacity nonvolatile memory able to conduct memory operations in large blocks without the aforementioned problems.

SUMMARY OF INVENTION

[0017] A non-volatile memory system is organized in physical groups of physical memory locations. Each physical group (metablock) is erasable as a unit and can be used to store a logical group of data. A memory management system allows for update of a logical group of data by allocating a metablock dedicated to recording the update data of the logical group. The update metablock records update data in the order received and has no restriction on whether the recording is in the correct logical order as originally stored (sequential) or not (chaotic). Eventually the update metablock is closed to further recording. One of several processes will take place, but will ultimately end up with a fully filled metablock in the correct order which replaces the original metablock. In the chaotic case, directory data is maintained in the non-volatile memory in a manner that is conducive to frequent updates. The system supports multiple logical groups being updated concurrently.

[0018] One feature of the invention allows data to be updated logical-group by logical-group. Thus, when a logical group is being updated, the distribution of logical units (and also the scatter of memory units that the updates obsolete) are limited in range. This is especially true when the logical group is normally contained within a physical block.

[0019] During updates of the logical group, typically one or two blocks need be assigned to buffer the updated logical units. Thus, garbage collection need only be performed over a relatively fewer number of blocks. Garbage collection of a chaotic block may be performed by either consolidation or compaction.

[0020] The economy of the update process is further evident in the generic treatment of the update blocks so that no additional block need be assigned for chaotic (non-sequential) updates as compared to the sequential ones. All update blocks are allocated as sequential update block, and any update block can change to a chaotic update block. Indeed, the change of an update block from sequential to chaotic is discretionary.

[0021] The efficient use of system resource allows multiple logical groups to be updated concurrently. This further increases efficiency and reduces overheads.

[0022] According to one aspect of the invention, in a nonvolatile memory with a block management system, an improved block replacement scheme is implemented for a system supporting up to a first predetermined maximum number of update blocks that are concurrently opened for recording data. The update blocks are mainly sequential update blocks where data are recorded in logically sequential order but up to a second predetermined maximum number of which are allowed to be chaotic update blocks where data are not recorded in logically sequential order. Whenever a new allocation of an update block may cause the pool of update blocks to exceed either the first or second predetermined maximum number, one of the existing update blocks in the pool will be closed and removed in order to comply with the limitation. Prior to closing the update block, its data are consolidated into a sequential block. The improved scheme is to avoid the situation where a sequential update can cause an excessive number of chaotic block consolidations. This is accomplished by separating sequential update blocks and chaotic update blocks into respective replacement or consolidation pools. In particular, when a sequential update causes the allocation of a new update block to exceed the first predetermined maximum number, a least recently used sequential update block of the pool is preferentially to make room.

[0023] In the current system, generally there are two types of data: user data and control data. The user data are sent from a host to the memory system typically in logical sequential order. Sequential update blocks are allocated to optimally handle sequential write operations from the host. The user data can also be in logically non-sequential order especially when there are subsequent updates to the logical data. Chaotic update blocks are created to optimally handle the data in non-sequential order. Another source of chaotic or non-sequential data is control data maintained by the file system or the memory system such as file and directory information which are generated in the course of storing user data.

[0024] Previous scheme of complying to a practical system limitation of supporting up to a maximum number of concurrently opened update blocks has been to close the least recently used update block in the pool, irrespective of whether it is sequential or chaotic.

[0025] The present scheme improves over the previous scheme where essentially, if during a sequential write operation an update block among a pool thereof needs to be closed to make room for a new allocation, the least recently used sequential update block in the pool is closed. This ensures the various update blocks are effectively used to handle sequential write operations and random write operations. In particular, it avoids the inefficient situation where a large sequential write operation by the host may force a premature closure of a chaotic update block containing FAT and Directory information. Another chaotic block will in effect be created very soon to store the FAT and Directory information which will be updated again once the large sequential write operation is done. The creation of the improved replacement policy mandates separation of the replacement and consolidation pool to prevent the added overhead in consolidating the chaotic block during sequential write and consolidation of potentially an open sequential or an open chaotic block to manage the subsequent FAT and Directory update.

[0026] A generalization of the present scheme is to classify the update blocks based on a set of attributes, such as whether the update block is storing sequential or non-sequential data or whether it is storing some predefined type of system data. In implementing a pool of update blocks of limited number, each class of update blocks will have its own rule for replacement when the maximum number supported for that class will be exceeded.

[0027] For example, sequential update block and non-sequential update blocks are two different classes. The replacement rules for each of these classes are the same, namely to replace a least active one with a new one. Thus, when the pool of sequential update blocks will be exceeded, a least active one in the pool will be closed and removed before a new one is introduced to the pool. Similarly for the pool of non-sequential update blocks.

[0028] In general each class has its own replacement rule independent of the other classes. Examples of replacement rules are to replace the least recently accessed, the most recently accessed, the least frequently accessed, the most frequently accessed, etc, depending on the corresponding classes.

[0029] Additional features and advantages of the present invention will be understood

from the following description of its preferred embodiments, which description should be taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0030] FIG. 1 illustrates schematically the main hardware components of a memory system suitable for implementing the present invention.

[0031] FIG. 2 illustrates the memory being organized into physical groups of sectors (or metablocks) and managed by a memory manager of the controller, according to a preferred embodiment of the invention.

[0032] FIGs. 3A(i) – 3A(iii) illustrate schematically the mapping between a logical group and a metablock, according to a preferred embodiment of the present invention.

[0033] FIG. 3B illustrates schematically the mapping between logical groups and metablocks.

[0034] FIG. 4 illustrates the alignment of a metablock with structures in physical memory.

[0035] FIG. 5A illustrates metablocks being constituted from linking of minimum erase units of different planes.

[0036] FIG. 5B illustrates one embodiment in which one minimum erase unit (MEU) is selected from each plane for linking into a metablock.

[0037] FIG. 5C illustrates another embodiment in which more than one MEU are selected from each plane for linking into a metablock.

[0038] FIG. 6 is a schematic block diagram of the metablock management system as implemented in the controller and flash memory.

[0039] FIG. 7A illustrates an example of sectors in a logical group being written in sequential order to a sequential update block.

[0040] FIG. 7B illustrates an example of sectors in a logical group being

written in chaotic order to a chaotic update block.

[0041] **FIG. 8** illustrates an example of sectors in a logical group being written in sequential order to a sequential update block as a result of two separate host write operations that has a discontinuity in logical addresses.

[0042] **FIG. 9** is a flow diagram illustrating a process by the update block manager to update a logical group of data, according a general embodiment of the invention.

[0043] **FIG. 10** is a flow diagram illustrating a process by the update block manager to update a logical group of data, according a preferred embodiment of the invention.

[0044] **FIG. 11A** is a flow diagram illustrating in more detail the consolidation process of closing a chaotic update block shown in **FIG. 10**.

[0045] **FIG. 11B** is a flow diagram illustrating in more detail the compaction process for closing a chaotic update block shown in **FIG. 10**.

[0046] **FIG. 12A** illustrates all possible states of a Logical Group, and the possible transitions between them under various operations.

[0047] **FIG. 12B** is a table listing the possible states of a Logical Group.

[0048] **FIG. 13A** illustrates all possible states of a metablock, and the possible transitions between them under various operations. A metablock is a Physical Group corresponding to a Logical Group.

[0049] **FIG. 13B** is a table listing the possible states of a metablock.

[0050] **FIGs. 14(A) – 14(J)** are state diagrams showing the effect of various operations on the state of the logical group and also on the physical metablock.

[0051] **FIG. 15** illustrates a preferred embodiment of the structure of an allocation block list (ABL) for keeping track of opened and closed update blocks and erased blocks for allocation.

[0052] FIG. 16A illustrates the data fields of a chaotic block index (CBI) sector.

[0053] FIG. 16B illustrates an example of the chaotic block index (CBI) sectors being recorded in a dedicated metablock.

[0054] FIG. 16C is a flow diagram illustrating access to the data of a logical sector of a given logical group undergoing chaotic update.

[0055] FIG. 16D is a flow diagram illustrating access to the data of a logical sector of a given logical group undergoing chaotic update, according to an alternative embodiment in which logical group has been partitioned into subgroups.

[0056] FIG. 16E illustrates examples of Chaotic Block Indexing (CBI) sectors and their functions for the embodiment where each logical group is partitioned into multiple subgroups.

[0057] FIG. 17A illustrates the data fields of a group address table (GAT) sector.

[0058] FIG. 17B illustrates an example of the group address table (GAT) sectors being recorded in a GAT block.

[0059] FIG. 18 is a schematic block diagram illustrating the distribution and flow of the control and directory information for usage and recycling of erased blocks.

[0060] FIG. 19 is a flow chart showing the process of logical to physical address translation.

[0061] FIG. 20 illustrates the hierarchy of the operations performed on control data structures in the course of the operation of the memory management.

[0062] FIG. 21 illustrates schematically the two prescribed limits on the number of update blocks for a block managing system.

[0063] FIG. 22 illustrates typical examples of combinations of the two limits optimized for various memory devices.

[0064] FIG. 23A illustrates schematically an update pool with a “5 – 2” configuration as described in FIG. 22.

[0065] FIG. 23B illustrates schematically the closing of the least active update block in order to make room for a new update block, according to the previous scheme.

[0066] FIG. 23C illustrates schematically introducing a newly allocated update block into the pool after a closed update block has been removed to make room.

[0067] FIG. 24A illustrates schematically an update pool with a “5 – 2” configuration as described in FIG. 22.

[0068] FIG. 24B illustrates schematically the closing of the least active update block in order to make room for a new update block, according to the previous scheme.

[0069] FIG. 24C illustrates schematically introducing a newly allocated update block into the pool after a closed update block has been removed to make room.

[0070] FIG. 25A illustrates the scheme previously illustrated in FIG. 10, STEP 410 and also in FIG. 23B and FIG. 24B where the least recently accessed update block is closed whenever a new allocation would exceed a predetermined limit.

[0071] FIG. 25B illustrates the scheme previously illustrated in FIG. 10, STEP 370 where the least recently accessed chaotic (non-sequential) update block is closed whenever the number of chaotic update blocks exceeds a predetermined limit.

[0072] FIG. 26A illustrates schematically an update pool with a “5 – 2” configuration as described in FIG. 22.

[0073] FIG. 26B illustrates schematically the closing of one among the pool of an update blocks in order to make room for a new update block, according to the present improved scheme.

[0074] FIG. 26C illustrates schematically introducing a newly allocated update block into the pool after a closed update block has been removed to make room.

[0075] FIG. 27A illustrates schematically an update pool with a “5 – 2” configuration as described in FIG. 22.

[0076] FIG. 27B illustrates schematically the closing of one among the pool of an update blocks in order to make room for a new update block, according to the present improved scheme.

[0077] FIG. 27C illustrates schematically introducing a new chaotic update block into the pool after another chaotic update block has been closed and removed to make room.

[0078] FIG. 28 is a flow chart illustrating the present improved scheme of managing a limited set of update blocks during a sequential update, according to a first embodiment.

[0079] FIG. 29 is a flow chart illustrating the present improved scheme of managing a limited set of update blocks having two predetermined limits, according to a second embodiment.

[0080] FIG. 30 is a flow chart illustrating the present improved scheme of managing a limited set of update blocks having class-based replacement rules.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0081] FIG.1 to FIG. 20 illustrate examples of memory systems with block management in which the various aspects of the present invention may be implemented. Similar memory systems have been disclosed in U.S. Patent Application Publication No. US-2005-0144365-A1, entitled “Non-Volatile Memory and Method with Control Data Management,” by Gorobets et al.

[0082] FIG. 1 illustrates schematically the main hardware components of a memory system suitable for implementing the present invention. The memory system 20 typically operates with a host 10 through a host interface. The memory system is

typically in the form of a memory card or an embedded memory system. The memory system 20 includes a memory 200 whose operations are controlled by a controller 100. The memory 200 comprises of one or more array of non-volatile memory cells distributed over one or more integrated circuit chip. The controller 100 includes an interface 110, a processor 120, an optional coprocessor 121, ROM 122 (read-only-memory), RAM 130 (random access memory) and optionally programmable nonvolatile memory 124. The interface 110 has one component interfacing the controller to a host and another component interfacing to the memory 200. Firmware stored in nonvolatile ROM 122 and/or the optional nonvolatile memory 124 provides codes for the processor 120 to implement the functions of the controller 100. Error correction codes may be processed by the processor 120 or the optional coprocessor 121. In an alternative embodiment, the controller 100 is implemented by a state machine (not shown.) In yet another embodiment, the controller 100 is implemented within the host.

LOGICAL AND PHYSICAL BLOCK STRUCTURES

[0083] FIG. 2 illustrates the memory being organized into physical groups of sectors (or metablocks) and managed by a memory manager of the controller, according to a preferred embodiment of the invention. The memory 200 is organized into metablocks, where each metablock is a group of physical sectors S_0, \dots, S_{N-1} that are erasable together.

[0084] The host 10 accesses the memory 200 when running an application under a file system or operating system. Typically, the host system addresses data in units of logical sectors where, for example, each sector may contain 512 bytes of data. Also, it is usual for the host to read or write to the memory system in unit of logical clusters, each consisting of one or more logical sectors. In some host systems, an optional host-side memory manager may exist to perform lower level memory management at the host. In most cases during read or write operations, the host 10 essentially issues a command to the memory system 20 to read or write a segment containing a string of logical sectors of data with contiguous addresses.

[0085] A memory-side memory manager is implemented in the controller 100 of the memory system 20 to manage the storage and retrieval of the data of host

logical sectors among metablocks of the flash memory 200. In the preferred embodiment, the memory manager contains a number of software modules for managing erase, read and write operations of the metablocks. The memory manager also maintains system control and directory data associated with its operations among the flash memory 200 and the controller RAM 130.

[0086] **FIGs. 3A(i) – 3A(iii)** illustrate schematically the mapping between a logical group and a metablock, according to a preferred embodiment of the present invention. The metablock of the physical memory has N physical sectors for storing N logical sectors of data of a logical group. **FIG. 3A(i)** shows the data from a logical group LG_i , where the logical sectors are in contiguous logical order $0, 1, \dots, N-1$. **FIG. 3A(ii)** shows the same data being stored in the metablock in the same logical order. The metablock when stored in this manner is said to be “sequential.” In general, the metablock may have data stored in a different order, in which case the metablock is said to be “non-sequential” or “chaotic.”

[0087] There may be an offset between the lowest address of a logical group and the lowest address of the metablock to which it is mapped. In this case, logical sector address wraps round as a loop from bottom back to top of the logical group within the metablock. For example, in **FIG. 3A(iii)**, the metablock stores in its first location beginning with the data of logical sector k . When the last logical sector $N-1$ is reached, it wraps around to sector 0 and finally storing data associated with logical sector $k-1$ in its last physical sector. In the preferred embodiment, a page tag is used to identify any offset, such as identifying the starting logical sector address of the data stored in the first physical sector of the metablock. Two blocks will be considered to have their logical sectors stored in similar order when they only differ by a page tag.

[0088] **FIG. 3B** illustrates schematically the mapping between logical groups and metablocks. Each logical group is mapped to a unique metablock, except for a small number of logical groups in which data is currently being updated. After a logical group has been updated, it may be mapped to a different metablock. The mapping information is maintained in a set of logical to physical directories, which will be described in more detail later.

[0089] Other types of logical group to metablock mapping are also

contemplated. For example, metablocks with variable size are disclosed in co-pending and co-owned United States Patent application, entitled, "Adaptive Metablocks," filed by Alan Sinclair, on the same day as the present application. The entire disclosure of the co-pending application is hereby incorporated herein by reference.

[0090] One feature of the invention is that the system operates with a single logical partition, and groups of logical sectors throughout the logical address range of the memory system are treated identically. For example, sectors containing system data and sectors containing user data can be distributed anywhere among the logical address space.

[0091] Unlike prior art systems, there is no special partitioning or zoning of system sectors (i.e., sectors relating to file allocation tables, directories or sub-directories) in order to localize in logical address space sectors that are likely to contain data with high-frequency and small-size updates. Instead, the present scheme of updating logical groups of sectors will efficiently handle the patterns of access that are typical of system sectors, as well as those typical of file data.

[0092] FIG. 4 illustrates the alignment of a metablock with structures in physical memory. Flash memory comprises blocks of memory cells which are erasable together as a unit. Such erase blocks are the minimum unit of erasure of flash memory or minimum erasable unit (MEU) of the memory. The minimum erase unit is a hardware design parameter of the memory, although in some memory systems that supports multiple MEUs erase, it is possible to configure a "super MEU" comprising more than one MEU. For flash EEPROM, a MEU may comprise one sector but preferably multiple sectors. In the example shown, it has M sectors. In the preferred embodiment, each sector can store 512 bytes of data and has a user data portion and a header portion for storing system or overhead data. If the metablock is constituted from P MEUs, and each MEU contains M sectors, then, each metablock will have $N = P * M$ sectors.

[0093] The metablock represents, at the system level, a group of memory locations, e.g., sectors that are erasable together. The physical address space of the flash memory is treated as a set of metablocks, with a metablock being the minimum

unit of erasure. Within this specification, the terms “metablock” and “block” are used synonymously to define the minimum unit of erasure at the system level for media management, and the term “minimum erase unit” or MEU is used to denote the minimum unit of erasure of flash memory.

LINKING OF MINIMUM ERASE UNITS (MEUs) TO FORM A METABLOCK

[0094] In order to maximize programming speed and erase speed, parallelism is exploited as much as possible by arranging for multiple pages of information, located in multiple MEUs, to be programmed in parallel, and for multiple MEUs to be erased in parallel.

[0095] In flash memory, a page is a grouping of memory cells that may be programmed together in a single operation. A page may comprise one or more sector. Also, a memory array may be partitioned into more than one plane, where only one MEU within a plane may be programmed or erased at a time. Finally, the planes may be distributed among one or more memory chips.

[0096] In flash memory, the MEUs may comprise one or more page. MEUs within a flash memory chip may be organized in planes. Since one MEU from each plane may be programmed or erased concurrently, it is expedient to form a multiple MEU metablock by selecting one MEU from each plane (see **FIG. 5B** below.)

[0097] **FIG. 5A** illustrates metablocks being constituted from linking of minimum erase units of different planes. Each metablock, such as MB0, MB1, ..., is constituted from MEUs from different planes of the memory system, where the different planes may be distributed among one or more chips. The metablock link manager 170 shown in **FIG. 2** manages the linking of the MEUs for each metablock. Each metablock is configured during an initial formatting process, and retains its constituent MEUs throughout the life of the system, unless there is a failure of one of the MEUs.

[0098] **FIG. 5B** illustrates one embodiment in which one minimum erase unit (MEU) is selected from each plane for linking into a metablock.

[0099] **FIG. 5C** illustrates another embodiment in which more than one MEU are

selected from each plane for linking into a metablock. In another embodiment, more than one MEU may be selected from each plane to form a super MEU. For example, a super MEU may be formed from two MEUs. In this case, it may take more than one pass for read or write operation.

[00100] The linking and re-linking of MEUs into metablocks is also disclosed in co-pending and co-owned United States Patent application, entitled "Adaptive Deterministic Grouping of Blocks into Multi-Block Structures," filed by Carlos Gonzales et al, on the same day as the present application. The entire disclosure of the co-pending application is hereby incorporated herein by reference.

METABLOCK MANAGEMENT

[00101] FIG. 6 is a schematic block diagram of the metablock management system as implemented in the controller and flash memory. The metablock management system comprises various functional modules implemented in the controller 100 and maintains various control data (including directory data) in tables and lists hierarchically distributed in the flash memory 200 and the controller RAM 130. The function modules implemented in the controller 100 includes an interface module 110, a logical-to-physical address translation module 140, an update block manager module 150, an erase block manager module 160 and a metablock link manager 170.

[00102] The interface 110 allows the metablock management system to interface with a host system. The logical to physical address translation module 140 maps the logical address from the host to a physical memory location. The update block Manager module 150 manages data update operations in memory for a given logical group of data. The erased block manager 160 manages the erase operation of the metablocks and their allocation for storage of new information. A metablock link manager 170 manages the linking of subgroups of minimum erasable blocks of sectors to constitute a given metablock. Detailed description of these modules will be given in their respective sections.

[00103] During operation the metablock management system generates and works with control data such as addresses, control and status information. Since

much of the control data tends to be frequently changing data of small size, it can not be readily stored and maintained efficiently in a flash memory with a large block structure. A hierarchical and distributed scheme is employed to store the more static control data in the nonvolatile flash memory while locating the smaller amount of the more varying control data in controller RAM for more efficient update and access. In the event of a power shutdown or failure, the scheme allows the control data in the volatile controller RAM to be rebuilt quickly by scanning a small set of control data in the nonvolatile memory. This is possible because the invention restricts the number of blocks associated with the possible activity of a given logical group of data. In this way, the scanning is confined. In addition, some of the control data that requires persistence are stored in a nonvolatile metablock that can be updated sector-by-sector, with each update resulting in a new sector being recorded that supercedes a previous one. A sector indexing scheme is employed for control data to keep track of the sector-by-sector updates in a metablock.

[00104] The non-volatile flash memory 200 stores the bulk of control data that are relatively static. This includes group address tables (GAT) 210, chaotic block indices (CBI) 220, erased block lists (EBL) 230 and MAP 240. The GAT 210 keeps track of the mapping between logical groups of sectors and their corresponding metablocks. The mappings do not change except for those undergoing updates. The CBI 220 keeps track of the mapping of logically non-sequential sectors during an update. The EBL 230 keeps track of the pool of metablocks that have been erased. MAP 240 is a bitmap showing the erase status of all metablocks in the flash memory.

[00105] The volatile controller RAM 130 stores a small portion of control data that are frequently changing and accessed. This includes an allocation block list (ABL) 134 and a cleared block list (CBL) 136. The ABL 134 keeps track of the allocation of metablocks for recording update data while the CBL 136 keeps track of metablocks that have been deallocated and erased. In the preferred embodiment, the RAM 130 acts as a cache for control data stored in flash memory 200.

UPDATE BLOCK MANAGER

[00106] The update block manager 150 (shown in **FIG. 2**) handles the update of logical groups. According to one aspect of the invention, each logical group of

sectors undergoing an update is allocated a dedicated update metablock for recording the update data. In the preferred embodiment, any segment of one or more sectors of the logical group will be recorded in the update block. An update block can be managed to receive updated data in either sequential order or non-sequential (also known as chaotic) order. A chaotic update block allows sector data to be updated in any order within a logical group, and with any repetition of individual sectors. In particular, a sequential update block can become a chaotic update block, without need for relocation of any data sectors. No predetermined allocation of blocks for chaotic data update is required; a non-sequential write at any logical address is automatically accommodated. Thus, unlike prior art systems, there is no special treatment whether the various update segments of the logical group is in logical sequential or non-sequential order. The generic update block will simply be used to record the various segments in the order they are requested by the host. For example, even if host system data or system control data tends to be updated in chaotic fashion, regions of logical address space corresponding to host system data do not need to be treated differently from regions with host user data.

[00107] Data of a complete logical group of sectors is preferably stored in logically sequential order in a single metablock. In this way, the index to the stored logical sectors is predefined. When the metablock has in store all the sectors of a given logical group in a predefined order it is said to be "intact." As for an update block, when it eventually fills up with update data in logically sequential order, then the update block will become an updated intact metablock that readily replace the original metablock. On the other hand, if the update block fills up with update data in a logically different order from that of the intact block, the update block is a non-sequential or chaotic update block and the out of order segments must be further processed so that eventually the update data of the logical group is stored in the same order as that of the intact block. In the preferred case, it is in logically sequential order in a single metablock. The further processing involves consolidating the updated sectors in the update block with unchanged sectors in the original block into yet another update metablock. The consolidated update block will then be in logically sequential order and can be used to replace the original block. Under some predetermined condition, the consolidation process is preceded by one or more

compaction processes. The compaction process simply re-records the sectors of the chaotic update block into a replacing chaotic update block while eliminating any duplicate logical sector that has been rendered obsolete by a subsequent update of the same logical sector.

[00108] The update scheme allows for multiple update threads running concurrently, up to a predefined maximum. Each thread is a logical group undergoing updates using its dedicated update metablock.

Sequential data update

[00109] When data belonging to a logical group is first updated, a metablock is allocated and dedicated as an update block for the update data of the logical group. The update block is allocated when a command is received from the host to write a segment of one or more sectors of the logical group for which an existing metablock has been storing all its sectors intact. For the first host write operation, a first segment of data is recorded on the update block. Since each host write is a segment of one or more sector with contiguous logical address, it follows that the first update is always sequential in nature. In subsequent host writes, update segments within the same logical group are recorded in the update block in the order received from the host. A block continues to be managed as a sequential update block whilst sectors updated by the host within the associated logical group remain logically sequential. All sectors updated in this logical group are written to this sequential update block, until the block is either closed or converted to a chaotic update block.

[00110] FIG. 7A illustrates an example of sectors in a logical group being written in sequential order to a sequential update block as a result of two separate host write operations, whilst the corresponding sectors in the original block for the logical group become obsolete. In host write operation #1, the data in the logical sectors LS5-LS8 are being updated. The updated data as LS5'-LS8' are recorded in a newly allocated dedicated update block.

[00111] For expediency, the first sector to be updated in the logical group is recorded in the dedicated update block starting from the first physical sector location. In general, the first logical sector to be updated is not necessarily the logical first

sector of the group, and there may therefore be an offset between the start of the logical group and the start of the update block. This offset is known as page tag as described previously in connection with **FIG. 3A**. Subsequent sectors are updated in logically sequential order. When the last sector of the logical group is written, group addresses wrap around and the write sequence continues with the first sector of the group.

[00112] In host write operation #2, the segment of data in the logical sectors LS9-LS12 are being updated. The updated data as LS9'-LS12' are recorded in the dedicated update block in a location directly following where the last write ends. It can be seen that the two host writes are such that the update data has been recorded in the update block in logically sequential order, namely LS5'-LS12'. The update block is regarded as a sequential update block since it has been filled in logically sequential order. The update data recorded in the update block obsoletes the corresponding ones in the original block.

Chaotic data update

[00113] Chaotic update block management may be initiated for an existing sequential update block when any sector updated by the host within the associated logical group is logically non-sequential. A chaotic update block is a form of data update block in which logical sectors within an associated logical group may be updated in any order and with any amount of repetition. It is created by conversion from a sequential update block when a sector written by a host is logically non-sequential to the previously written sector within the logical group being updated. All sectors subsequently updated in this logical group are written in the next available sector location in the chaotic update block, whatever their logical sector address within the group.

[00114] **FIG. 7B** illustrates an example of sectors in a logical group being written in chaotic order to a chaotic update block as a result of five separate host write operations, whilst superseded sectors in the original block for the logical group and duplicated sectors in the chaotic update block become obsolete. In host write operation #1, the logical sectors LS10-LS11 of a given logical group stored in an original metablock is updated. The updated logical sectors LS10'-LS11' are stored in

a newly allocated update block. At this point, the update block is a sequential one. In host write operation #2, the logical sectors LS5-LS6 are updated as LS5'-LS6' and recorded in the update block in the location immediately following the last write. This converts the update block from a sequential to a chaotic one. In host write operation #3, the logical sector LS10 is being updated again and is recorded in the next location of the update block as LS10''. At this point LS10'' in the update block supersedes LS10' in a previous recording which in turn supersedes LS10 in the original block. In host write operation #4, the data in the logical sector LS10 is again updated and is recorded in the next location of the update block as LS10'''. Thus, LS10''' is now the latest and only valid data for the logical sector LS10. In host write operation #5, the data in logical sector LS30 is being updated and recorded in the update block as LS30'. Thus, the example illustrates that sectors within a logical group can be written in a chaotic update block in any order and with any repetition.

Forced Sequential Update

[00115] **FIG. 8** illustrates an example of sectors in a logical group being written in sequential order to a sequential update block as a result of two separate host write operations that has a discontinuity in logical addresses. In host write #1, the update data in the logical sectors LS5-LS8 is recorded in a dedicated update block as LS5'-LS8'. In host write #2, the update data in the logical sectors LS14-LS16 is being recorded in the update block following the last write as LS14'-LS16'. However, there is an address jump between LS8 and LS14 and the host write #2 would normally render the update block non-sequential. Since the address jump is not substantial, one option is to first perform a padding operation (#2A) by copying the data of the intervening sectors from the original block to the update block before executing host write #2. In this way, the sequential nature of the update block is preserved.

[00116] **FIG. 9** is a flow diagram illustrating a process by the update block manager to update a logical group of data, according a general embodiment of the invention. The update process comprises the following steps:

[00117] **STEP 260:** The memory is organized into blocks, each block partitioned into memory units that are erasable together, each memory unit for storing

a logical unit of data.

[00118] **STEP 262:** The data is organized into logical groups, each logical group partitioned into logical units.

[00119] **STEP 264:** In the standard case, all logical units of a logical group is stored among the memory units of an original block according to a first prescribed order, preferably, in logically sequential order. In this way, the index for accessing the individual logical units in the block is known.

[00120] **STEP 270:** For a given logical group (e.g., LG_X) of data, a request is made to update a logical unit within LG_X . (A logical unit update is given as an example. In general the update will be a segment of one or more contiguous logical units within LG_X .)

[00121] **STEP 272:** The requested update logical unit is to be stored in a second block, dedicated to recording the updates of LG_X . The recording order is according to a second order, typically, the order the updates are requested. One feature of the invention allows an update block to be set up initially generic to recording data in logically sequential or chaotic order. So depending on the second order, the second block can be a sequential one or a chaotic one.

[00122] **STEP 274:** The second block continues to have requested logical units recorded as the process loops back to **STEP 270**. The second block will be closed to receiving further update when a predetermined condition for closure materializes. In that case, the process proceeds to **STEP 276**.

[00123] **STEP 276:** Determination is made whether or not the closed, second block has its update logical units recorded in a similar order as that of the original block. The two blocks are considered to have similar order when they recorded logical units differ by only a page tag, as described in connection with **FIG. 3A**. If the two blocks have similar order the process proceeds to **STEP 280**, otherwise, some sort of garbage collection need to be performed in **STEP 290**.

[00124] **STEP 280:** Since the second block has the same order as the first block, it is used to replace the original, **first** block. The update process then ends at

STEP 299.

[00125] **STEP 290:** The latest version of each logical units of the given logical group are gathered from among the second block (update block) and the first block (original block). The consolidated logical units of the given logical group are then written to a third block in an order similar to the first block.

[00126] **STEP 292:** Since the third block (consolidated block) has a similar order to the first block, it is used to replace the original, first block. The update process then ends at **STEP 299**.

[00127] **STEP 299:** When a closeout process creates an intact update block, it becomes the new standard block for the given logical group. The update thread for the logical group will be terminated.

[00128] **FIG. 10** is a flow diagram illustrating a process by the update block manager to update a logical group of data, according a preferred embodiment of the invention. The update process comprises the following steps:

[00129] **STEP 310:** For a given logical group (e.g., LG_X) of data, a request is made to update a logical sector within LG_X . (A sector update is given as an example. In general the update will be a segment of one or more contiguous logical sectors within LG_X .)

[00130] **STEP 312:** If an update block dedicated to LG_X does not already exist, proceed to **STEP 410** to initiate a new update thread for the logical group. This will be accomplished by allocating an update block dedicated to recording update data of the logical group. If there is already an update block open, proceed to **STEP 314** to begin recording the update sector onto the update block.

[00131] **STEP 314:** If the current update block is already chaotic (i.e., non-sequential) then simply proceed to **STEP 510** for recording the requested update sector onto the chaotic update block. If the current update block is sequential, proceed to **STEP 316** for processing of a sequential update block.

[00132] **STEP 316:** One feature of the invention allows an update block to be

set up initially generic to recording data in logically sequential or chaotic order. However, since the logical group ultimately has its data stored in a metablock in a logically sequential order, it is desirable to keep the update block sequential as far as possible. Less processing will then be required when an update block is closed to further updates as garbage collection will not be needed.

[00133] Thus determination is made whether the requested update will follow the current sequential order of the update block. If the update follows sequentially, then proceed to **STEP 510** to perform a sequential update, and the update block will remain sequential. On the other hand, if the update does not follow sequentially (chaotic update), it will convert the sequential update block to a chaotic one if no other actions are taken.

[00134] In one embodiment, nothing more is done to salvage the situation and the process proceeds directly to **STEP 370** where the update is allowed to turn the update block into a chaotic one.

Optional Forced Sequential Process

[00135] In another embodiment, a forced sequential process **STEP 320** is optionally performed to preserve the sequential update block as far as possible in view of a pending chaotic update. There are two situations, both of which require copying missing sectors from the original block to maintain the sequential order of logical sectors recorded on the update block. The first situation is where the update creates a short address jump. The second situation is to prematurely close out an update block in order to keep it sequential. The forced sequential process **STEP 320** comprises the following substeps:

[00136] **STEP 330:** If the update creates a logical address jump not greater a predetermined amount, C_B , the process proceeds to a forced sequential update process in **STEP 350**, otherwise the process proceeds to **STEP 340** to consider if it qualifies for a forced sequential closeout.

[00137] **STEP 340:** If the number of unfilled physical sectors exceeds a predetermined design parameter, C_C , whose typical value is half of the size of the

update block, then the update block is relatively unused and will not be prematurely closed. The process proceeds to **STEP 370** and the update block will become chaotic. On the other hand, if the update block is substantially filled, it is considered to have been well utilized already and therefore is directed to **STEP 360** for forced sequential closeout.

[00138] **STEP 350:** Forced sequential update allows current sequential update block to remain sequential as long as the address jump does not exceed a predetermined amount, C_B . Essentially, sectors from the update block's associated original block are copied to fill the gap spanned by the address jump. Thus, the sequential update block will be padded with data in the intervening addresses before proceeding to **STEP 510** to record the current update sequentially.

[00139] **STEP 360:** Forced sequential closeout allows the currently sequential update block to be closed out if it is already substantially filled rather than converted to a chaotic one by the pending chaotic update. A chaotic or non-sequential update is defined as one with a forward address transition not covered by the address jump exception described above, a backward address transition, or an address repetition. To prevent a sequential update block to be converted by a chaotic update, the unwritten sector locations of the update block are filled by copying sectors from the update block's associated original partly-obsolete block. The original block is then fully obsolete and may be erased. The current update block now has the full set of logical sectors and is then closed out as an intact metablock replacing the original metablock. The process then proceeds to **STEP 430** to have a new update block allocated in its place to accept the recording of the pending sector update that was first requested in **STEP 310**.

Conversion to Chaotic Update Block

[00140] **STEP 370:** When the pending update is not in sequential order and optionally, if the forced sequential conditions are not satisfied, the sequential update block is allowed to be converted to a chaotic one by virtue of allowing the pending update sector, with non-sequential address, to be recorded on the update block when the process proceeds to **STEP 510**. If the maximum number of chaotic update blocks exist, it is necessary to close the least recently accessed chaotic update block before

allowing the conversion to proceed; thus preventing the maximum number of chaotic blocks from being exceeded. The identification of the least recently accessed chaotic update block is the same as the general case described in **STEP 420**, but is constrained to chaotic update blocks only. Closing a chaotic update block at this time is achieved by consolidation as described in **STEP 550**.

Allocation of New Update Block subject to System Restriction

[00141] **STEP 410:** The process of allocating an erase metablock as an update block begins with the determination whether a predetermined system limitation is exceeded or not. Due to finite resources, the memory management system typically allows a predetermined maximum number of update blocks, U_{MAX} , to exist concurrently. This limit is the aggregate of sequential update blocks and chaotic update blocks, and is a design parameter. In a preferred embodiment, the limit is, for example, a maximum of 8 update blocks. Also, due to the higher demand on system resources, there may also be a corresponding predetermined limit on the maximum number of chaotic update blocks that can be open concurrently (e.g., 4.)

[00142] Thus, when U_{MAX} update blocks have already been allocated, then the next allocation request could only be satisfied after closing one of the existing allocated ones. The process proceeds to **STEP 420**. When the number of open update blocks is less than C_A , the process proceeds directly to **STEP 430**.

[00143] **STEP 420:** In the event the maximum number of update blocks, C_A , is exceeded, the least-recently accessed update block is closed and garbage collection is performed. The least recently accessed update block is identified as the update block associated with the logical block that has been accessed least recently. For the purpose of determining the least recently accessed blocks, an access includes writes and optionally reads of logical sectors. A list of open update blocks is maintained in order of access; at initialization, no access order is assumed. The closure of an update block follows along the similar process described in connection with **STEP 360** and **STEP 530** when the update block is sequential, and in connection with **STEP 540** when the update block is chaotic. The closure makes room for the allocation of a new update block in **STEP 430**.

[00144] **STEP 430:** The allocation request is fulfilled with the allocation of a new metablock as an update block dedicated to the given logical group LG_X. The process then proceeds to **STEP 510**.

Record Update Data onto Update Block

[00145] **STEP 510:** The requested update sector is recorded onto next available physical location of the update block. The process then proceeds to **STEP 520** to determine if the update block is ripe for closeout.

Update Block Closeout

[00146] **STEP 520:** If the update block still has room for accepting additional updates, proceed to **STEP 570**. Otherwise proceed to **STEP 522** to closeout the update block. There are two possible implementations of filling up an update block when the current requested write attempts to write more logical sectors than the block has room for. In the first implementation, the write request is split into two portions, with the first portion writing up to the last physical sector of the block. The block is then closed and the second portion of the write will be treated as the next requested write. In the other implementation, the requested write is withheld while the block has its remaining sectors padded and is then closed. The requested write will be treated as the next requested write.

[00147] **STEP 522:** If the update block is sequential, proceed to **STEP 530** for sequential closure. If the update block is chaotic, proceed to **STEP 540** for chaotic closure.

Sequential Update Block Closeout

[00148] **STEP 530:** Since the update block is sequential and fully filled, the logical group stored in it is intact. The metablock is intact and replaces the original one. At this time, the original block is fully obsolete and may be erased. The process then proceeds to **STEP 570** where the update thread for the given logical group ends.

Chaotic Update Block Closeout

[00149] **STEP 540:** Since the update block is non-sequentially filled and may

contain multiple updates of some logical sectors, garbage collection is performed to salvage the valid data in it. The chaotic update block will either be compacted or consolidated. Which process to perform will be determined in **STEP 542**.

[00150] **STEP 542:** To perform compaction or consolidation will depend on the degeneracy of the update block. If a logical sector is updated multiple times, its logical address is highly degenerate. There will be multiple versions of the same logical sector recorded on the update block and only the last recorded version is the valid one for that logical sector. In an update block containing logical sectors with multiple versions, the number of distinct logical sectors will be much less than that of a logical group.

[00151] In the preferred embodiment, when the number of distinct logical sectors in the update block exceeds a predetermined design parameter, C_D , whose typical value is half of the size of a logical group, the closeout process will perform a consolidation in **STEP 550**, otherwise the process will proceed to compaction in **STEP 560**.

[00152] **STEP 550:** If the chaotic update block is to be consolidated, the original block and the update block will be replaced by a new standard metablock containing the consolidated data. After consolidation the update thread will end in **STEP 570**.

[00153] **STEP 560:** If the chaotic update block is to be compacted, it will be replaced by a new update block carrying the compacted data. After compaction the processing of the compacted update block will end in **STEP 570**. Alternatively, compaction can be delayed until the update block is written to again, thus removing the possibility of compaction being followed by consolidation without intervening updates. The new update block will then be used in further updating of the given logical block when a next request for update in LG_X appears in **STEP 502**.

[00154] **STEP 570:** When a closeout process creates an intact update block, it becomes the new standard block for the given logical group. The update thread for the logical group will be terminated. When a closeout process creates a new update block replacing an existing one, the new update block will be used to record the next

update requested for the given logical group. When an update block is not closed out, the processing will continue when a next request for update in LG_X appears in **STEP 310**.

[00155] As can be seen from the process described above, when a chaotic update block is closed, the update data recorded on it is further processed. In particular its valid data is garbage collected either by a process of compaction to another chaotic block, or by a process of consolidation with its associated original block to form a new standard sequential block.

[00156] **FIG. 11A** is a flow diagram illustrating in more detail the consolidation process of closing a chaotic update block shown in **FIG. 10**. Chaotic update block consolidation is one of two possible processes performed when the update block is being closed out, e.g., when the update block is full with its last physical sector location written. Consolidation is chosen when the number of distinct logical sectors written in the block exceeds a predetermined design parameter, C_D . The consolidation process **STEP 550** shown in **FIG. 10** comprises the following substeps:

[00157] **STEP 551:**When a chaotic update block is being closed, a new metablock replacing it will be allocated.

[00158] **STEP 552:** Gather the latest version of each logical sector among the chaotic update block and its associated original block, ignoring all the obsolete sectors.

[00159] **STEP 554:** Record the gathered valid sectors onto the new metablock in logically sequential order to form an intact block, i.e., a block with all the logical sectors of a logical group recorded in sequential order.

[00160] **STEP 556:** Replace the original block with the new intact block.

[00161] **STEP 558:** Erase the closed out update block and the original block.

[00162] **FIG. 11B** is a flow diagram illustrating in more detail the compaction process for closing a chaotic update block shown in **FIG. 10**. Compaction is chosen

when the number of distinct logical sectors written in the block is below a predetermined design parameter, C_D . The compaction process **STEP 560** shown in **FIG. 10** comprises the following substeps:

[00163] **STEP 561:** When a chaotic update block is being compacted, a new metablock replacing it will be allocated.

[00164] **STEP 562:** Gather the latest version of each logical sector among the existing chaotic update block to be compacted.

[00165] **STEP 564:** Record the gathered sectors onto the new update block to form a new update block having compacted sectors.

[00166] **STEP 566:** Replace the existing update block with the new update block having compacted sectors.

[00167] **STEP 568:** Erase the closed out update block

Logical and Metablock States

[00168] **FIG. 12A** illustrates all possible states of a Logical Group, and the possible transitions between them under various operations.

[00169] **FIG. 12B** is a table listing the possible states of a Logical Group. The Logical Group states are defined as follows:

1. *Intact:* All logical sectors in the Logical Group have been written in logically sequential order, possibly using page tag wrap around, in a single metablock.
2. *Unwritten:* No logical sector in the Logical Group has ever been written. The Logical Group is marked as unwritten in a group address table and has no allocated metablock. A predefined data pattern is returned in response to a host read for every sector within this group.
3. *Sequential Update:* Some sectors within the Logical Group have been written in logically sequential order in a metablock, possibly using page tag, so that they supersede the corresponding logical sectors from any previous Intact state of the group.

4. *Chaotic Update*: Some sectors within the Logical Group have been written in logically non-sequential order in a metablock, possibly using page tag, so that they supersede the corresponding logical sectors from any previous Intact state of the group. A sector within the group may be written more than once, with the latest version superseding all previous versions.

[00170] FIG. 13A illustrates all possible states of a metablock, and the possible transitions between them under various operations.

[00171] FIG. 13B is a table listing the possible states of a metablock. The metablock states are defined as follows:

1. *Erased*: All the sectors in the metablock are erased.
2. *Sequential Update*: The metablock is partially written with sectors in logically sequential order, possibly using page tag. All the sectors belong to the same Logical Group.
3. *Chaotic Update*: The metablock is partially or fully written with sectors in logically non-sequential order. Any sector can be written more than once. All sectors belong to the same Logical Group.
4. *Intact*: The metablock is fully written in logically sequential order, possibly using page tag.
5. *Original*: The metablock was previously Intact but at least one sector has been made obsolete by a host data update.

[00172] FIGS. 14(A) – 14(J) are state diagrams showing the effect of various operations on the state of the logical group and also on the physical metablock.

[00173] FIG. 14(A) shows state diagrams corresponding to the logical group and the metablock transitions for a first write operation. The host writes one or more sectors of a previously unwritten Logical Group in logically sequential order to a newly allocated Erased metablock. The Logical Group and the metablock go to the Sequential Update state.

[00174] FIG. 14(B) shows state diagrams corresponding to the logical group and the metablock transitions for a first intact operation. A previously unwritten

Sequential Update Logical Group becomes Intact as all the sectors are written sequentially by the host. The transition can also happen if the card fills up the group by filling the remaining unwritten sectors with a predefined data pattern. The metablock becomes Intact.

[00175] **FIG. 14(C)** shows state diagrams corresponding to the logical group and the metablock transitions for a first chaotic operation. A previously unwritten Sequential Update Logical Group becomes Chaotic when at least one sector has been written non-sequentially by the host.

[00176] **FIG. 14(D)** shows state diagrams corresponding to the logical group and the metablock transitions for a first compaction operation. All valid sectors within a previously unwritten Chaotic Update Logical Group are copied to a new Chaotic metablock from the old block, which is then erased.

[00177] **FIG. 14(E)** shows state diagrams corresponding to the logical group and the metablock transitions for a first consolidation operation. All valid sectors within a previously unwritten Chaotic Update Logical Group are moved from the old Chaotic block to fill a newly allocated Erased block in logically sequential order. Sectors unwritten by the host are filled with a predefined data pattern. The old chaotic block is then erased.

[00178] **FIG. 14(F)** shows state diagrams corresponding to the logical group and the metablock transitions for a sequential write operation. The host writes one or more sectors of an Intact Logical Group in logically sequential order to a newly allocated Erased metablock. The Logical Group and the metablock go to Sequential Update state. The previously Intact metablock becomes an Original metablock.

[00179] **FIG. 14(G)** shows state diagrams corresponding to the logical group and the metablock transitions for a sequential fill operation. A Sequential Update Logical Group becomes Intact when all its sectors are written sequentially by the host. This may also occur during garbage collection when the Sequential Update Logical Group is filled with valid sectors from the original block in order to make it Intact, after which the original block is erased.

[00180] FIG. 14(H) shows state diagrams corresponding to the logical group and the metablock transitions for a non-sequential write operation. A Sequential Update Logical Group becomes Chaotic when at least one sector is written non-sequentially by the host. The non-sequential sector writes may cause valid sectors in either the Update block or the corresponding Original block to become obsolete.

[00181] FIG. 14(I) shows state diagrams corresponding to the logical group and the metablock transitions for a compaction operation. All valid sectors within a Chaotic Update Logical Group are copied into a new chaotic metablock from the old block, which is then erased. The Original block is unaffected.

[00182] FIG. 14(J) shows state diagrams corresponding to the logical group and the metablock transitions for a consolidation operation. All valid sectors within a Chaotic Update Logical Group are copied from the old chaotic block and the Original block to fill a newly allocated Erased block in logically sequential order. The old chaotic block and the Original block are then erased.

UPDATE BLOCK TRACKING AND MANAGEMENT

[00183] FIG. 15 illustrates a preferred embodiment of the structure of an allocation block list (ABL) for keeping track of opened and closed update blocks and erased blocks for allocation. The allocation block list (ABL) 610 is held in controller RAM 130, to allow management of allocation of erased blocks, allocated update blocks, associated blocks and control structures, and to enable correct logical to physical address translation. In the preferred embodiment, the ABL includes a list of erased blocks, an open update block list 614 and a closed update block list 616.

[00184] The open update block list 614 is the set of block entries in the ABL with the attributes of Open Update Block. The open update block list has one entry for each data update block currently open. Each entry holds the following information. LG is the logical group address the current update metablock is dedicated to. Sequential/Chaotic is a status indicating whether the update block has been filled with sequential or chaotic update data. MB is the metablock address of the update block. Page tag is the starting logical sector recorded at the first physical location of the update block. Number of sectors written indicates the number of

sectors currently written onto the update block. MB_0 is the metablock address of the associated original block. Page Tag₀ is the page tag of the associated original block.

[00185] The closed update block list 616 is a subset of the Allocation Block List (ABL). It is the set of block entries in the ABL with the attributes of Closed Update Block. The closed update block list has one entry for each data update block which has been closed, but whose entry has not been updated in a logical to a main physical directory. Each entry holds the following information. LG is the logical group address the current update block is dedicated to. MB is the metablock address of the update block. Page tag is the starting logical sector recorded at the first physical location of the update block. MB_0 is the metablock address of the associated original block.

CHAOTIC BLOCK INDEXING

[00186] A sequential update block has the data stored in logically sequential order, thus any logical sector among the block can be located easily. A chaotic update block has its logical sectors stored out of order and may also store multiple update generations of a logical sector. Additional information must be maintained to keep track of where each valid logical sector is located in the chaotic update block.

[00187] In the preferred embodiment, chaotic block indexing data structures allow tracking and fast access of all valid sectors in a chaotic block. Chaotic block indexing independently manages small regions of logical address space, and efficiently handles system data and hot regions of user data. The indexing data structures essentially allow indexing information to be maintained in flash memory with infrequent update requirement so that performance is not significantly impacted. On the other hand, lists of recently written sectors in chaotic blocks are held in a chaotic sector list in controller RAM. Also, a cache of index information from flash memory is held in controller RAM in order to minimize the number of flash sector accesses for address translation. Indexes for each chaotic block are stored in chaotic block index (CBI) sectors in flash memory.

[00188] **FIG. 16A** illustrates the data fields of a chaotic block index (CBI) sector. A Chaotic Block Index Sector (CBI sector) contains an index for each sector

in a logical group mapped to a chaotic update block, defining the location of each sector of the logical group within the chaotic update block or its associated original block. A CBI sector includes a chaotic block index field for keeping track of valid sectors within the chaotic block, a chaotic block info field for keeping track of address parameters for the chaotic block, and a sector index field for keeping track of the valid CBI sectors within the metablock (CBI block) storing the CBI sectors.

[00189] FIG. 16B illustrates an example of the chaotic block index (CBI) sectors being recorded in a dedicated metablock. The dedicated metablock will be referred to as a CBI block 620. When a CBI sector is updated, it is written in the next available physical sector location in the CBI block 620. Multiple copies of a CBI sector may therefore exist in the CBI block, with only the last written copy being valid. For example the CBI sector for the logical group LG₁ has been updated three times with the latest version being the valid one. The location of each valid sector in the CBI block is identified by a set of indices in the last written CBI sector in the block. In this example, the last written CBI sector in the block is CBI sector for LG₁₃₆ and its set of indices is the valid one superceding all previous ones. When the CBI block eventually becomes fully filled with CBI sectors, the block is compacted during a control write operation by rewriting all valid sectors to a new block location. The full block is then erased.

[00190] The chaotic block index field within a CBI sector contains an index entry for each logical sector within a logical group or sub-group mapped to a chaotic update block. Each index entry signifies an offset within the chaotic update block at which valid data for the corresponding logical sector is located. A reserved index value indicates that no valid data for the logical sector exists in the chaotic update block, and that the corresponding sector in the associated original block is valid. A cache of some chaotic block index field entries is held in controller RAM.

[00191] The chaotic block info field within a CBI sector contains one entry for each chaotic update block that exists in the system, recording address parameter information for the block. Information in this field is only valid in the last written sector in the CBI block. This information is also present in data structures in RAM.

[00192] The entry for each chaotic update block includes three address

parameters. The first is the logical address of the logical group (or logical group number) associated with the chaotic update block. The second is the metablock address of the chaotic update block. The third is the physical address offset of the last sector written in the chaotic update block. The offset information sets the start point for scanning of the chaotic update block during initialization, to rebuild data structures in RAM.

[00193] The sector index field contains an entry for each valid CBI sector in the CBI block. It defines the offsets within the CBI block at which the most recently written CBI sectors relating to each permitted chaotic update block are located. A reserved value of an offset in the index indicates that a permitted chaotic update block does not exist.

[00194] **FIG. 16C** is a flow diagram illustrating access to the data of a logical sector of a given logical group undergoing chaotic update. During the update process, the update data is recorded in the chaotic update block while the unchanged data remains in the original metablock associated with the logical group. The process of accessing a logical sector of the logical group under chaotic update is as follows:

[00195] **STEP 650:** Begin locating a given logical sector of a given logical group.

[00196] **STEP 652:** Locate last written CBI sector in the CBI block

[00197] **STEP 654:** Locate the chaotic update block or original block associated with the given logical group by looking up the Chaotic Block Info field of the last written CBI sector. This step can be performed any time just before **STEP 662**.

[00198] **STEP 658:** If the last written CBI sector is directed to the given logical group, the CBI sector is located. Proceed to **STEP 662**. Otherwise, proceed to **STEP 660**.

[00199] **STEP 660:** Locate the CBI sector for the given logical group by looking up the sector index field of the last written CBI sector.

[00200] **STEP 662:** Locate the given logical sector among either the chaotic block or the original block by looking up the Chaotic Block Index field of the located CBI sector.

[00201] **FIG. 16D** is a flow diagram illustrating access to the data of a logical sector of a given logical group undergoing chaotic update, according to an alternative embodiment in which logical group has been partitioned into subgroups. The finite capacity of a CBI sector can only keep track of a predetermined maximum number of logical sectors. When the logical group has more logical sectors than a single CBI sector can handle, the logical group is partitioned into multiple subgroups with a CBI sector assigned to each subgroup. In one example, each CBI sector has enough capacity for tracking a logical group consisting of 256 sectors and up to 8 chaotic update blocks. If the logical group has a size exceeding 256 sectors, a separate CBI sector exists for each 256-sector sub-group within the logical group. CBI sectors may exist for up to 8 sub-groups within a logical group, giving support for logical groups up to 2048 sectors in size.

[00202] In the preferred embodiment, an indirect indexing scheme is employed to facilitate management of the index. Each entry of the sector index has direct and indirect fields.

[00203] The direct sector index defines the offsets within the CBI block at which all possible CBI sectors relating to a specific chaotic update block are located. Information in this field is only valid in the last written CBI sector relating to that specific chaotic update block. A reserved value of an offset in the index indicates that the CBI sector does not exist because the corresponding logical subgroup relating to the chaotic update block either does not exist, or has not been updated since the update block was allocated.

[00204] The indirect sector index defines the offsets within the CBI block at which the most recently written CBI sectors relating to each permitted chaotic update block are located. A reserved value of an offset in the index indicates that a permitted chaotic update block does not exist.

[00205] **FIG. 16D** shows the process of accessing a logical sector of the

logical group under chaotic update as follows:

[00206] **STEP 670:** Partition each Logical Group into multiple subgroups and assign a CBI sector to each subgroup

[00207] **STEP 680:** Begin locating a given logical sector of a given subgroup of a given logical group.

[00208] **STEP 682:** Locate the last written CBI sector in the CBI block.

[00209] **STEP 684:** Locate the chaotic update block or original block associated with the given subgroup by looking up the Chaotic Block Info field of the last written CBI sector. This step can be performed any time just before **STEP 696**.

[00210] **STEP 686:** If the last written CBI sector is directed to the given logical group, proceed to **STEP 691**. Otherwise, proceed to **STEP 690**.

[00211] **STEP 690:** Locate the last written of the multiple CBI sectors for the given logical group by looking up the Indirect Sector Index field of the last written CBI sector.

[00212] **STEP 691:** At least a CBI sector associate with one of the subgroups for the given logical group has been located. Continue.

[00213] **STEP 692:** If the located CBI sector directed to the given subgroup, the CBI sector for the given subgroup is located. Proceed to **STEP 696**. Otherwise, proceed to **STEP 694**.

[00214] **STEP 694:** Locate the CBI sector for the given subgroup by looking up the direct sector index field of the currently located CBI sector.

[00215] **STEP 696:** Locate the given logical sector among either the chaotic block or the original block by looking up the Chaotic Block Index field of the CBI sector for the given subgroup.

[00216] **FIG. 16E** illustrates examples of Chaotic Block Indexing (CBI) sectors and their functions for the embodiment where each logical group is partitioned into multiple subgroups. A logical group 700 originally has its intact data stored in an

original metablock 702. The logical group is then undergoing updates with the allocation of a dedicated chaotic update block 704. In the present examples, the logical group 700 is partitioned into subgroups, such subgroups A, B, C, D, each having 256 sectors.

[00217] In order to locate the i th sector in the subgroup B, the last written CBI sector in the CBI block 620 is first located. The chaotic block info field of the last written CBI sector provides the address to locate the chaotic update block 704 for the given logical group. At the same time it provides the location of the last sector written in the chaotic block. This information is useful in the event of scanning and rebuilding indices.

[00218] If the last written CBI sector turns out to be one of the four CBI sectors of the given logical group, it will be further determined if it is exactly the CBI sector for the given subgroup B that contains the i th logical sector. If it is, then the CBI sector's chaotic block index will point to the metablock location for storing the data for the i th logical sector. The sector location could be either in the chaotic update block 704 or the original block 702.

[00219] If the last written CBI sector turns out to be one of the four CBI sectors of the given logical group but is not exactly for the subgroup B, then its direct sector index is looked up to locate the CBI sector for the subgroup B. Once this exact CBI sector is located, its chaotic block index is looked up to locate the i th logical sector among the chaotic update block 704 and the original block 702.

[00220] If the last written CBI sector turns out not to be anyone of the four CBI sectors of the given logical group, its indirect sector index is looked up to locate one of the four. In the example shown in **FIG. 16E**, the CBI sector for subgroup C is located. Then this CBI sector for subgroup C has its direct sector index looked up to locate the exact CBI sector for the subgroup B. The example shows that when its chaotic block index is looked up, the i th logical sector is found to be unchanged and its valid data will be located in the original block.

[00221] Similar consideration applies to locating the j th logical sector in subgroup C of the given logical group. The example shows that the last written CBI

sector turns out not to be any one of the four CBI sectors of the given logical group. Its indirect sector index points to one of the four CBI sectors for the given group. The last written of four pointed to also turns out to be exactly the CBI sector for the subgroup C. When its chaotic block index is looked up, the j th logical sector is found to be located at a designated location in the chaotic update block 704.

[00222] A list of chaotic sectors exists in controller RAM for each chaotic update block in the system. Each list contains a record of sectors written in the chaotic update block since a related CBI sector was last updated in flash memory. The number of logical sector addresses for a specific chaotic update block, which can be held in a chaotic sector list, is a design parameter with a typical value of 8 to 16. The optimum size of the list is determined as a tradeoff between its effects on overhead for chaotic data-write operations and sector scanning time during initialization.

[00223] During system initialization, each chaotic update block is scanned as necessary to identify valid sectors written since the previous update of one of its associated CBI sectors. A chaotic sector list in controller RAM for each chaotic update block is constructed. Each block need only be scanned from the last sector address defined in its chaotic block info field in the last written CBI sector.

[00224] When a chaotic update block is allocated, a CBI sector is written to correspond to all updated logical sub-groups. The logical and physical addresses for the chaotic update block are written in an available chaotic block info field in the sector, with null entries in the chaotic block index field. A chaotic sector list is opened in controller RAM.

[00225] When a chaotic update block is closed, a CBI sector is written with the logical and physical addresses of the block removed from the chaotic block info field in the sector. The corresponding chaotic sector list in RAM becomes unused.

[00226] The corresponding chaotic sector list in controller RAM is modified to include records of sectors written to a chaotic update block. When a chaotic sector list in controller RAM has no available space for records of further sector writes to a chaotic update block, updated CBI sectors are written for logical sub-groups relating

to sectors in the list, and the list is cleared.

[00227] When the CBI block 620 becomes full, valid CBI sectors are copied to an allocated erased block, and the previous CBI block is erased.

ADDRESS TABLES

[00228] The logical to physical address translation module 140 shown in **FIG. 2** is responsible for relating a host's logical address to a corresponding physical address in flash memory. Mapping between logical groups and physical groups (metablocks) are stored in a set of table and lists distributed among the nonvolatile flash memory 200 and the volatile but more agile RAM 130 (see **FIG. 1**.) An address table is maintained in flash memory, containing a metablock address for every logical group in the memory system. In addition, logical to physical address records for recently written sectors are temporarily held in RAM. These volatile records can be reconstructed from block lists and data sector headers in flash memory when the system is initialized after power-up. Thus, the address table in flash memory need be updated only infrequently, leading to a low percentage of overhead write operations for control data.

[00229] The hierarchy of address records for logical groups includes the open update block list, the closed update block list in RAM and the group address table (GAT) maintained in flash memory.

[00230] The open update block list is a list in controller RAM of data update blocks which are currently open for writing updated host sector data. The entry for a block is moved to the closed update block list when the block is closed. The closed update block list is a list in controller RAM of data update blocks which have been closed. A subset of the entries in the list is moved to a sector in the Group Address Table during a control write operation.

[00231] The Group Address Table (GAT) is a list of metablock addresses for all logical groups of host data in the memory system. The GAT contains one entry for each logical group, ordered sequentially according to logical address. The n th entry in the GAT contains the metablock address for the logical group with address n . In

the preferred embodiment, it is a table in flash memory, comprising a set of sectors (referred to as GAT sectors) with entries defining metablock addresses for every logical group in the memory system. The GAT sectors are located in one or more dedicated control blocks (referred to as GAT blocks) in flash memory.

[00232] **FIG. 17A** illustrates the data fields of a group address table (GAT) sector. A GAT sector may for example have sufficient capacity to contain GAT entries for a set of 128 contiguous logical groups. Each GAT sector includes two components, namely a set of GAT entries for the metablock address of each logical group within a range, and a GAT sector index. The first component contains information for locating the metablock associated with the logical address. The second component contains information for locating all valid GAT sectors within the GAT block. Each GAT entry has three fields, namely, the metablock number, the page tag as defined earlier in connection with **FIG. 3A(iii)**, and a flag indicating whether the metablock has been relinked. The GAT sector index lists the positions of valid GAT sectors in a GAT block. This index is in every GAT sector but is superseded by the version of the next written GAT sector in the GAT block. Thus only the version in the last written GAT sector is valid.

[00233] **FIG. 17B** illustrates an example of the group address table (GAT) sectors being recorded in one or more GAT block. A GAT block is a metablock dedicated to recording GAT sectors. When a GAT sector is updated, it is written in the next available physical sector location in the GAT block 720. Multiple copies of a GAT sector may therefore exist in the GAT block, with only the last written copy being valid. For example the GAT sector 255 (containing pointers for the logical groups LG₃₉₆₈ - LG₄₀₉₈) has been updated at least two times with the latest version being the valid one. The location of each valid sector in the GAT block is identified by a set of indices in the last written GAT sector in the block. In this example, the last written GAT sector in the block is GAT sector 236 and its set of indices is the valid one superseding all previous ones. When the GAT block eventually becomes fully filled with GAT sectors, the block is compacted during a control write operation by rewriting all valid sectors to a new block location. The full block is then erased.

[00234] As described earlier, a GAT block contains entries for a logically

contiguous set of groups in a region of logical address space. GAT sectors within a GAT block each contain logical to physical mapping information for 128 contiguous logical groups. The number of GAT sectors required to store entries for all logical groups within the address range spanned by a GAT block occupy only a fraction of the total sector positions in the block. A GAT sector may therefore be updated by writing it at the next available sector position in the block. An index of all valid GAT sectors and their position in the GAT block is maintained in an index field in the most recently written GAT sector. The fraction of the total sectors in a GAT block occupied by valid GAT sectors is a system design parameter, which is typically 25%. However, there is a maximum of 64 valid GAT sectors per GAT block. In systems with large logical capacity, it may be necessary to store GAT sectors in more than one GAT block. In this case, each GAT block is associated with a fixed range of logical groups.

[00235] A GAT update is performed as part of a control write operation, which is triggered when the ABL runs out of blocks for allocation (see **FIG. 18.**) It is performed concurrently with ABL fill and CBL empty operations. During a GAT update operation, one GAT sector has entries updated with information from corresponding entries in the closed update block list. When a GAT entry is updated, any corresponding entries are removed from the closed update block list (CUBL). For example, the GAT sector to be updated is selected on the basis of the first entry in the closed update block list. The updated sector is written to the next available sector location in the GAT block.

[00236] A GAT rewrite operation occurs during a control write operation when no sector location is available for an updated GAT sector. A new GAT block is allocated, and valid GAT sectors as defined by the GAT index are copied in sequential order from the full GAT block. The full GAT block is then erased.

[00237] A GAT cache is a copy in controller RAM 130 of entries in a subdivision of the 128 entries in a GAT sector. The number of GAT cache entries is a system design parameter, with typical value 32. A GAT cache for the relevant sector subdivision is created each time an entry is read from a GAT sector. Multiple GAT caches are maintained. The number is a design parameter with a typical value of 4. A

GAT cache is overwritten with entries for a different sector subdivision on a least-recently-used basis.

ERASED METABLOCK MANAGEMENT

[00238] The erase block manager 160 shown in **FIG. 2** manages erase blocks using a set of lists for maintaining directory and system control information. These lists are distributed among the controller RAM 130 and flash memory 200. When an erased metablock must be allocated for storage of user data, or for storage of system control data structures, the next available metablock number in the allocation block list (ABL) (see **FIG. 15**) held in controller RAM is selected. Similarly, when a metablock is erased after it has been retired, its number is added to a cleared block list (CBL) also held in controller RAM. Relatively static directory and system control data are stored in flash memory. These include erased block lists and a bitmap (MAP) listing the erased status of all metablocks in the flash memory. The erased block lists and MAP are stored in individual sectors and are recorded to a dedicated metablock, known as a MAP block. These lists, distributed among the controller RAM and flash memory, provide a hierarchy of erased block records to efficiently manage erased metablock usage.

[00239] **FIG. 18** is a schematic block diagram illustrating the distribution and flow of the control and directory information for usage and recycling of erased blocks. The control and directory data are maintained in lists which are held either in controller RAM 130 or in a MAP block 750 residing in flash memory 200.

[00240] In the preferred embodiment, the controller RAM 130 holds the allocation block list (ABL) 610 and a cleared block list (CBL) 740. As described earlier in connection with **FIG. 15**, the allocation block list (ABL) keeps track of which metablocks have recently been allocated for storage of user data, or for storage of system control data structures. When a new erased metablock need be allocated, the next available metablock number in the allocation block list (ABL) is selected. Similarly, the cleared block list (CBL) is used to keep track of update metablocks that have been de-allocated and erased. The ABL and CBL are held in controller RAM 130 (see **FIG. 1**) for speedy access and easy manipulation when tracking the relatively active update blocks.

[00241] The allocation block list (ABL) keeps track of a pool of erased metablocks and the allocation of the erased metablocks to be an update block. Thus, each of these metablocks that may be described by an attribute designating whether it is an erased block in the ABL pending allocation, an open update block, or a closed update block. **FIG. 18** shows the ABL containing an erased ABL list 612, the open update block list 614 and the closed update block list 616. In addition, associated with the open update block list 614 is the associated original block list 615. Similarly, associated with the closed update block list is the associated erased original block list 617. As shown previously in **FIG. 15**, these associated lists are subset of the open update block list 614 and the closed update block list 616 respectively. The erased ABL block list 612, the open update block list 614, and the closed update block list 616 are all subsets of the allocation block list (ABL) 610, the entries in each having respectively the corresponding attribute.

[00242] The MAP block 750 is a metablock dedicated to storing erase management records in flash memory 200. The MAP block stores a time series of MAP block sectors, with each MAP sector being either an erase block management (EBM) sector 760 or a MAP sector 780. As erased blocks are used up in allocation and recycled when a metablock is retired, the associated control and directory data is preferably contained in a logical sector which may be updated in the MAP block, with each instance of update data being recorded to a new block sector. Multiple copies of EBM sectors 760 and MAP sectors 780 may exist in the MAP block 750, with only the latest version being valid. An index to the positions of valid MAP sectors is contained in a field in the EMB block. A valid EMB sector is always written last in the MAP block during a control write operation. When the MAP block 750 is full, it is compacted during a control write operation by rewriting all valid sectors to a new block location. The full block is then erased.

[00243] Each EBM sector 760 contains erased block lists (EBL) 770, which are lists of addresses of a subset of the population of erased blocks. The erased block lists (EBL) 770 act as a buffer containing erased metablock numbers, from which metablock numbers are periodically taken to re-fill the ABL, and to which metablock numbers are periodically added to re-empty the CBL. The EBL 770 serves as buffers for the available block buffer (ABB) 772, the erased block buffer (EBB) 774 and the

cleared block buffer (CBB) 776.

[00244] The available block buffer (ABB) 772 contains a copy of the entries in the ABL 610 immediately following the previous ABL fill operation. It is in effect a backup copy of the ABL just after an ABL fill operation..

[00245] The erased block buffer (EBB) 774 contains erased block addresses which have been previously transferred either from MAP sectors 780 or from the CBB list 776 (described below), and which are available for transfer to the ABL 610 during an ABL fill operation.

[00246] The cleared block buffer (CBB) 776 contains addresses of erased blocks which have been transferred from the CBL 740 during a CBL empty operation and which will be subsequently transferred to MAP sectors 780 or to the EBB list 774.

[00247] Each of the MAP sectors 780 contains a bitmap structure referred to as MAP. The MAP uses one bit for each metablock in flash memory, which is used to indicate the erase status of each block. Bits corresponding to block addresses listed in the ABL, CBL, or erased block lists in the EBM sector are not set to the erased state in the MAP.

[00248] Any block which does not contain valid data structures and which is not designated as an erased block within the MAP, erased block lists, ABL or CBL is never used by the block allocation algorithm and is therefore inaccessible for storage of host or control data structures. This provides a simple mechanism for excluding blocks with defective locations from the accessible flash memory address space.

[00249] The hierarchy shown in **FIG. 18** allows erased block records to be managed efficiently and provides full security of the block address lists stored in the controller's RAM. Erased block entries are exchanged between these block address lists and one or more MAP sectors 780, on an infrequent basis. These lists may be reconstructed during system initialization after a power-down, via information in the erased block lists and address translation tables stored in sectors in flash memory, and limited scanning of a small number of referenced data blocks in flash memory.

[00250] The algorithms adopted for updating the hierarchy of erased metablock records results in erased blocks being allocated for use in an order which interleaves bursts of blocks in address order from the MAP block 750 with bursts of block addresses from the CBL 740 which reflect the order blocks were updated by the host. For most metablock sizes and system memory capacities, a single MAP sector can provide a bitmap for all metablocks in the system. In this case, erased blocks are always allocated for use in address order as recorded in this MAP sector.

Erase Block Management Operations

[00251] As described earlier, the ABL 610 is a list with address entries for erased metablocks which may be allocated for use, and metablocks which have recently been allocated as data update blocks. The actual number of block addresses in the ABL lies between maximum and minimum limits, which are system design variables. The number of ABL entries formatted during manufacturing is a function of the card type and capacity. In addition, the number of entries in the ABL may be reduced near the end of life of the system, as the number of available erased blocks is reduced by failure of blocks during life. For example, after a fill operation, entries in the ABL may designate blocks available for the following purposes. Entries for Partially written data update blocks with one entry per block, not exceeding a system limit for a maximum of concurrently opened update blocks. Between one to twenty entries for Erased blocks for allocation as data update blocks. Four entries for erased blocks for allocation as control blocks.

ABL Fill Operation

[00252] As the ABL 610 becomes depleted through allocations, it will need to be refilled. An operation to fill the ABL occurs during a control write operation. This is triggered when a block must be allocated, but the ABL contains insufficient erased block entries available for allocation as a data update block, or for some other control data update block. During a control write, the ABL fill operation is concurrent with a GAT update operation.

[00253] The following actions occur during an ABL fill operation.

1. ABL entries with attributes of current data update blocks are retained.
2. ABL entries with attributes of closed data update blocks are retained, unless an entry for the block is being written in the concurrent GAT update operation, in which case the entry is removed from the ABL.
3. ABL entries for unallocated erase blocks are retained.
4. The ABL is compacted to remove gaps created by removal of entries, maintaining the order of entries.
5. The ABL is completely filled by appending the next available entries from the EBB list.
6. The ABB list is over-written with the current entries in the ABL.

CBL Empty Operation

[00254] The CBL is a list of erased block addresses in controller RAM with the same limitation on the number of erased block entries as the ABL. An operation to empty the CBL occurs during a control write operation. It is therefore concurrent with an ABL fill / GAT update operations, or CBI block write operations. In a CBL empty operation, entries are removed from the CBL 740 and written to the CBB list 776.

MAP Exchange Operation

[00255] A MAP exchange operation between the erase block information in the MAP sectors 780 and the EBM sectors 760 may occur periodically during a control write operation, when the EBB list 774 is empty. If all erased metablocks in the system are recorded in the EBM sector 760, no MAP sector 780 exists and no MAP exchange is performed. During a MAP exchange operation, a MAP sector feeding the EBB 774 with erased blocks is regarded as a source MAP sector 782. Conversely, a MAP sector receiving erased blocks from the CBB 776 is regarded as a destination MAP sector 784. If only one MAP sector exists, it acts as both source and destination MAP sector, as defined below.

[00256] The following actions are performed during a MAP exchange.

1. A source MAP sector is selected, on the basis of an incremental pointer.
2. A destination MAP sector is selected, on the basis of the block address in the first CBB entry that is not in the source MAP sector.
3. The destination MAP sector is updated, as defined by relevant entries in the CBB, and the entries are removed from the CBB.
4. The updated destination MAP sector is written in the MAP block, unless no separate source MAP sector exists.
5. The source MAP sector is updated, as defined by relevant entries in the CBB, and the entries are removed from the CBB.
6. Remaining entries in the CBB are appended to the EBB.
7. The EBB is filled to the extent possible with erased block addresses defined from the source MAP sector.
8. The updated source MAP sector is written in the MAP block.
9. An updated EBM sector is written in the MAP block.

List Management

[00257] **FIG. 18** shows the distribution and flow of the control and directory information between the various lists. For expediency, operations to move entries between elements of the lists or to change the attributes of entries, identified in **FIG. 18** as [A] to [O], are as follows.

[A] When an erased block is allocated as an update block for host data, the attributes of its entry in the ABL are changed from Erased ABL Block to Open Update Block.

[B] When an erased block is allocated as a control block, its entry in the ABL is removed.

[C] When an ABL entry is created with Open Update Block attributes, an Associated Original Block field is added to the entry to record the original metablock address for the logical group being updated. This information is obtained from the GAT.

[D] When an update block is closed, the attributes of its entry in the ABL are changed from Open Update Block to Closed Update Block.

[E] When an update block is closed, its associated original block is erased and the attributes of the Associated Original Block field in its entry in the ABL are changed to Erased Original Block.

[F] During an ABL fill operation, any closed update block whose address is updated in the GAT during the same control write operation has its entry removed from the ABL.

[G] During an ABL fill operation, when an entry for a closed update block is removed from the ABL, an entry for its associated erased original block is moved to the CBL.

[H] When a control block is erased, an entry for it is added to the CBL.

[I] During an ABL fill operation, erased block entries are moved to the ABL from the EBB list, and are given attributes of Erased ABL Blocks.

[J] After modification of all relevant ABL entries during an ABL fill operation, the block addresses in the ABL replace the block addresses in the ABB list.

[K] Concurrently with an ABL fill operation during a control write, entries for erased blocks in the CBL are moved to the CBB list.

[L] During a MAP exchange operation, all relevant entries are moved from the CBB list to the MAP destination sector.

[M] During a MAP exchange operation, all relevant entries are moved from the CBB list to the MAP source sector.

[N] Subsequent to [L] and [M] during a MAP exchange operation, all remaining

entries are moved from the CBB list to the EBB list.

[O] Subsequent to [N] during a MAP exchange operation, entries other than those moved in [M] are moved from the MAP source sector to fill the EBB list, if possible.

LOGICAL TO PHYSICAL ADDRESS TRANSLATION

[00258] To locate a logical sector's physical location in flash memory, the logical to physical address translation module 140 shown in **FIG. 2** performs a logical to physical address translation. Except for those logical groups that have recently been updated, the bulk of the translation could be performed using the group address table (GAT) residing in the flash memory 200 or the GAT cache in controller RAM 130. Address translations for the recently updated logical groups will require looking up address lists for update blocks which reside mainly in controller RAM 130. The process for logical to physical address translation for a logical sector address is therefore dependent on the type of block associated with the logical group within which the sector is located. The types of blocks are: intact block, sequential data update block, chaotic data update block, closed data update block.

[00259] **FIG. 19** is a flow chart showing the process of logical to physical address translation. Essentially, the corresponding metablock and the physical sector is located by using the logical sector address first to lookup the various update directories such as the open update block list and the close update block list. If the associated metablock is not part of an update process, then directory information is provided by the GAT. The logical to physical address translation includes the following steps:

[00260] **STEP 800:** A logical sector address is given.

[00261] **STEP 810:** Look up given logical address in the open update blocks list 614 (see **FIGs. 15** and **18**) in controller RAM. If lookup fails, proceed to **STEP 820**, otherwise proceed to **STEP 830**.

[00262] **STEP 820:** Look up given logical address in the closed update block list 616. If lookup fails, the given logical address is not part of any update process; proceed to **STEP 870** for GAT address translation. Otherwise proceed to **STEP 860**

for closed update block address translation.

[00263] **STEP 830:** If the update block containing the given logical address is sequential, proceed to **STEP 840** for sequential update block address translation. Otherwise proceed to **STEP 850** for chaotic update block address translation.

[00264] **STEP 840:** Obtain the metablock address using sequential update block address translation. Proceed to **STEP 880**.

[00265] **STEP 850:** Obtain the metablock address using chaotic update block address translation. Proceed to **STEP 880**.

[00266] **STEP 860:** Obtain the metablock address using closed update block address translation. Proceed to **STEP 880**.

[00267] **STEP 870:** Obtain the metablock address using group address table (GAT) translation. Proceed to **STEP 880**.

[00268] **STEP 880:** Convert the Metablock Address to a physical address. The translation method depends on whether the metablock has been relinked.

[00269] **STEP 890:** Physical sector address obtained.

[00270] The various address translation processes are described in more detail as follows:

Sequential Update Block Address Translation (STEP 840)

[00271] Address translation for a target logical sector address in a logical group associated with a sequential update block can be accomplished directly from information in the open update block list 614 (**FIGs. 15 and 18**), as follows.

1. It is determined from the “page tag” and “number of sectors written” fields in the list whether the target logical sector is located in the update block or its associated original block.
2. The metablock address appropriate to the target logical sector is read from the list.

3. The sector address within the metablock is determined from the appropriate “page tag” field.

Chaotic Update Block Address Translation (STEP 850)

[00272] The address translation sequence for a target logical sector address in a logical group associated with a chaotic update block is as follows.

1. If it is determined from the chaotic sector list in RAM that the sector is a recently written sector, address translation may be accomplished directly from its position in this list.

2. The most recently written sector in the CBI block contains, within its chaotic block data field, the physical address of the chaotic update block relevant to the target logical sector address. It also contains, within its indirect sector index field, the offset within the CBI block of the last written CBI sector relating to this chaotic update block (see **FIGs. 16A-16E**).

3. The information in these fields is cached in RAM, eliminating the need to read the sector during subsequent address translation.

4. The CBI sector identified by the indirect sector index field at step 3 is read.

5. The direct sector index field for the most recently accessed chaotic update sub-group is cached in RAM, eliminating the need to perform the read at step 4 for repeated accesses to the same chaotic update block.

6. The direct sector index field read at step 4 or step 5 identifies in turn the CBI sector relating to the logical sub-group containing the target logical sector address.

7. The chaotic block index entry for the target logical sector address is read from the CBI sector identified in step 6.

8. The most recently read chaotic block index field may be cached in controller RAM, eliminating the need to perform the reads at step 4 and step 7 for repeated accesses to the same logical sub-group.

9. The chaotic block index entry defines the location of the target logical sector either in the chaotic update block or in the associated original block. If the valid copy of the target logical sector is in the original block, it is located by use of the original metablock and page tag information.

Closed Update Block Address Translation (STEP 860)

[00273] Address translation for a target logical sector address in a logical group associated with a closed update block can be accomplished directly from information in the closed block update list (see **FIGs. 18**), as follows.

1. The metablock address assigned to the target logical group is read from the list.
2. The sector address within the metablock is determined from the “page tag” field in the list.

GAT Address Translation (STEP 870)

[00274] If a logical group is not referenced by either the open or closed block update lists, its entry in the GAT is valid. The address translation sequence for a target logical sector address in a logical group referenced by the GAT is as follows.

1. The ranges of the available GAT caches in RAM are evaluated to determine if an entry for the target logical group is contained in a GAT cache.
2. If the target logical group is found in step 1, the GAT cache contains full group address information, including both metablock address and page tag, allowing translation of the target logical sector address.
3. If the target address is not in a GAT cache, the GAT index must be read for the target GAT block, to identify the location of the GAT sector relating to the target logical group address.
4. The GAT index for the last accessed GAT block is held in controller RAM, and may be accessed without need to read a sector from flash memory.
5. A list of metablock addresses for every GAT block, and the number of

sectors written in each GAT block, is held in controller RAM. If the required GAT index is not available at step 4, it may therefore be read immediately from flash memory.

6. The GAT sector relating to the target logical group address is read from the sector location in the GAT block defined by the GAT index obtained at step 4 or step 6. A GAT cache is updated with the subdivision of the sector containing the target entry.

7. The target sector address is obtained from the metablock address and “page tag” fields within the target GAT entry.

Metablock to Physical Address Translation (STEP 880)

[00275] If a flag associated with the metablock address indicates that the metablock has been re-linked, the relevant LT sector is read from the BLM block, to determine the erase block address for the target sector address. Otherwise, the erase block address is determined directly from the metablock address.

CONTROL DATA MANAGEMENT

[00276] FIG. 20 illustrates the hierarchy of the operations performed on control data structures in the course of the operation of the memory management. Data Update Management Operations act on the various lists that reside in RAM. Control write operations act on the various control data sectors and dedicated blocks in flash memory and also exchange data with the lists in RAM.

[00277] Data update management operations are performed in RAM on the ABL, the CBL and the chaotic sector list. The ABL is updated when an erased block is allocated as an update block or a control block, or when an update block is closed. The CBL is updated when a control block is erased or when an entry for a closed update block is written to the GAT. The update chaotic sector list is updated when a sector is written to a chaotic update block.

[00278] A control write operation causes information from control data structures in RAM to be written to control data structures in flash memory, with

consequent update of other supporting control data structures in flash memory and RAM, if necessary. It is triggered either when the ABL contains no further entries for erased blocks to be allocated as update blocks, or when the CBI block is rewritten.

[00279] In the preferred embodiment, the ABL fill operation, the CBL empty operation and the EBM sector update operation are performed during every control write operation. When the MAP block containing the EBM sector becomes full, valid EBM and MAP sectors are copied to an allocated erased block, and the previous MAP block is erased.

[00280] One GAT sector is written, and the Closed Update Block List is modified accordingly, during every control write operation. When a GAT block becomes full, a GAT rewrite operation is performed.

[00281] A CBI sector is written, as described earlier, after certain chaotic sector write operations. When the CBI block becomes full, valid CBI sectors are copied to an allocated erased block, and the previous CBI block is erased.

[00282] A MAP exchange operation, as described earlier, is performed when there are no further erased block entries in the EBB list in the EBM sector.

[00283] A MAP Address (MAPA) sector, which records the current address of the MAP block, is written in a dedicated MAPA block on each occasion the MAP block is rewritten. When the MAPA block becomes full, the valid MAPA sector is copied to an allocated erased block, and the previous MAPA block is erased.

[00284] A Boot sector is written in a current Boot block on each occasion the MAPA block is rewritten. When the boot block becomes full, the valid Boot sector is copied from the current version of the Boot block to the backup version, which then becomes the current version. The previous current version is erased and becomes the backup version, and the valid Boot sector is written back to it.

[00285] Example of control data are the directory information and block allocation information associated with the memory block management system, such as those described in connection with **FIG. 20**. As described earlier, the control data is maintained in both high speed RAM and the slower nonvolatile memory blocks. Any frequently changing control data is maintained in RAM with periodic control writes to update equivalent information stored in a nonvolatile metablock. In this way, the control data is stored in nonvolatile, but slower flash memory without the need for frequent access. A hierarchy of control data structures such as GAT, CBI, MAP, and MAPA shown in **FIG. 20** is maintained in flash memory. Thus, a control write operation causes information from control data structures in RAM to update equivalent control data structures in flash memory.

[00286] As described in connection with **FIG. 20**, the block management system maintains a set of control data in flash memory during its operation. This set of control data is stored in the metablocks similar to host data. As such, the control data itself will be block managed and will be subject to updates and therefore garbage collection operations.

[00287] It has also been described that a hierarchy of control data exists, with the ones in the lower hierarchy being updated more often than those higher up. For example, assuming that every control block has N control sectors to write, the following sequence of control updates and control block relocations, normally happens. Referring to **FIG. 20** again, every N CBI updates fill up the CBI block and trigger a CBI relocation (rewrite) and a MAP update. If the Chaotic block gets closed then it also triggers GAT update. Every GAT update triggers a MAP update. Every N GAT updates fill up the block and trigger a GAT block relocation. In addition, when a MAP block gets full it also triggers a MAP block relocation and a MAPA Block (if exist, otherwise BOOT block points directly to MAP) update. In addition, when a MAPA block gets full, it also triggers a MAPA block relocation, a BOOT Block update and a MAP update. In addition, when a BOOT Block gets full, it triggers an active BOOT Block relocation to another BOOT Block.

UPDATE BLOCK REPLACEMENT SCHEME

[00288] According to another aspect of the invention, in a nonvolatile memory with a block management system, an improved block replacement scheme is implemented for a system supporting up to a first predetermined maximum number of update blocks that are concurrently opened for recording data. The update blocks are mainly sequential update blocks where data are recorded in logically sequential order but up to a second predetermined maximum number of which are allowed to be chaotic update blocks where data are not recorded in logically sequential order. Whenever a new allocation of an update block may cause the pool of update blocks to exceed either the first or second predetermined maximum number, one of the existing update blocks in the pool will be closed and removed in order to comply with the limitation. Prior to closing the update block, its data are consolidated into a sequential block. The improved scheme is to avoid the situation where a sequential update can cause an excessive number of chaotic block consolidations. This is accomplished by separating sequential update blocks and chaotic update blocks into respective replacement or consolidation pools. In particular, when a sequential update causes the allocation of a new update block to exceed the first predetermined maximum number, a least recently used sequential update block of the pool is preferentially to make room.

[00289] In the current system, generally there are two types of data: user data and control data. The user data are sent from a host to the memory system typically in logical sequential order. Sequential update blocks are allocated to optimally handle sequential write operations from the host. The user data can also be in logically non-sequential order especially when there are subsequent updates to the logical data. Chaotic update blocks are created to optimally handle the data in non-sequential order. Another source of chaotic or non-sequential data is control data maintained by the file system or memory system such as file and directory information which are generated in the course of storing user data.

[00290] Previous scheme of complying to a practical system limitation of supporting up to a maximum number of concurrently opened update blocks has been to close the least recently used update block in the pool, irrespective of whether it is sequential or chaotic.

[00291] The present scheme improves over the previous scheme where essentially, if during a sequential write operation an update block among a pool thereof needs to be closed to make room for a new allocation, the least recently used sequential update block in the pool is closed. This ensures the various update blocks are effectively used to handle sequential write operations and random write operations. In particular, it avoids the inefficient situation where a large sequential write operation by the host may force a premature closure of a chaotic update block containing FAT and Directory information. Another chaotic block will in effect be created very soon to store the FAT and Directory information which will be updated again once the large sequential write operation is done. The creation of the improved replacement policy mandates separation of the replacement and consolidation pool to prevent the added overhead in consolidating the chaotic block during sequential write and consolidation of potentially an open sequential or an open chaotic block to manage the subsequent FAT and Directory update.

[00292] The practical system limitation of supporting up to a maximum number of concurrently opened update blocks has been described earlier. For example, in one embodiment described in connection with **FIG. 10, STEP 410** tests if a new allocation will exceed the maximum number U_{MAX} of update blocks that can be concurrently opened for accepting update data. If U_{MAX} will be exceeded, the least active among the update blocks, irrespective of whether it is a sequential or chaotic update block, will be closed in **STEP 420** to keep the system within the prescribed limit.

[00293] **FIG. 21** illustrates schematically the two prescribed limits on the number of update blocks for a block managing system. There is an overall limit in the total number of update blocks (U_{MAX}), which is given by the sum of the number of chaotic update blocks N_C and the number of sequential update blocks N_S . Since a chaotic update block is more resource-intensive, requiring additional maintenance of a chaotic block index (CBI), preferably there is also a limit on the maximum number of chaotic update blocks (" U_{CMAX} "). Thus, the first limit requires that the total number of update blocks, $N_C + N_S \leq U_{MAX}$. The second limit requires that the number of chaotic update blocks $N_C \leq U_{CMAX}$.

[00294] FIG. 22 illustrates typical examples of combinations of the two limits optimized for various memory devices. A given combination is designated by U_{MAX} “dash” U_{CMAX} . For example, “3 – 1” designates a block managing system allowing up to a maximum of three update blocks in the update pool and of which only up to one is a chaotic update block. Similarly, “7 – 3” designates a block managing system supporting up to a maximum of seven update blocks and of which up to three can be chaotic update blocks. In general simpler memory systems having smaller memory capacity will be more restrictive, having smaller maximum numbers.

[00295] FIG. 23A, FIG. 23B and FIG. 23C illustrate schematically the sequence of event for introducing a new update block into a pool of update blocks, according to a first situation in a previous replacement scheme.

[00296] FIG. 23A illustrates schematically an update pool with a “5 – 2” configuration as described in FIG. 22. In this example, the update pool is fully populated with a maximum of five allowable update blocks. The update pool is further partitioned into a sequential pool 1200 that contains three sequential update blocks, S1, S2 and S3 and a chaotic pool 1300 that contains a maximum of two chaotic or non-sequential update blocks, C4 and C5. The example shows a first situation in which the least active block happens to be a sequential update block such as S3 1201.

[00297] In the event that a new update block needs to be allocated, one of the existing update blocks in the update pool will need to be closed to make room. For example in the event when the host writes sequential data for a logical group of sectors not serviced by the existing update blocks in the pool, a new update block will need to be allocated for recording the data.

[00298] FIG. 23B illustrates schematically the closing of the least active update block in order to make room for a new update block, according to the previous scheme. The least active update block, in this case happens to be S3 1201 and it will be closed and removed from the pool of the update blocks.

[00299] FIG. 23C illustrates schematically introducing a newly allocated update block into the pool after a closed update block has been removed to make

room. In this case, S6 1212 which is a newly allocated update block will be introduced into the sequential pool 1200 for recording data in logically sequential order. In this way, U_{MAX} , the maximum number of update blocks allowed is not exceeded.

[00300] **FIG. 24A, FIG. 24B** and **FIG. 24C** illustrate schematically the sequence of event for introducing a new update block into a pool of update blocks, according to a second situation in the previous replacement scheme.

[00301] **FIG. 24A** illustrates schematically an update pool with a “5 – 2” configuration as described in **FIG. 22**. In this example, the update pool is fully populated with a maximum of five allowable update blocks. The update pool is further partitioned into a sequential pool 1200 that contains three sequential update blocks, S1, S2 and S3 and a chaotic pool 1300 that contains a maximum of two chaotic or non-sequential update blocks, C4 and C5. The example shows a second situation in which the least active block happens to be a chaotic update block such as C4 1301.

[00302] **FIG. 24B** illustrates schematically the closing of the least active update block in order to make room for a new update block, according to the previous scheme. The least active update block, in this case is C4 1301 and it will be closed and removed from the pool of the update blocks.

[00303] **FIG. 24C** illustrates schematically introducing a newly allocated update block into the pool after a closed update block has been removed to make room. In this case, S6 1212 which is a newly allocated update block will be introduced into the sequential pool 1200 for recording data in logically sequential order. In this way, U_{MAX} , the maximum number of update blocks allowed is not exceeded.

[00304] **FIG. 25A** and **FIG. 25B**, respectively illustrate the maintenance of the U_{MAX} and U_{CMAX} limitation in the scheme previously described in **FIG. 10, FIG. 23B** and **FIG. 24B**. The two limitations are typically imposed at the same time.

[00305] **FIG. 25A** illustrates the scheme previously illustrated in **FIG.10,**

STEP 410 and also in **FIG. 23B** and **FIG. 24B** where the least recently accessed update block is closed whenever a new allocation would exceed a predetermined limit.

[00306] **STEP 1252:** Organizing a nonvolatile memory into blocks, each block for storing data that are erasable together.

[00307] **STEP 1254:** Allocating up to a first predetermined number of update blocks that are concurrently open for storing updates of logical units of data.

[00308] **STEP 1256:** Whenever allocating a new update block would exceed the predetermined number, closing one of the least recently accessed update blocks to make room for the new update block.

[00309] **FIG. 25B** illustrates the scheme previously illustrated in **FIG. 10**, **STEP 370** where the least recently accessed chaotic (non-sequential) update block is closed whenever the number of chaotic update blocks exceeds a predetermined limit.

[00310] **STEP 1354:** Allocating up to a second predetermined number of update blocks among the opened update blocks for storing logical units of data in logically non-sequential order.

[00311] **STEP 1356:** Whenever the number of non-sequential update blocks would exceed the second predetermined number, closing one of the least recently accessed non-sequential update blocks so as not to exceed the second predetermined number.

[00312] One disadvantage of this previous scheme is that it can under some circumstances lead to excessive closure of chaotic update blocks. It is particularly inefficient in the case when a sequential write causes a chaotic block containing control data to be prematurely closed-out to make room in the pool for a newly allocated sequential block. For example, if the closed-out chaotic update block, C4 1310, was recording FAT and Directory information, a replacement will immediately need to be allocated to serve that function as soon as the sequential write is done. This would entail yet another round of closure of the least active update block currently in the pool in order to make room for the replacement update block for

recording control data.

[00313] According to the current aspect of the invention, the closure of an update block so as not to exceed a predetermined maximum number of update blocks is further refined from merely selecting the least recently used update blocks to a scheme that reduces the instance of excessive closure of chaotic update blocks. In a preferred embodiment, if an update block is allocated to record sequential data and one in a pool of update blocks needs to be closed to make room, the least active sequential update block in the pool is closed.

[00314] FIG. 26A, FIG. 26B and FIG. 26C illustrate schematically the sequence of event for introducing a new update block into a pool of update blocks, according to the present improved replacement scheme.

[00315] FIG. 26A illustrates schematically an update pool with a "5 - 2" configuration as described in FIG. 22. In this example, the update pool is fully populated with a maximum of five allowable update blocks. The update pool is further partitioned into a sequential pool 1200 that contains three sequential update blocks, S1, S2 and S3 and a chaotic pool 1300 that contains a maximum of two chaotic or non-sequential update blocks, C4 and C5. The example shows a similar situation as in FIG. 24A in which the least active block happens to be a chaotic update block such as C4 1301. Furthermore it shows the sequential block S3 1202 as being the least active in the sequential pool 1200.

[00316] FIG. 26B illustrates schematically the closing of one among the pool of an update blocks in order to make room for a new update block, according to the present improved scheme. In the event of a new allocation associated with a sequential update and the number of update blocks in the pool is already at a maximum, one of the update blocks in the pool will have to be closed out to make room for the newly allocated update block. However, in this case the least active block is C4 1301, and because it is a chaotic block, it will be passed over. Instead, the least active update block in the sequential pool 1200 will be closed out. In this example, it is S3 1202 that will be closed out and removed from the update block pool.

[00317] **FIG. 26C** illustrates schematically introducing a newly allocated update block into the pool after a closed update block has been removed to make room. In this case, S6 1212 which is a newly allocated update block will be introduced into the sequential pool 1200 for recording data in logically sequential order. In this way, U_{MAX} , the maximum number of update blocks allowed is not exceeded.

[00318] **FIG. 27A, FIG. 27B** and **FIG. 27C** illustrate schematically the sequence of event for introducing a new chaotic update block into a pool of update blocks, according to the present improved replacement scheme.

[00319] **FIG. 27A** illustrates schematically an update pool with a “5 – 2” configuration as described in **FIG. 22**. In this example, the update pool is fully populated with a maximum of five allowable update blocks. The update pool is further partitioned into a sequential pool 1200 that contains three sequential update blocks, S1, S2 and S3 and a chaotic pool 1300 that contains a maximum of two chaotic or non-sequential update blocks, C4 and C5. The example shows that the least active block happens to be a sequential update block such as S6 1201. Furthermore it shows the chaotic block C4 1302 as being the least active in the chaotic pool 1300.

[00320] **FIG. 27B** illustrates schematically the closing of one among the pool of an update blocks in order to make room for a new update block, according to the present improved scheme. In the event of a new chaotic update block being introduced into an already full chaotic pool 1300, one of the update blocks in the chaotic pool will have to be closed out to make room. In the example, the chaotic pool 1300 already contains a maximum of two chaotic update blocks. When, another chaotic update block is created, as for example when an existing sequential update block S1 1220 has be converted to a chaotic block, the maximum number of chaotic blocks would be exceeded unless one of them is removed. In this case the least active chaotic block C4 1302 is closed out and removed from the chaotic pool 1300 to make room.

[00321] **FIG. 27C** illustrates schematically introducing a new chaotic update block into the pool after another chaotic update block has been closed and removed to

make room. In this case, S1 which has been converted from a sequential update block 1220 in the sequential pool 1200 to a chaotic update block C6 1320 in the chaotic pool 1300. In this way, U_{CMAX} , the maximum number of chaotic update blocks allowed is not exceeded.

[00322] **FIG. 28** is a flow chart illustrating the present improved scheme of managing a limited set of update blocks during a sequential update, according to a first embodiment.

[00323] **STEP 1400:** Organizing a nonvolatile memory into blocks, each block for storing data that are erasable together.

[00324] **STEP 1402:** Allocating up to a first predetermined number of update blocks that are concurrently open for storing updates of logical units of data

[00325] **STEP 1406:** Responsive to a write command to write sequential data, writing logical units of data in sequential order onto an update block.

[00326] **STEP 1408:** Responsive to a predetermined condition being satisfied for the update block to be closed to further writing of the sequential logical units of data, allocating a new update block to continue the writing, and if the new allocation would exceed the first predetermined number, preferentially closing a least recently accessed update block in sequential order over any least recently accessed update block in non-sequential order.

[00327] **FIG. 29** is a flow chart illustrating the present improved scheme of managing a limited set of update blocks having two predetermined limits, according to a second embodiment.

[00328] **STEP 1410:** Organizing a nonvolatile memory into blocks, each block for storing data that are erasable together.

[00329] **STEP 1412:** Allocating up to a first predetermined number of update blocks that are concurrently open for storing updates of logical units of data.

[00330] **STEP 1416:** Allocating up to a second predetermined number of update blocks among the opened update blocks for storing logical units of data in

logically non-sequential order.

[00331] **STEP 1418:** Whenever introduction of an update block for storing data in logically sequential order may exceed the first predetermined number, closing a least recently accessed update blocks containing data in logically sequential order to make room for the introduced update block.

[00332] **STEP 1420:** Whenever introduction of an update block for storing data in logically non-sequential order may exceed the second predetermined number, closing a least recently accessed update blocks containing data in logically non-sequential order to make room for the introduced update block.

[00333] A generalization of the present scheme is to classify the update blocks based on a set of attributes, such as whether the update block is storing sequential or non-sequential data or whether it is storing some predefined type of system data. In implementing a pool of update blocks of limited number, each class of update blocks will have its own rule for replacement when the maximum number supported for that class will be exceeded.

[00334] For example, sequential update block and non-sequential update blocks are two different classes. The replacement rules for each of these classes are the same, namely to replace a least active one with a new one. Thus, when the pool of sequential update blocks will be exceeded, a least active one in the pool will be closed and removed before a new one is introduced to the pool. Similarly for the pool of non-sequential update blocks.

[00335] In general each class has its own replacement rule independent of the other classes. Examples of replacement rules are to replace the least recently accessed, the most recently accessed, the least frequently accessed, the most frequently accessed, etc, depending on the corresponding classes.

[00336] **FIG. 30** is a flow chart illustrating the present improved scheme of managing a limited set of update blocks having class-based replacement rules.

[00337] **STEP 1430:** Organizing a nonvolatile memory into blocks, each block for storing data that are erasable together.

[00338] **STEP 1432:** Providing a pool of up to a first predetermined maximum number of update blocks that are concurrently open for storing updates of logical units of data.

[00339] **STEP 1436:** Providing a set of predefined classes for classifying update blocks based on a set of attributes, with each class supporting a sub-pool of up to an associated predetermined maximum number of update blocks;

[00340] **STEP 1438:** Providing a set of corresponding replacement rules to the set of predefined classes to specify the update block in the respective sub-pools to be replaced.

[00341] **STEP 1438:** Grouping the update blocks in the pool by class into corresponding sub-pools.

[00342] **STEP 1440:** Closing and removing a least active update block in a sub-pool containing the associated predetermined maximum number of update blocks whenever another update block of the same class is being introduced thereto.

[00343] All patents, patent applications, articles, books, specifications, other publications, documents and things referenced herein are hereby incorporated herein by this reference in their entirety for all purposes. To the extent of any inconsistency or conflict in the definition or use of a term between any of the incorporated publications, documents or things and the text of the present document, the definition or use of the term in the present document shall prevail.

[00344] Although the various aspects of the present invention have been described with respect to certain embodiments, it is understood that the invention is entitled to protection within the full scope of the appended claims.

IT IS CLAIMED:

1. In a non-volatile memory organized into a plurality of blocks, each block for storing logical units of data that are erasable together, a method of storing data in the memory, comprising:

allocating up to a first predetermined number of blocks as update blocks that are concurrently open for storing updates of logical units of data;

responsive to a write command to write data that is in logically sequential order, writing data onto an update block in logically sequential order; and

responsive to a predetermined condition being satisfied for the update block to be closed to further writing of the sequential logical units of data, allocating a new update block to continue the writing, and if the new allocation would exceed the first predetermined number, preferentially closing a least recently accessed update block storing data in sequential order over any least recently accessed one storing data in non-sequential order.

2. In a non-volatile memory organized into a plurality of blocks, each block for storing logical units of data that are erasable together, a method of storing data in the memory, comprising:

allocating up to a first predetermined number of blocks as update blocks that are concurrently open for storing updates of logical units of data;

allocating up to a second predetermined number of update blocks among the opened update blocks for storing logical units of data in logically non-sequential order;

whenever introduction of an update block for storing data in logically sequential order may exceed the first predetermined number, closing a least recently accessed update blocks containing data in logically sequential order to make room for the introduced update block; and

whenever introduction of an update block for storing data in logically non-sequential order may exceed the second predetermined number, closing a least recently accessed update blocks containing data in logically non-sequential order to make room for the introduced update block.

3. The method as in claim 2, wherein the update block for storing data in logically non-sequential order was converted from one storing data in logically sequential order.

4. In a non-volatile memory organized into a plurality of blocks, each block for storing logical units of data that are erasable together, a method of storing data in the memory, comprising:

allocating up to a first predetermined number of blocks as update blocks that are concurrently open for storing updates of logical units of data;

providing a pool of up to a first predetermined maximum number of update blocks that are concurrently open for storing updates of logical units of data;

providing a set of predefined classes for classifying update blocks based on a set of attributes, with each class supporting a sub-pool of up to an associated predetermined maximum number of update blocks;

providing a set of corresponding replacement rules to the set of predefined classes to specify the update block in the respective sub-pools to be replaced;

grouping the update blocks in the pool by class into corresponding sub-pools; and

closing and removing a least active update block in a sub-pool containing the associated predetermined maximum number of update blocks whenever another update block of the same class is being introduced thereto, the removed update block being selected in accordance with the corresponding replacement rule for the same class.

5. The method as in claim 4, wherein the set attributes includes a block storing data in logically sequential order.

6. The method as in claim 4, wherein the set attributes includes a block storing data in logically non-sequential order.

7. The method as in claim 4, wherein the set attributes includes a block storing system data associated with operating the memory.

8. The method as in claim 4, wherein the memory is a flash EEPROM.
9. The method as in claim 4, wherein the memory has a NAND structure.
10. The method as in claim 4, wherein the memory is on a removable memory card.
11. The method as in claim 4, wherein the non-volatile memory has memory cells with a floating gate structure.
12. The method as in claim 4, wherein the non-volatile memory has memory cells with a dielectric layer structure.
13. The method as in any one of claims 1-12, wherein the memory has memory cells that each stores one bit of data.
14. The method as in any one of claims 1-12, wherein the memory has memory cells that each stores more than one bit of data.
15. A nonvolatile memory, comprising:
 - a memory organized into blocks, each block partitioned in memory units that are erasable together, each memory unit for storing a logical unit of data; and
 - a controller for controlling operations of the said blocks including:
 - allocating up to a first predetermined number of blocks as update blocks that are concurrently open for storing updates of logical units of data;
 - responsive to a write command to write data that is in logically sequential order, writing data onto an update block in logically sequential order; and
 - responsive to a predetermined condition being satisfied for the update block to be closed to further writing of the sequential logical units of data, allocating a new update block to continue the writing, and if the new allocation would exceed the first predetermined number, preferentially closing a least recently accessed update block storing data in sequential order over any least recently accessed one storing data in non-sequential order.

16. A nonvolatile memory, comprising:
a memory organized into blocks, each block partitioned in memory units that are erasable together, each memory unit for storing a logical unit of data; and
a controller for controlling operations of the said blocks including:
allocating up to a first predetermined number of blocks as update blocks that are concurrently open for storing updates of logical units of data;
allocating up to a second predetermined number of update blocks among the opened update blocks for storing logical units of data in logically non-sequential order;
whenever introduction of an update block for storing data in logically sequential order may exceed the first predetermined number, closing a least recently accessed update blocks containing data in logically sequential order to make room for the introduced update block; and
whenever introduction of an update block for storing data in logically non-sequential order may exceed the second predetermined number, closing a least recently accessed update blocks containing data in logically non-sequential order to make room for the introduced update block.

17. The memory as in claim 16, wherein the update block for storing data in logically non-sequential order was converted from one storing data in logically sequential order.

18. A nonvolatile memory, comprising:
a memory organized into blocks, each block partitioned in memory units that are erasable together, each memory unit for storing a logical unit of data;
a pool of up to a first predetermined maximum number of update blocks that are concurrently open for storing updates of logical units of data;
a set of predefined classes for classifying update blocks based on a set of attributes, with each class supporting a sub-pool of up to an associated predetermined maximum number of update blocks;
a set of corresponding replacement rules to the set of predefined classes to specify the update block in the respective sub-pools to be replaced;

a set of subpools containing update blocks by class; and
a controller for controlling operations of the blocks including:
closing and removing an update block in a sub-pool containing the
associated predetermined maximum number of update blocks whenever another
update block of the same class is being introduced thereto, the removed update block
being selected in accordance with the corresponding replacement rule for the same
class.

19. The memory as in claim 18, wherein the set attributes includes a block
storing data in logically sequential order.

20. The memory as in claim 18, wherein the set attributes includes a block
storing data in logically non-sequential order.

21. The memory as in claim 18, wherein the set attributes includes a block
storing system data associated with operating the memory.

22. The memory as in claim 18, wherein the memory is a flash EEPROM.

23. The memory as in claim 18, wherein the memory has a NAND
structure.

24. The memory as in claim 18, wherein the memory is on a removable
memory card.

25. The memory as in claim 18, wherein the non-volatile memory has
memory cells with a floating gate structure.

26. The memory as in claim 18, wherein the non-volatile memory has
memory cells with a dielectric layer structure.

27. A nonvolatile memory, comprising:

a memory organized into blocks, each block partitioned in memory units that are erasable together, each memory unit for storing a logical unit of data;

a pool of up to a first predetermined maximum number of update blocks that are concurrently open for storing updates of logical units of data;

a set of predefined classes for classifying update blocks based on a set of attributes, with each class supporting a sub-pool of up to an associated predetermined maximum number of update blocks;

a set of corresponding replacement rules to the set of predefined classes to specify the update block in the respective sub-pools to be replaced;

a set of subpools containing update blocks by class; and

means for closing and removing an update block in a sub-pool containing the associated predetermined maximum number of update blocks whenever another update block of the same class is being introduced thereto, the removed update block being selected in accordance with the corresponding replacement rule for the same class

28. The memory as in any one of claims 15-27, wherein the memory has memory cells that each stores one bit of data.

29. The memory as in any one of claims 15-27, wherein the memory has memory cells that each stores more than one bit of data.


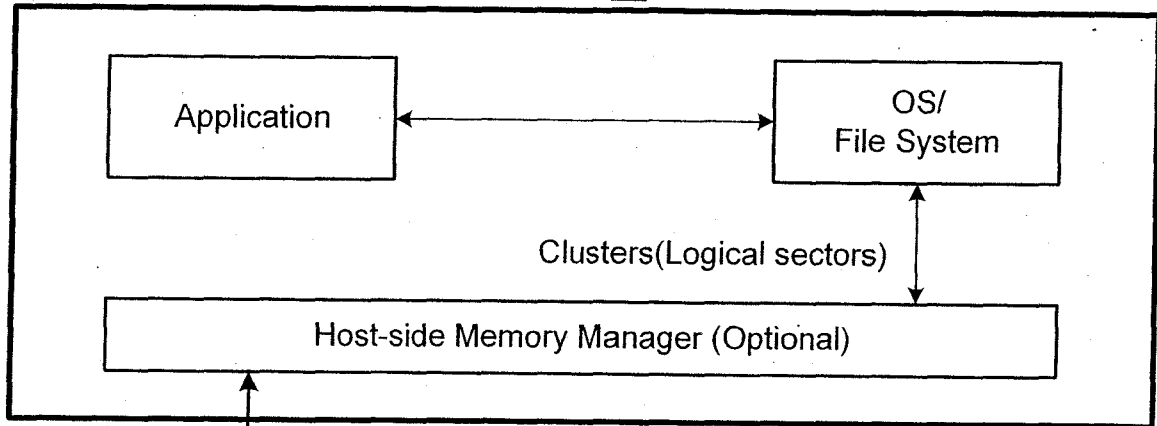


FIG. 1

2/38
HOST 10



MEMORY SYSTEM 20

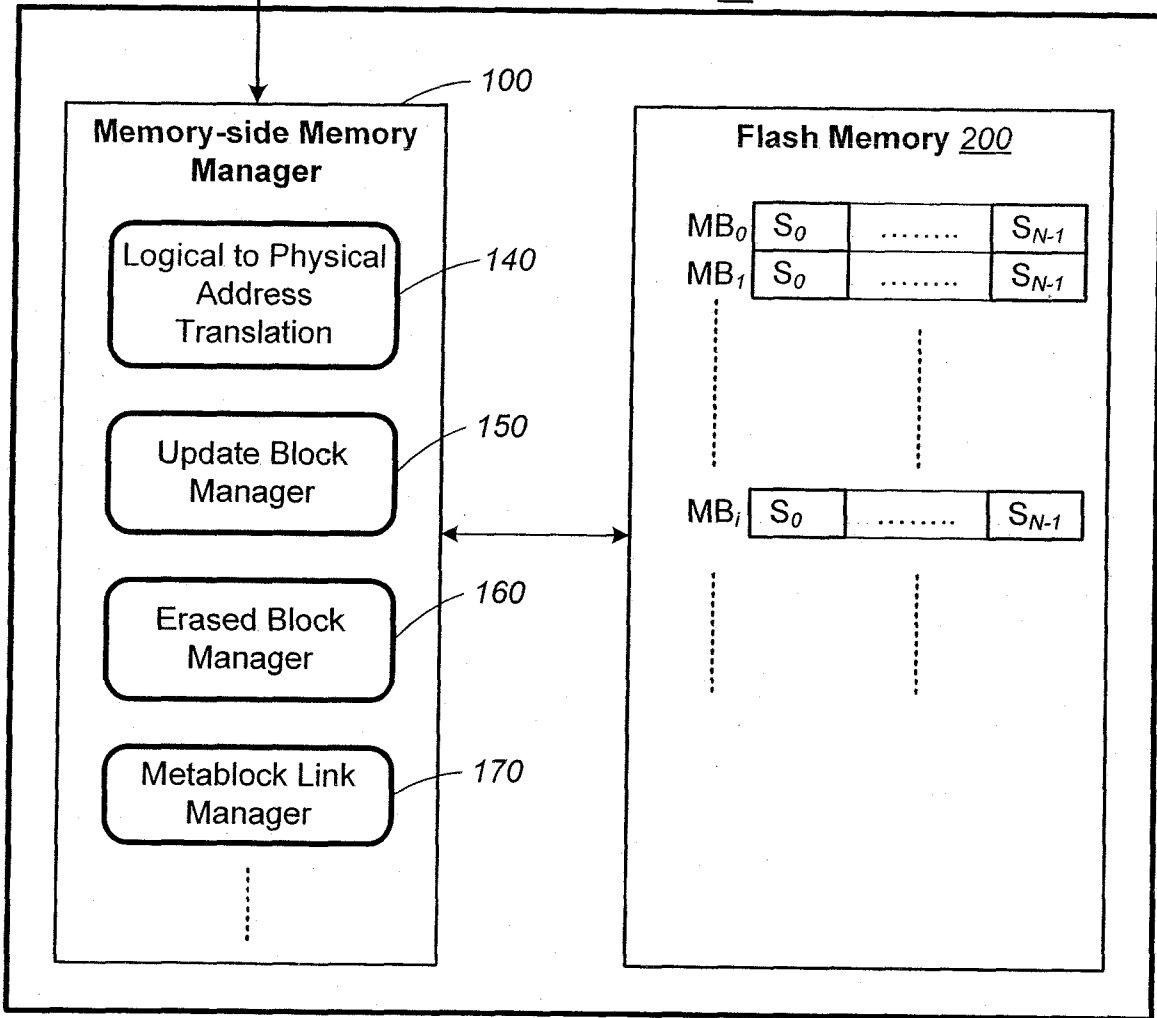


FIG. 2

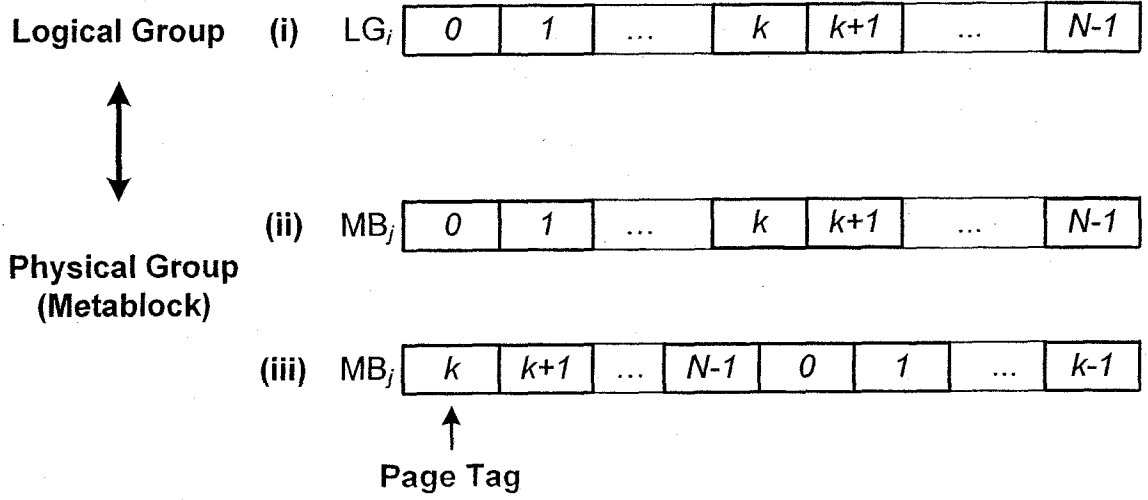


FIG. 3A




FIG. 3B

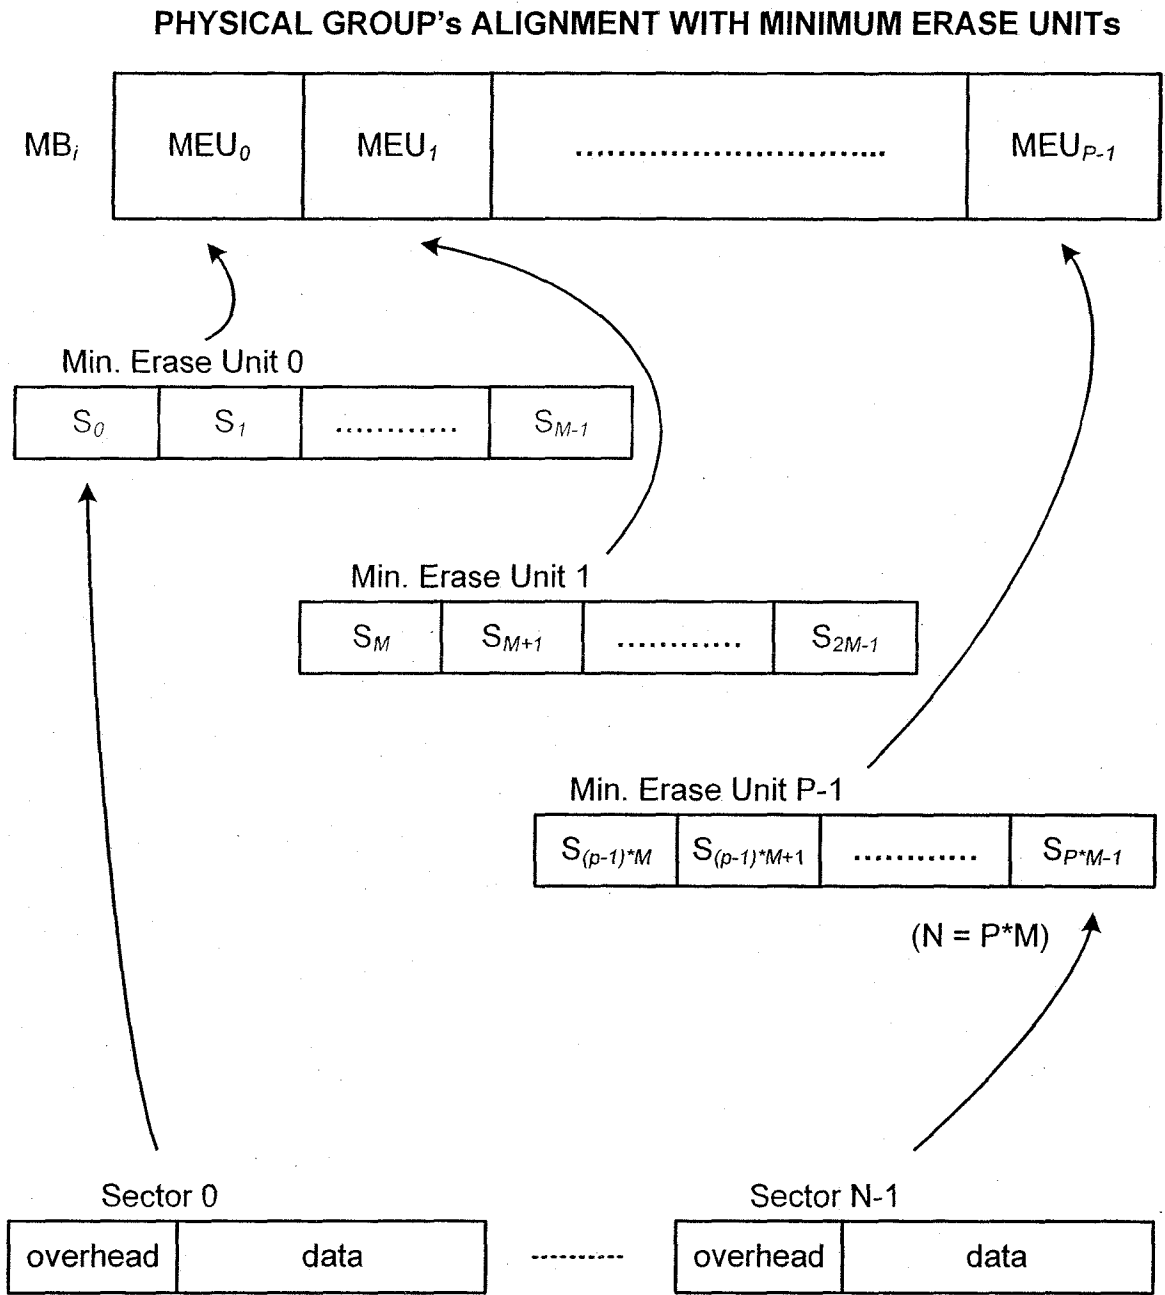


FIG. 4




FIG. 5A



FIG. 5B




FIG. 5C

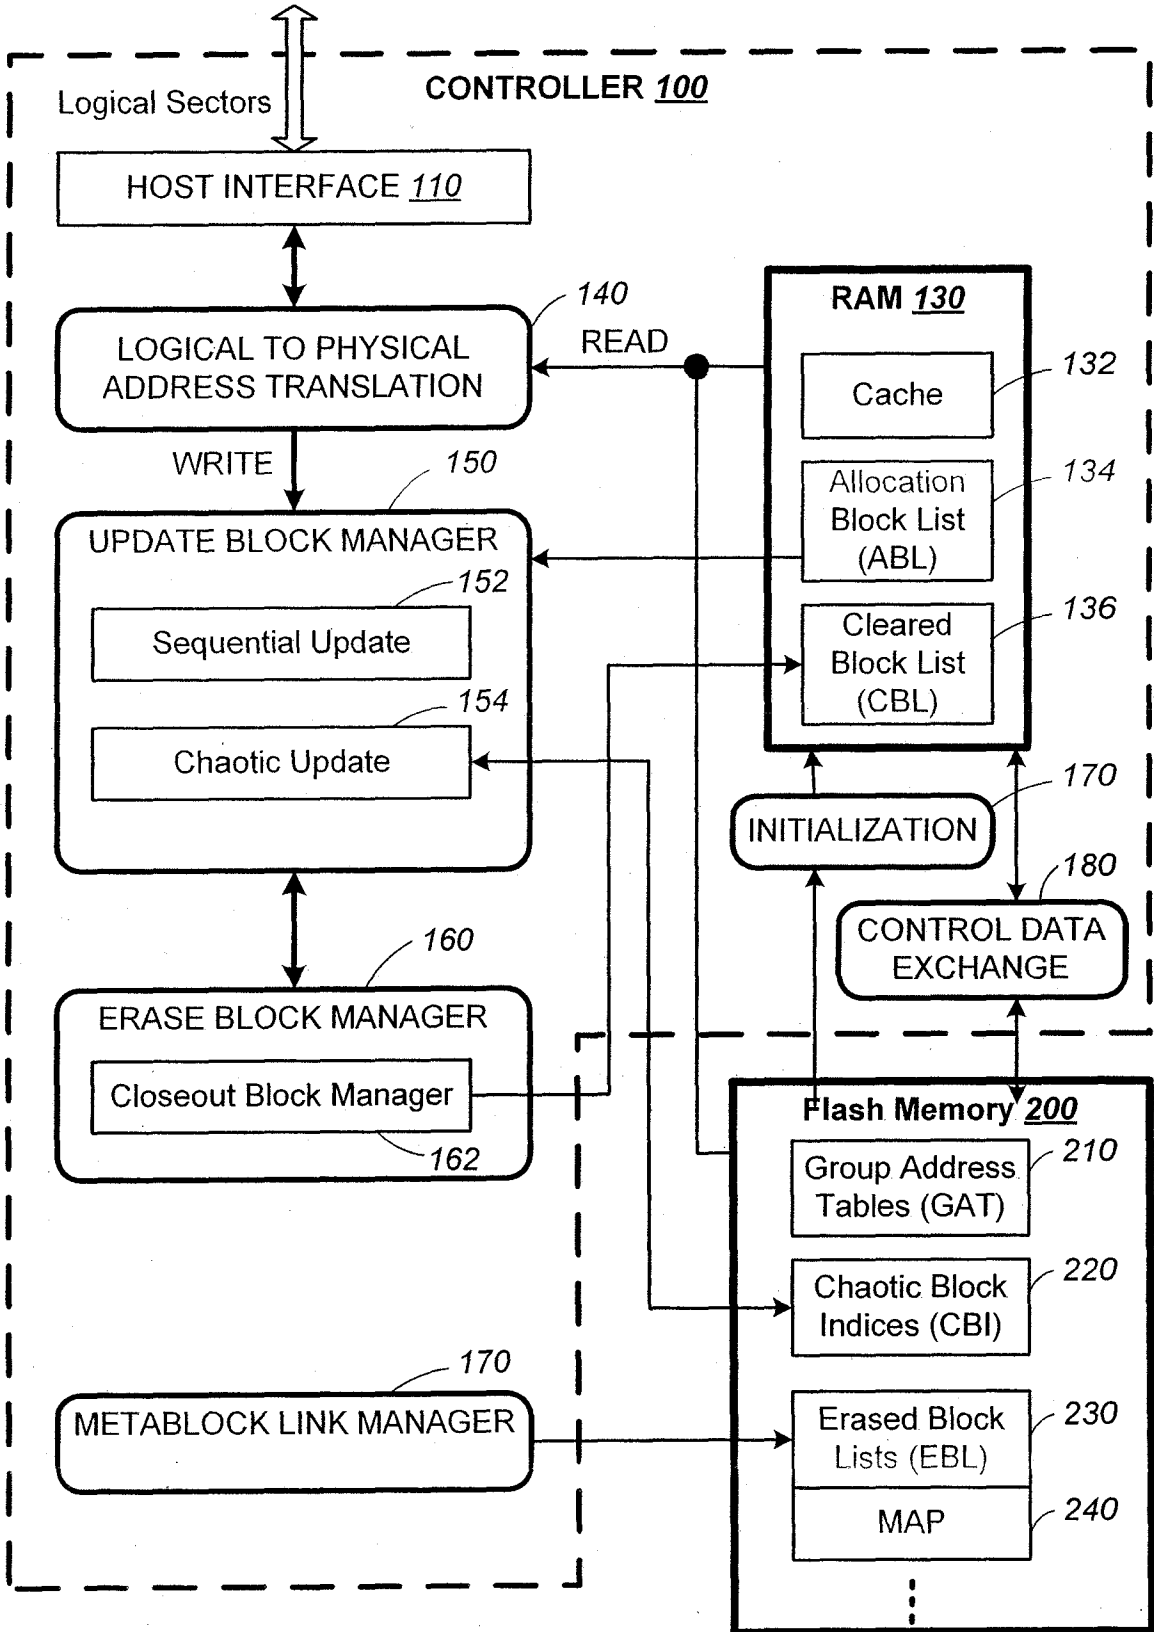
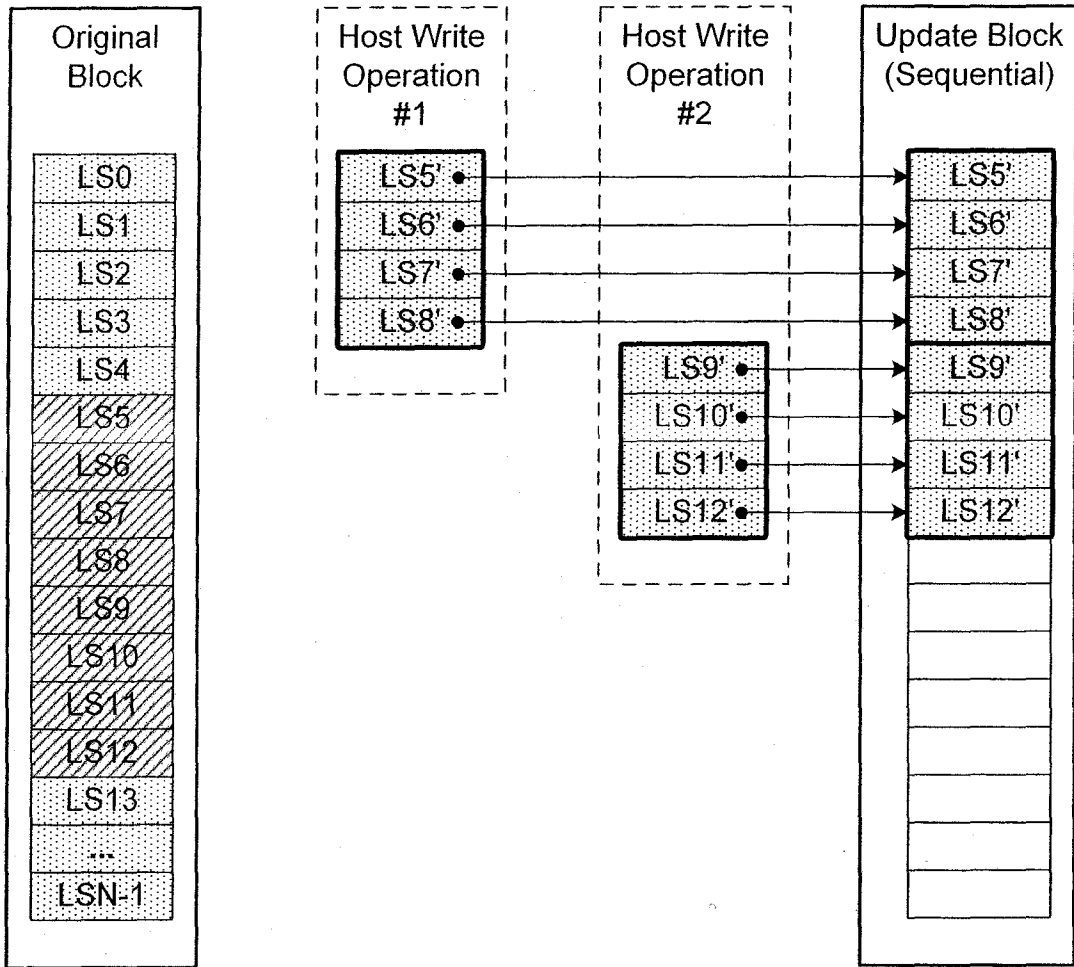

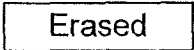



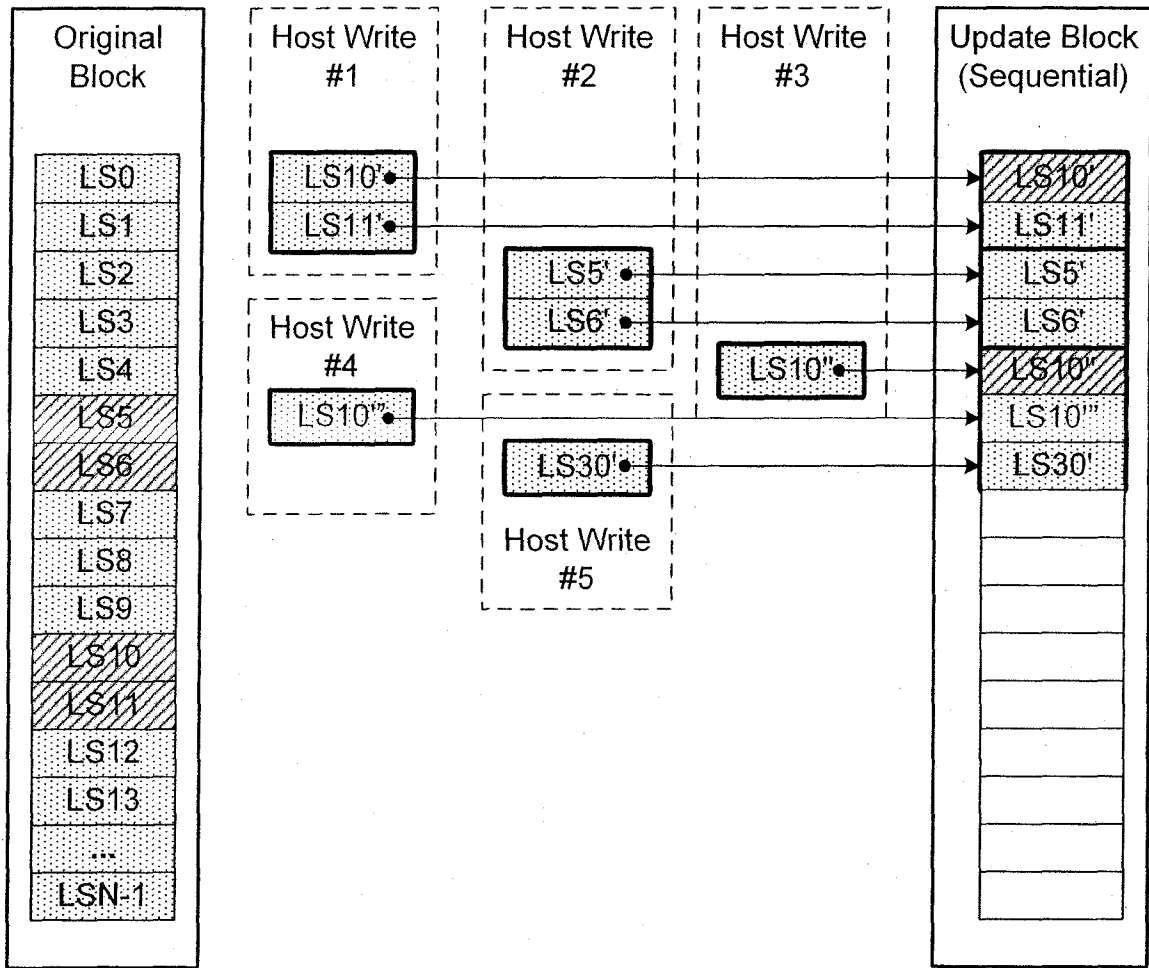
FIG. 6


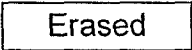



Key:  Obsolete  Erased  Valid

SEQUENTIAL UPDATE
Example

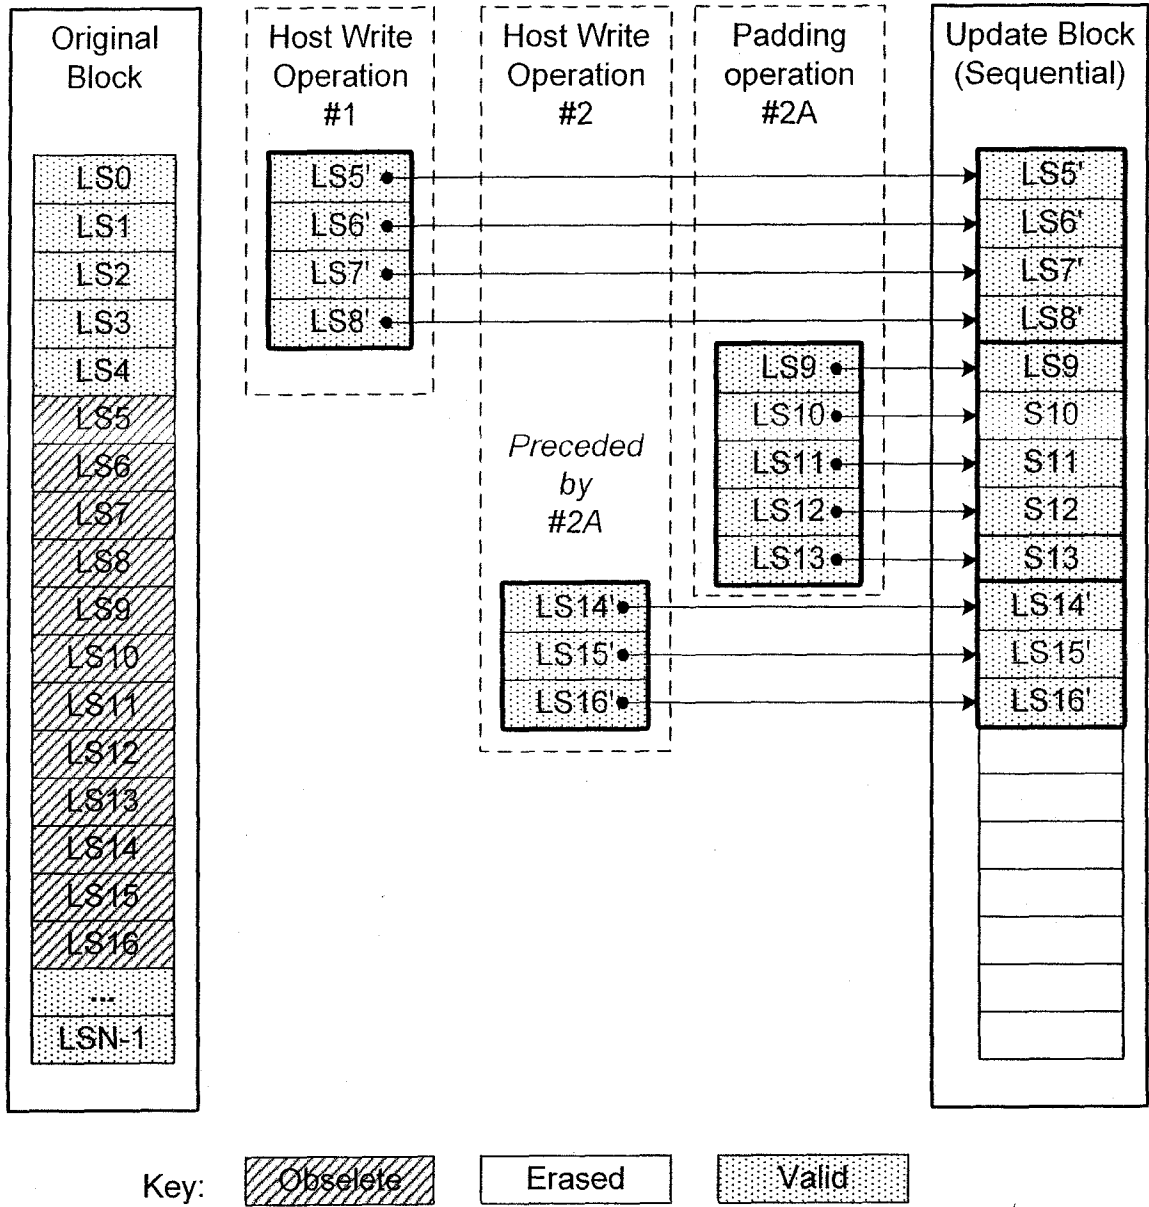
FIG. 7A



Key:  Obsolete  Erased  Valid

**CHAOTIC UPDATE
(NONSEQUENTIAL)
Example**

FIG. 7B



**FORCED SEQUENTIAL
UPDATE
Example**

FIG. 8

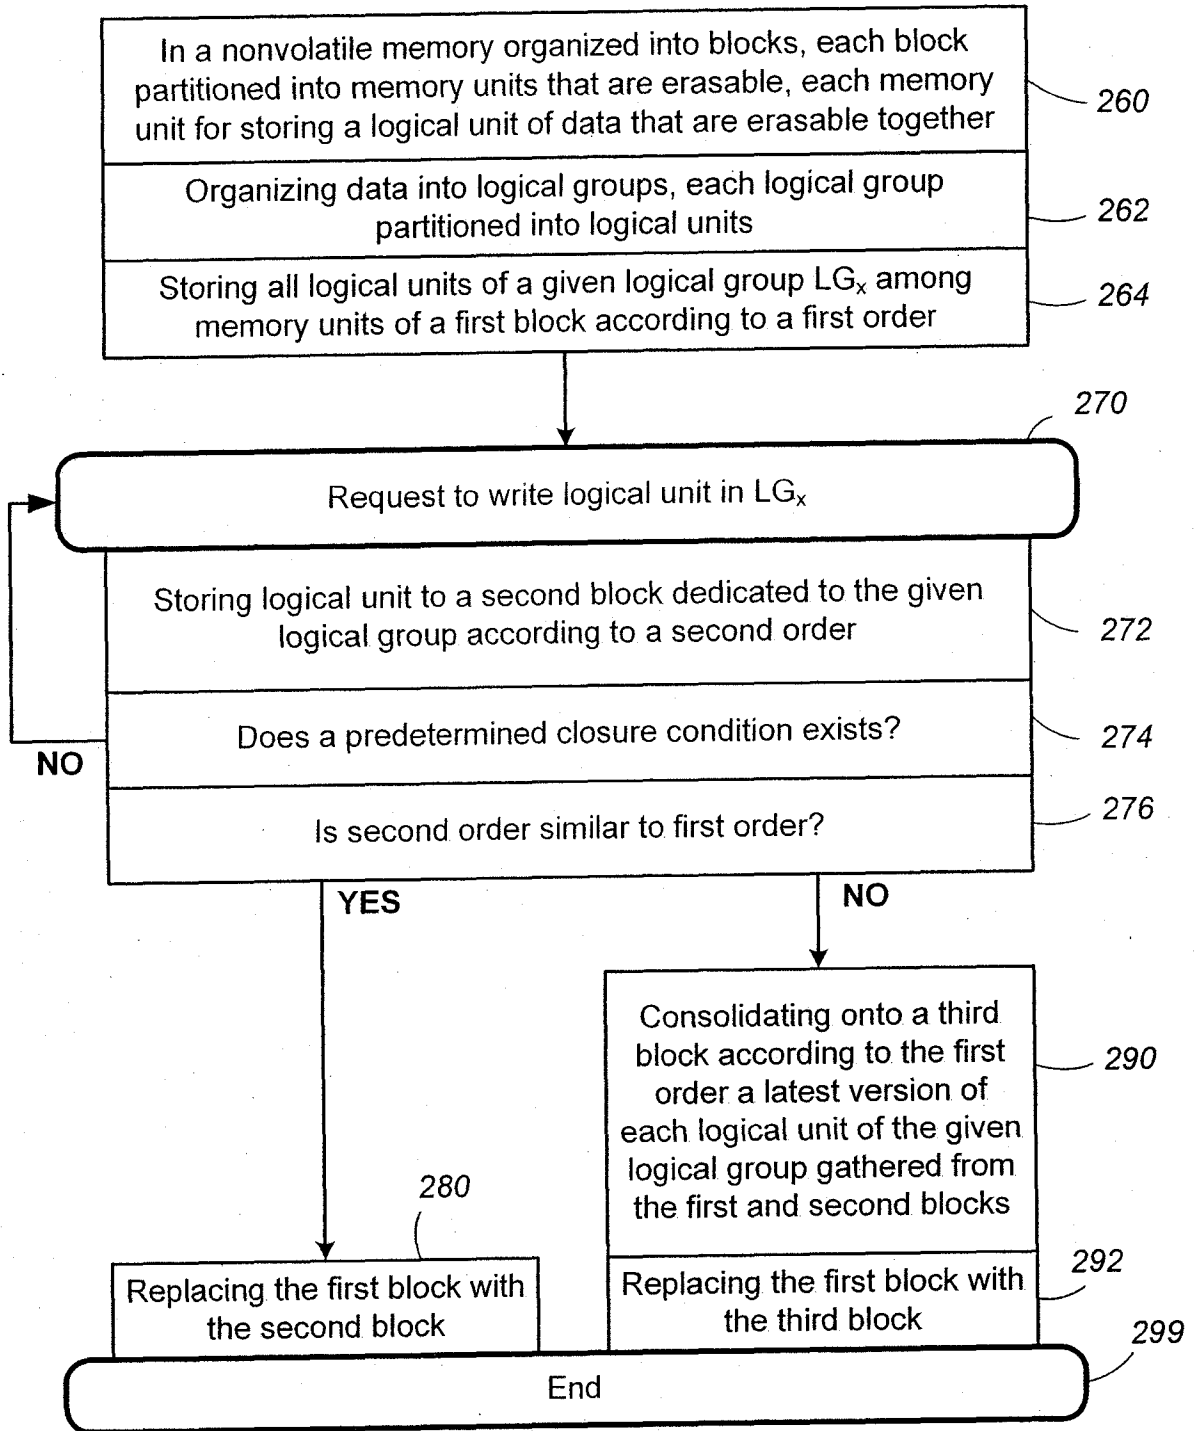


FIG. 9

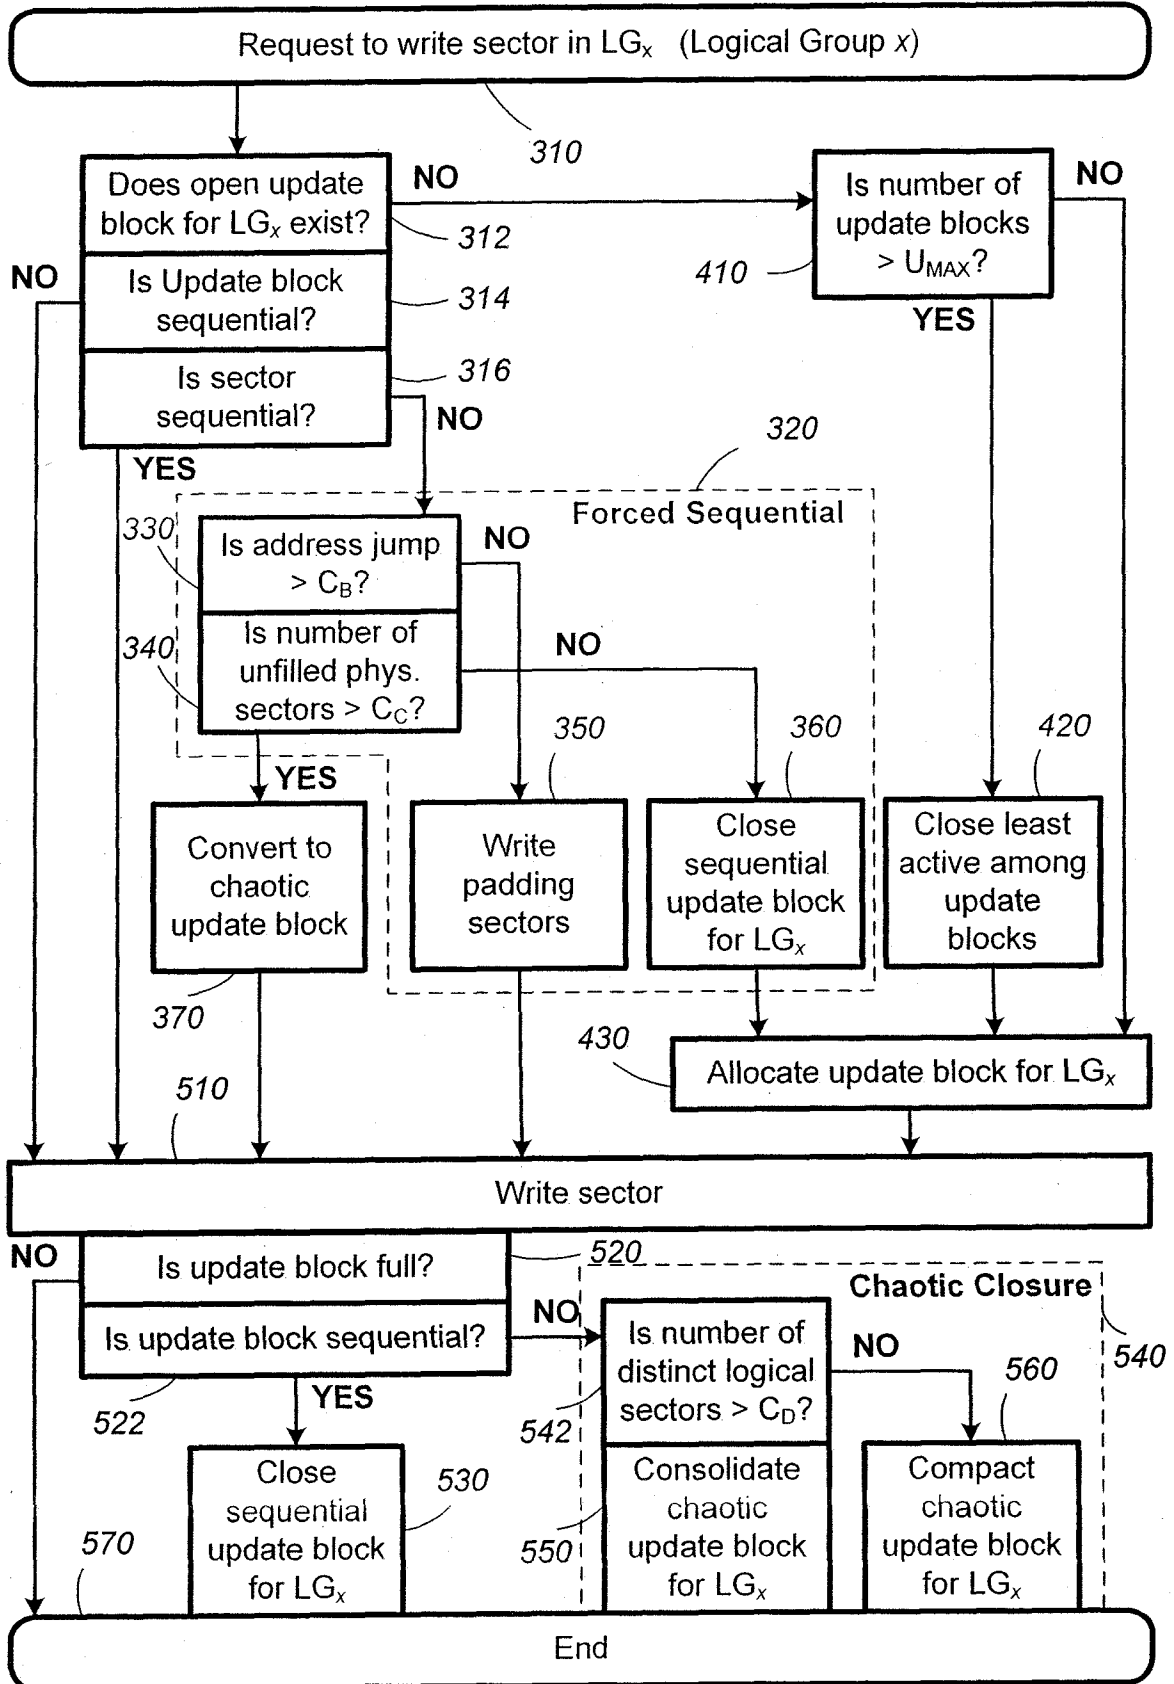


FIG. 10 (Prior Art)

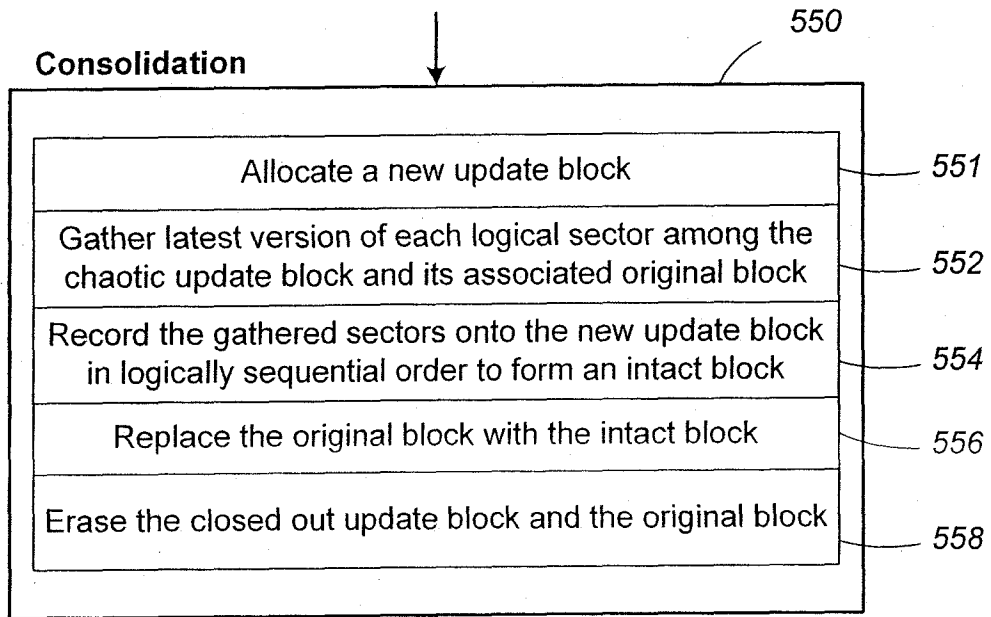


FIG. 11A

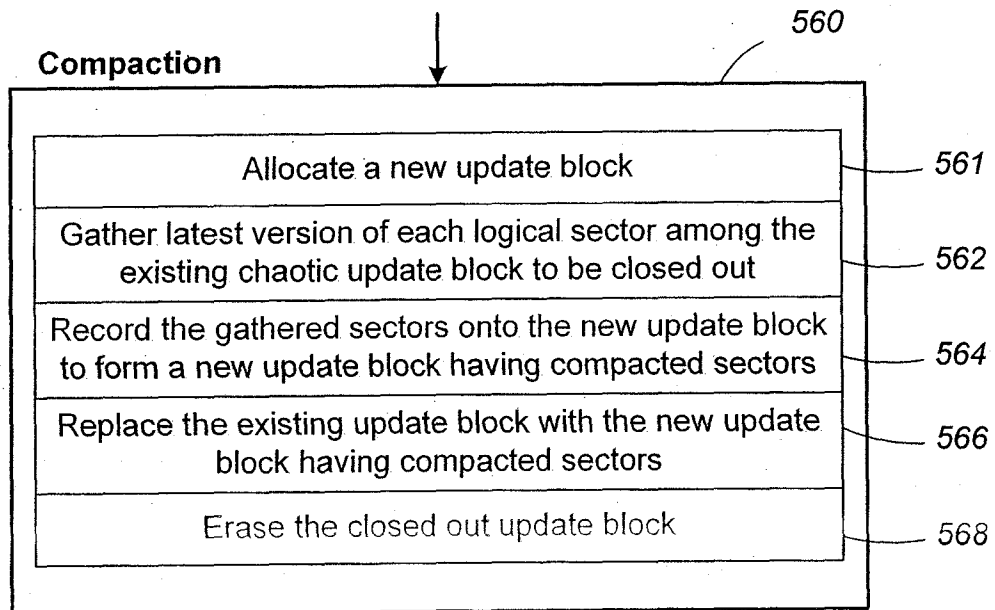
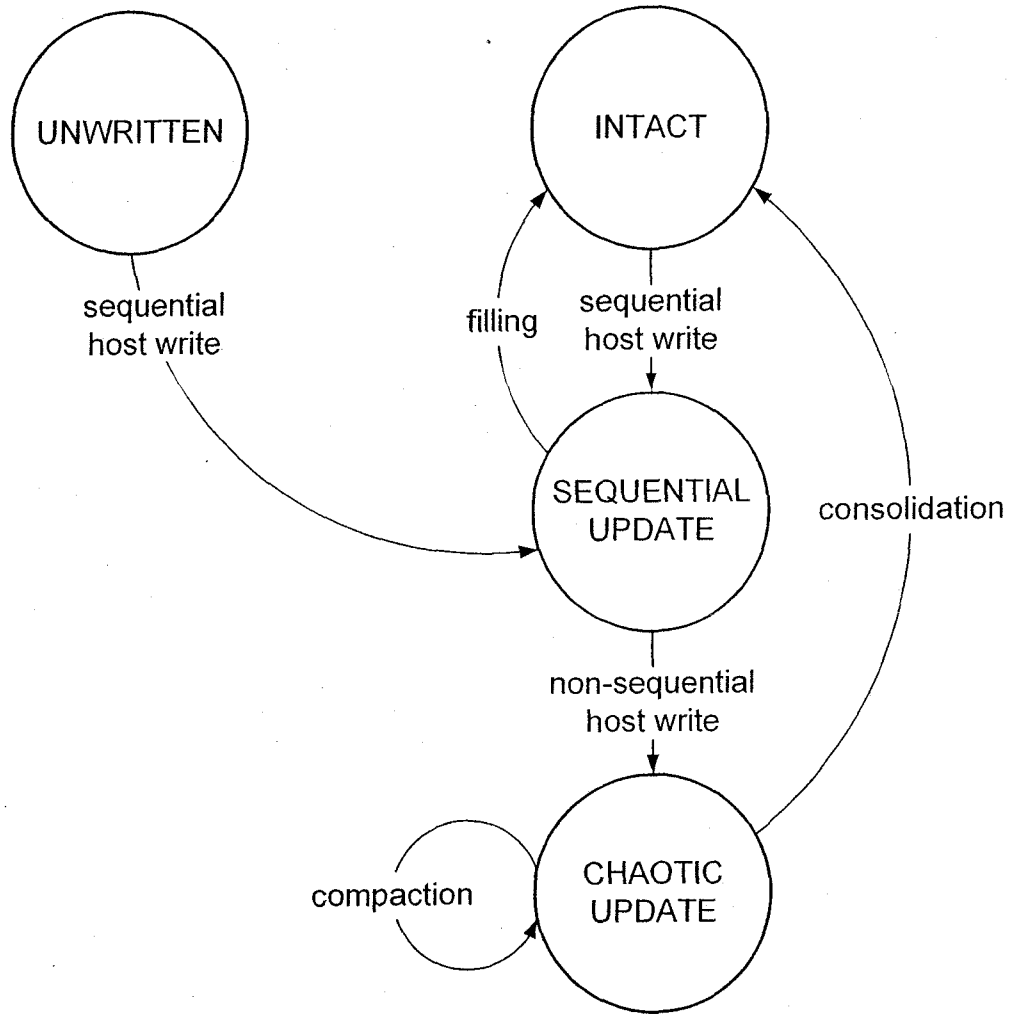


FIG. 11B

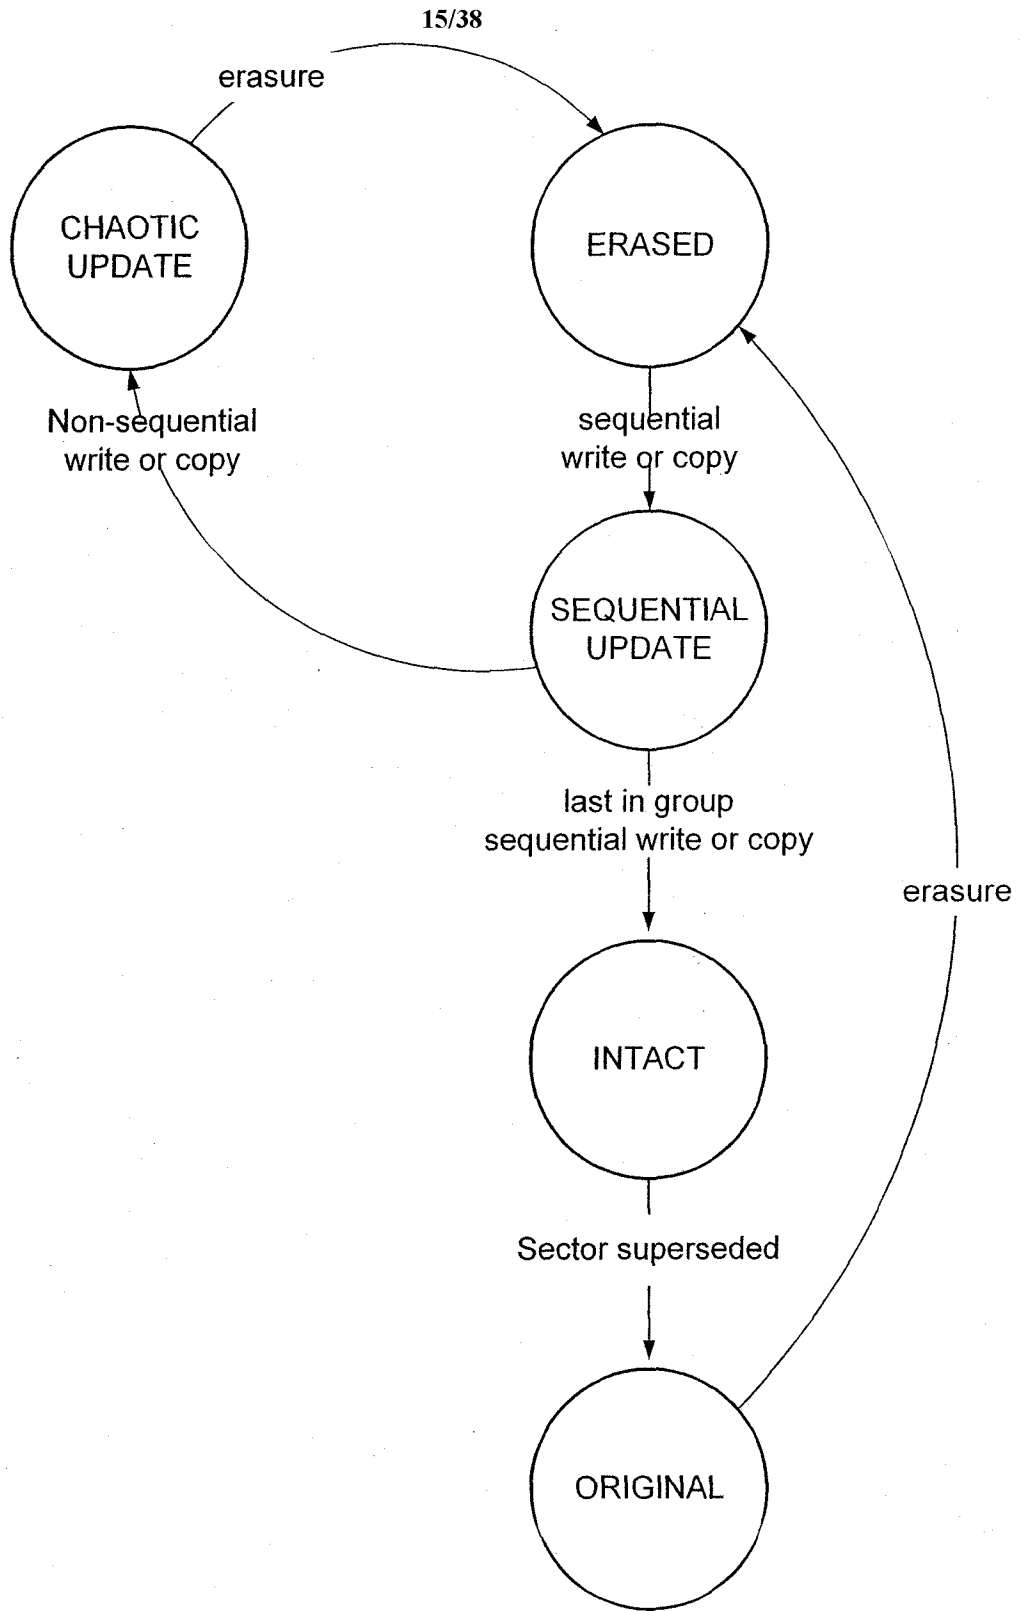


LOGICAL GROUP STATES

FIG. 12A

LOGICAL GROUP STATE	DESCRIPTION
Intact	Complete Logical Group in logically sequential order in a single metablock
Unwritten	No portion of the Logical Group has ever been written
Sequential Update	An update segment of the Logical Group is written to an update block without changing its existing sequential nature
Chaotic Update	An update segment of the Logical Group is written to an update block in logically non-sequential order

FIG. 12B

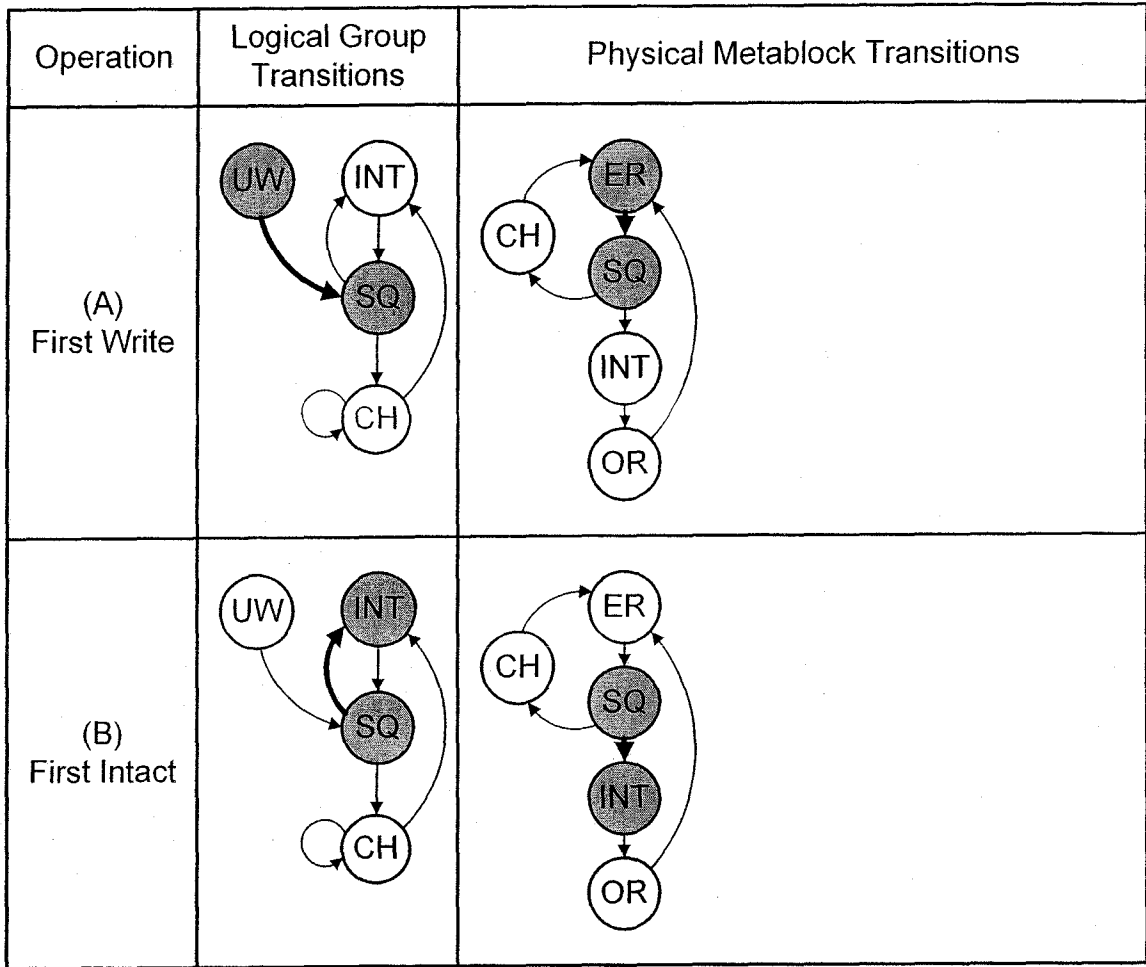


PHYSICAL METABLOCK STATES


FIG. 13A

METABLOCK STATE	DESCRIPTION
Erased	All the sectors in the metablock are erased
Sequential Update	The metablock is partially written with sectors in logically sequential order, possibly using page tag. All the sectors belong to the same logical group.
Non-sequential (Chaotic) Update	The metablock is partially or fully written with sectors in logically non-sequential order. Any sector can be written more than once. All sectors belong to the same logical group
Intact	The metablock is fully written in logically sequential order, possibly using page tag
Original	The metablock was previously Intact but at least one sector has been made obsolete by a host data update

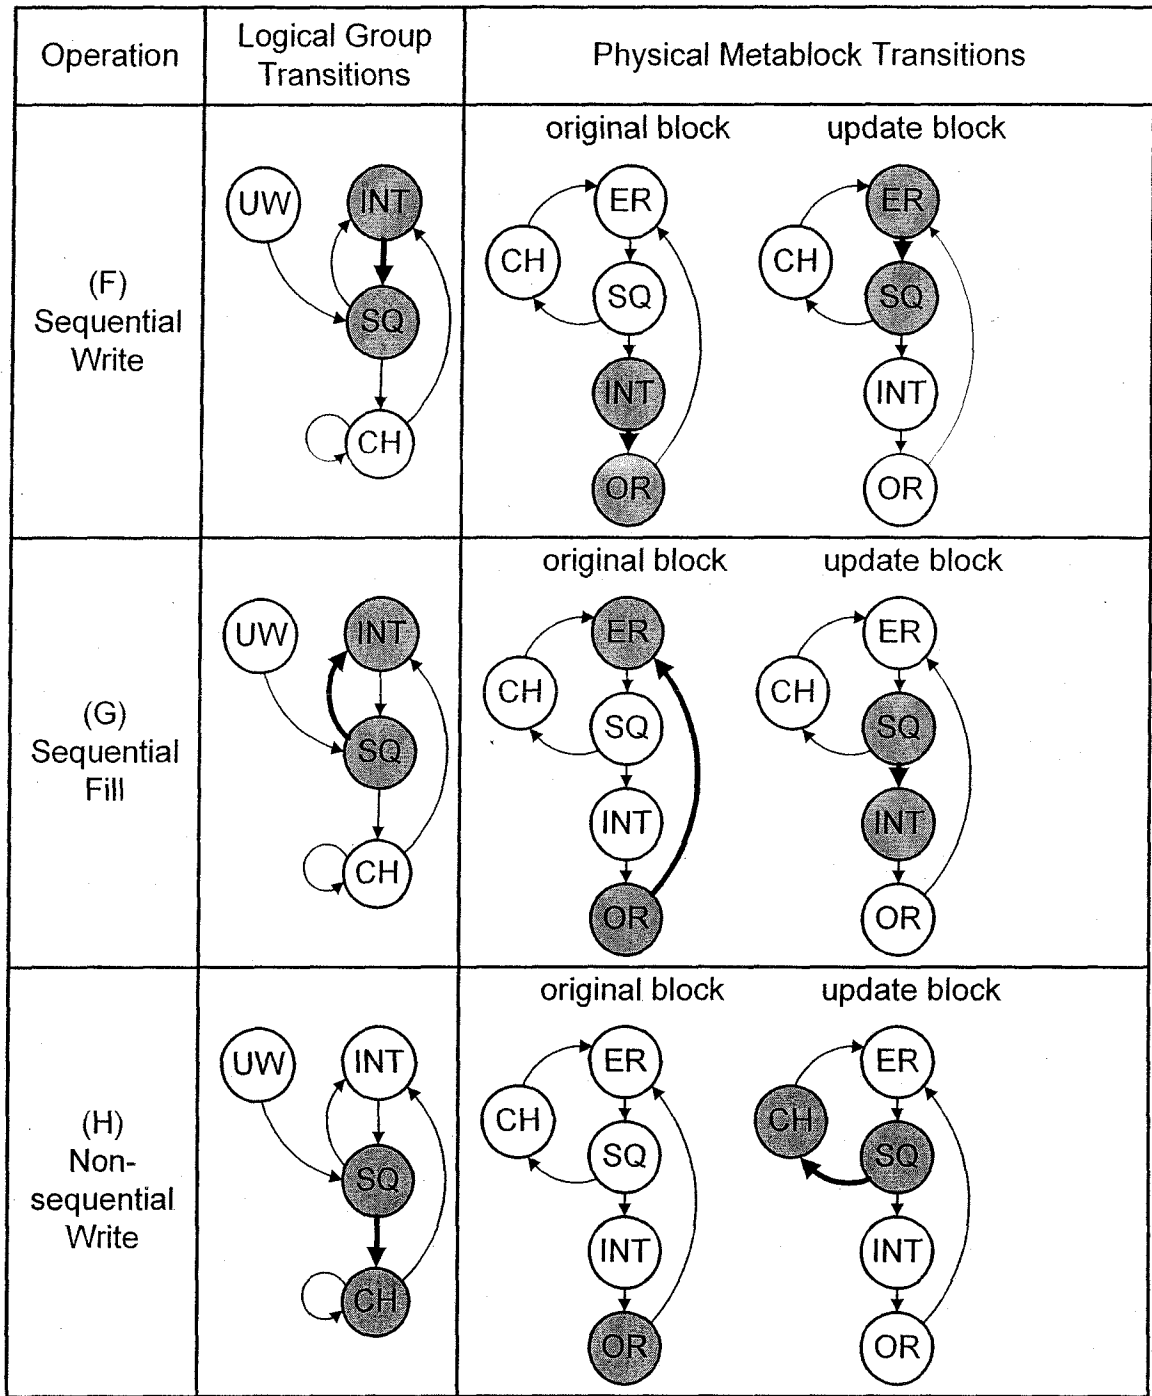
FIG. 13B



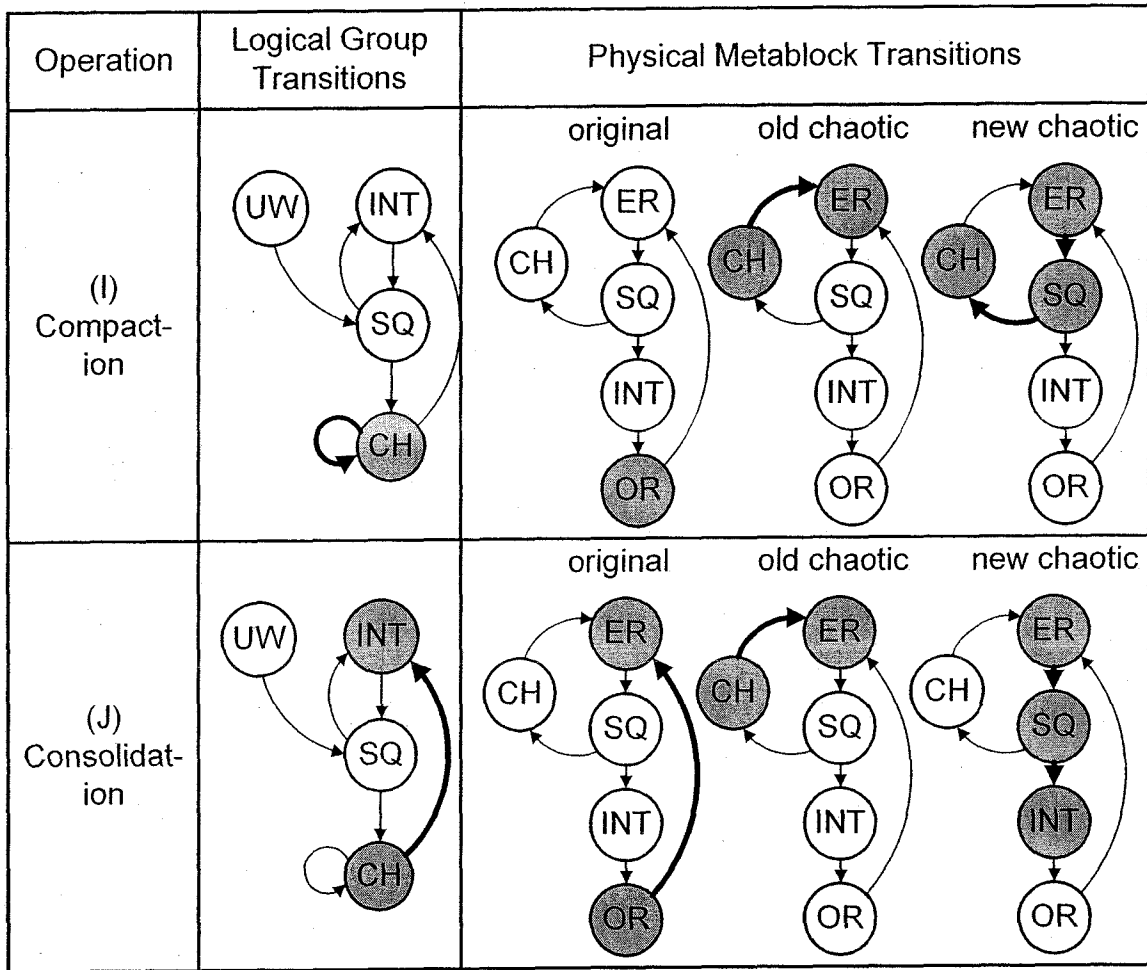
FIGs. 14(A)-14(B)



FIGs. 14(C)-14(E)



FIGs. 14(F)-14(H)



FIGs. 14(I)-14(J)

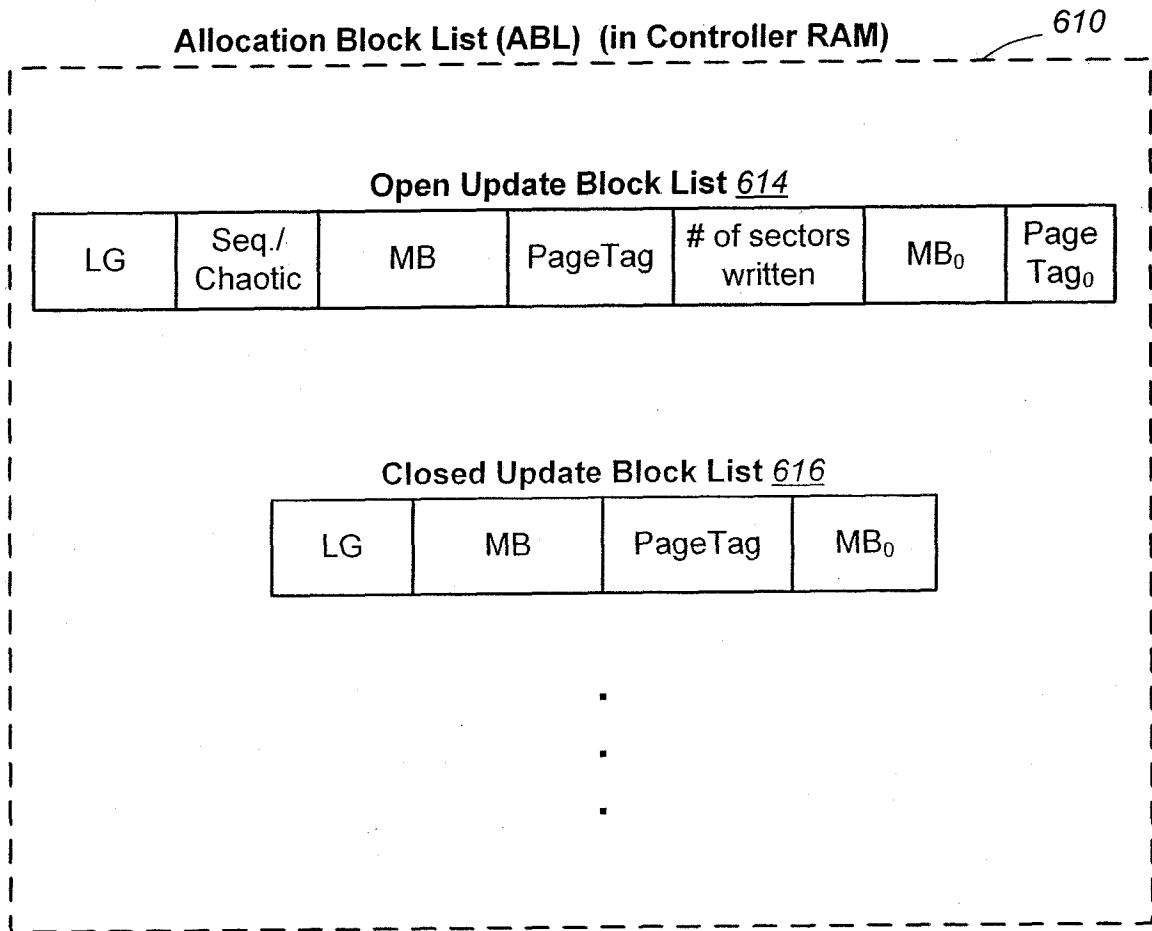


FIG. 15

Chaotic Block Index (CBI) Sector

Chaotic Block Index	Chaotic Block Info	CBI Sector Index
---------------------	--------------------	------------------

FIG. 16A

CBI Block 620

CBI Sector for LG ₀
CBI Sector for LG ₁
CBI Sector for LG ₂
⋮
CBI Sector for LG ₁
⋮
CBI Sector for LG ₁ (latest version)
CBI Sector for LG ₂₀₂
CBI Sector for LG ₁₃₅
CBI Sector for LG ₁₃₆ (last written))
unwritten

FIG. 16B

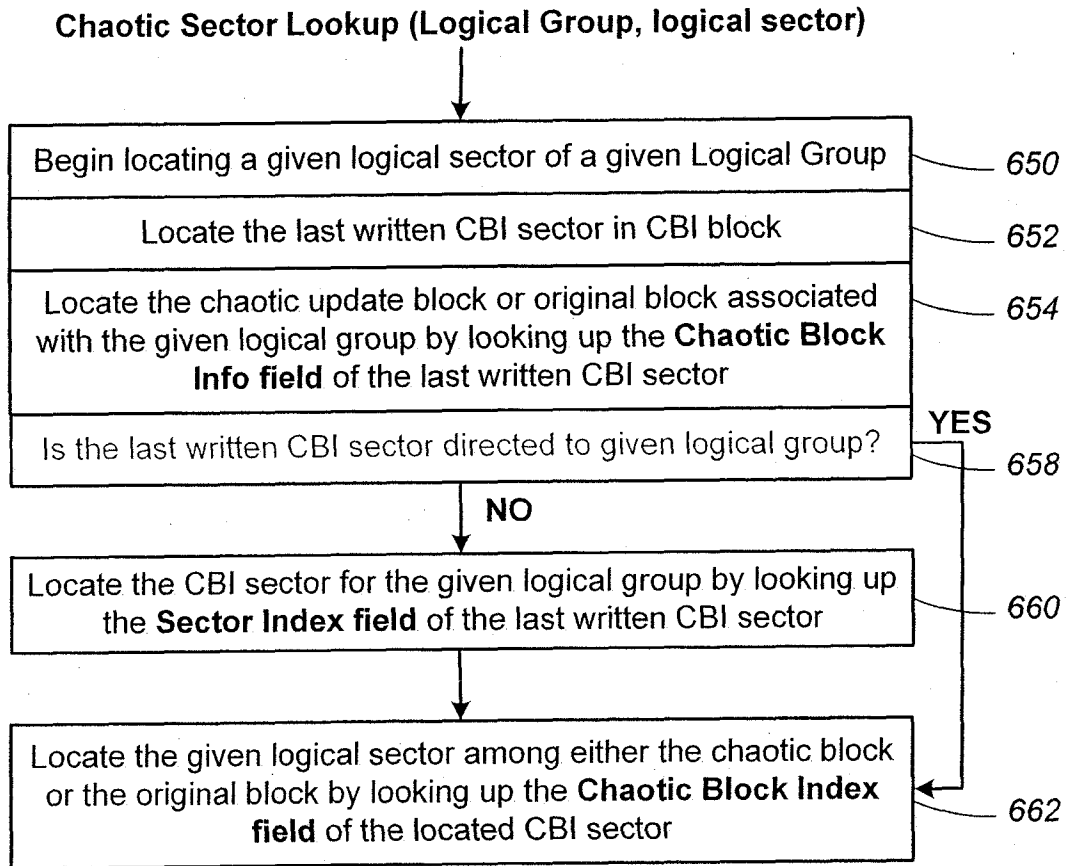


FIG. 16C

Chaotic Sector Lookup (Logical Group, Subgroup, logical sector)

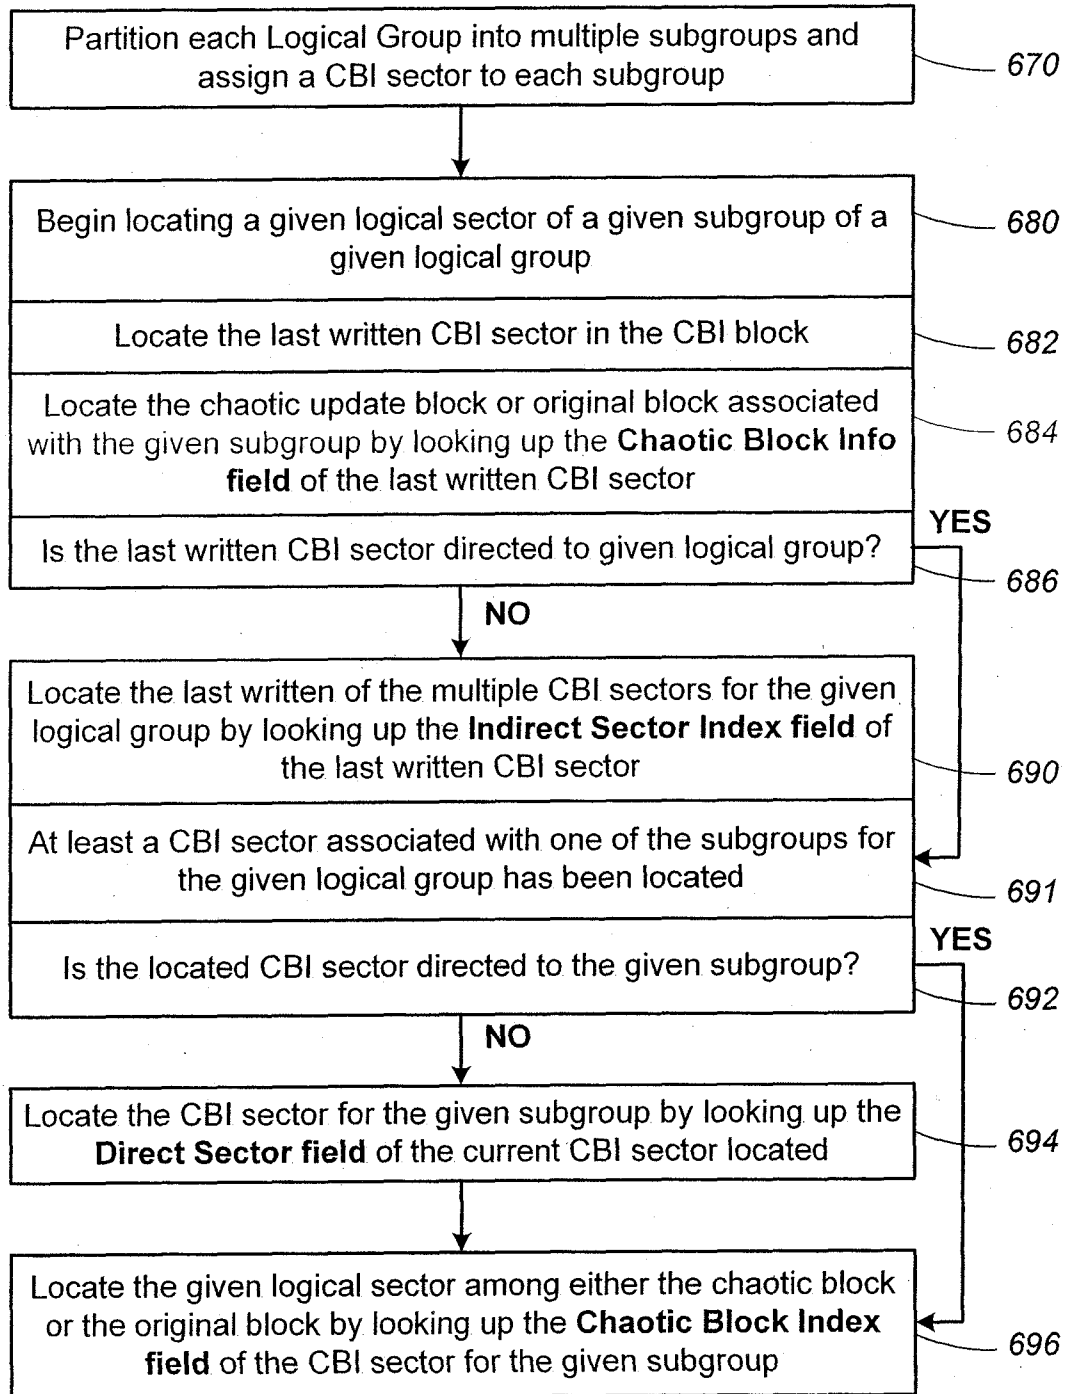



FIG. 16D

CBI Block



Chaotic Block Indexing Structure

FIG. 16E

Group Address Table (GAT) sector




FIG. 17A

GAT Block 720




FIG. 17B

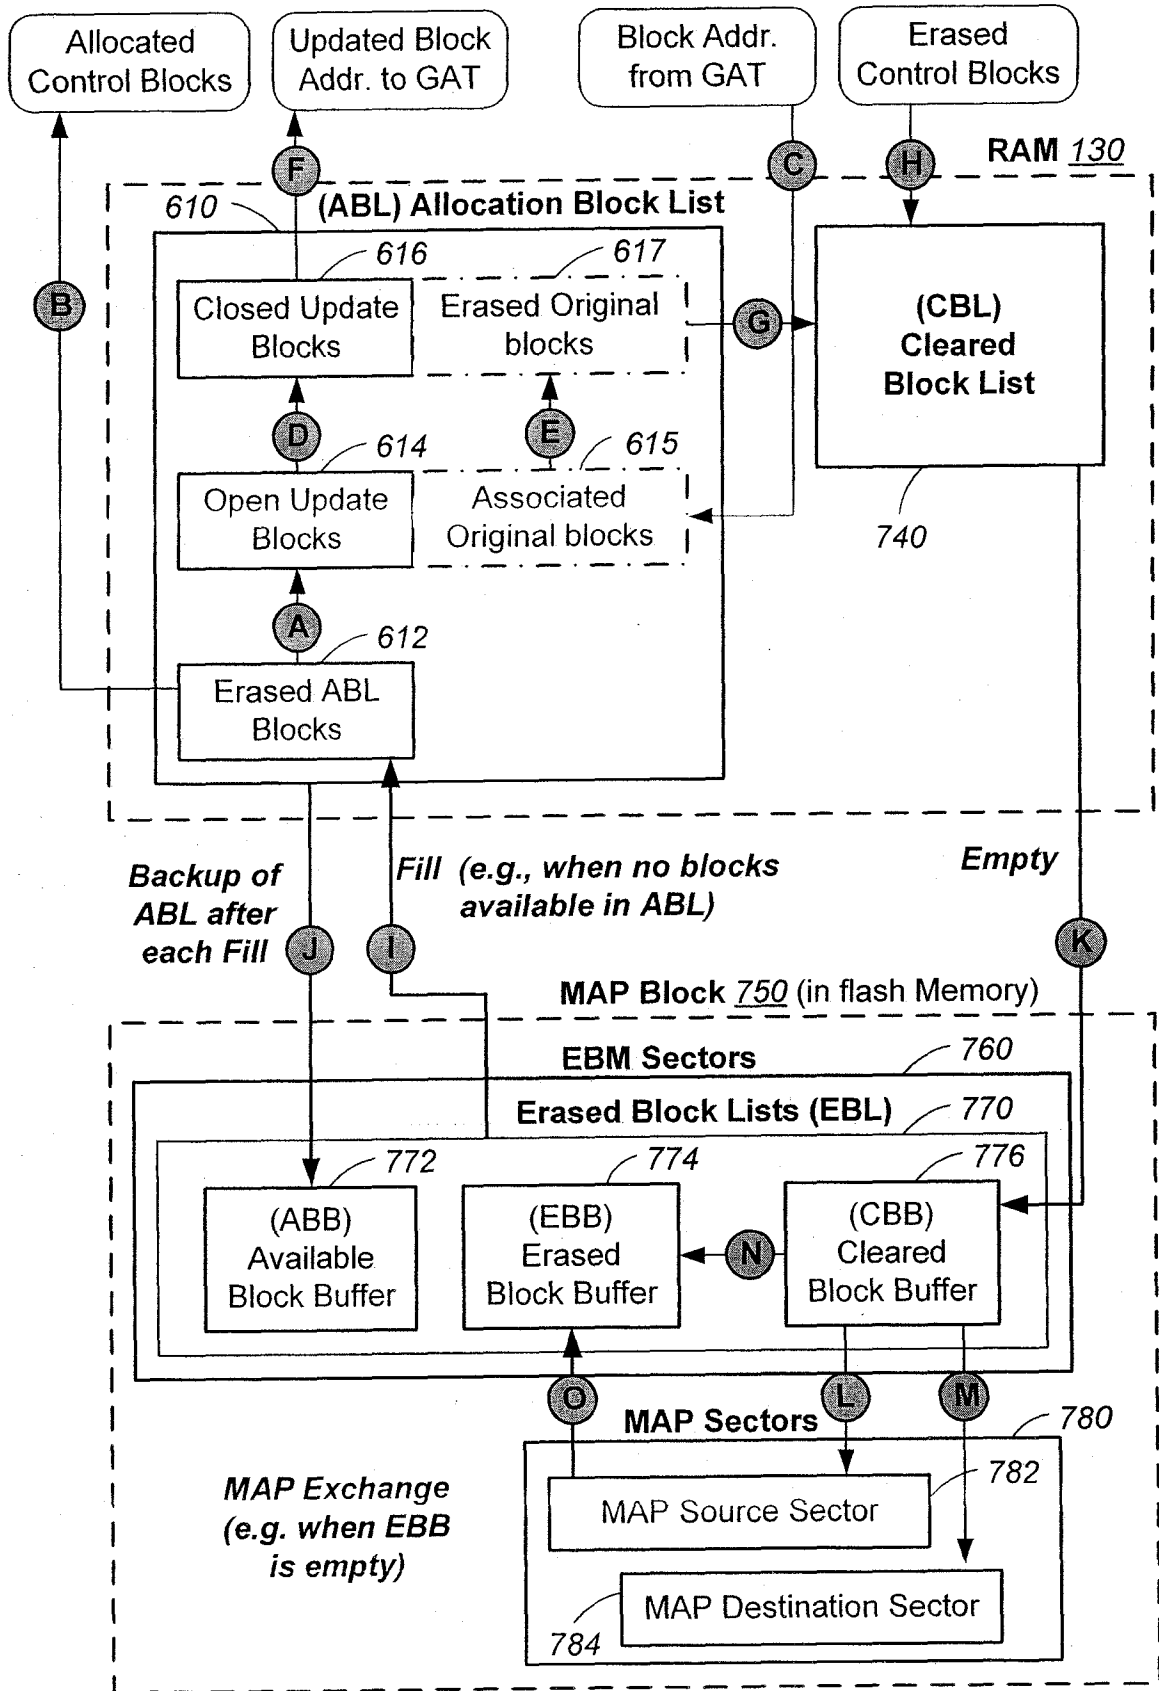
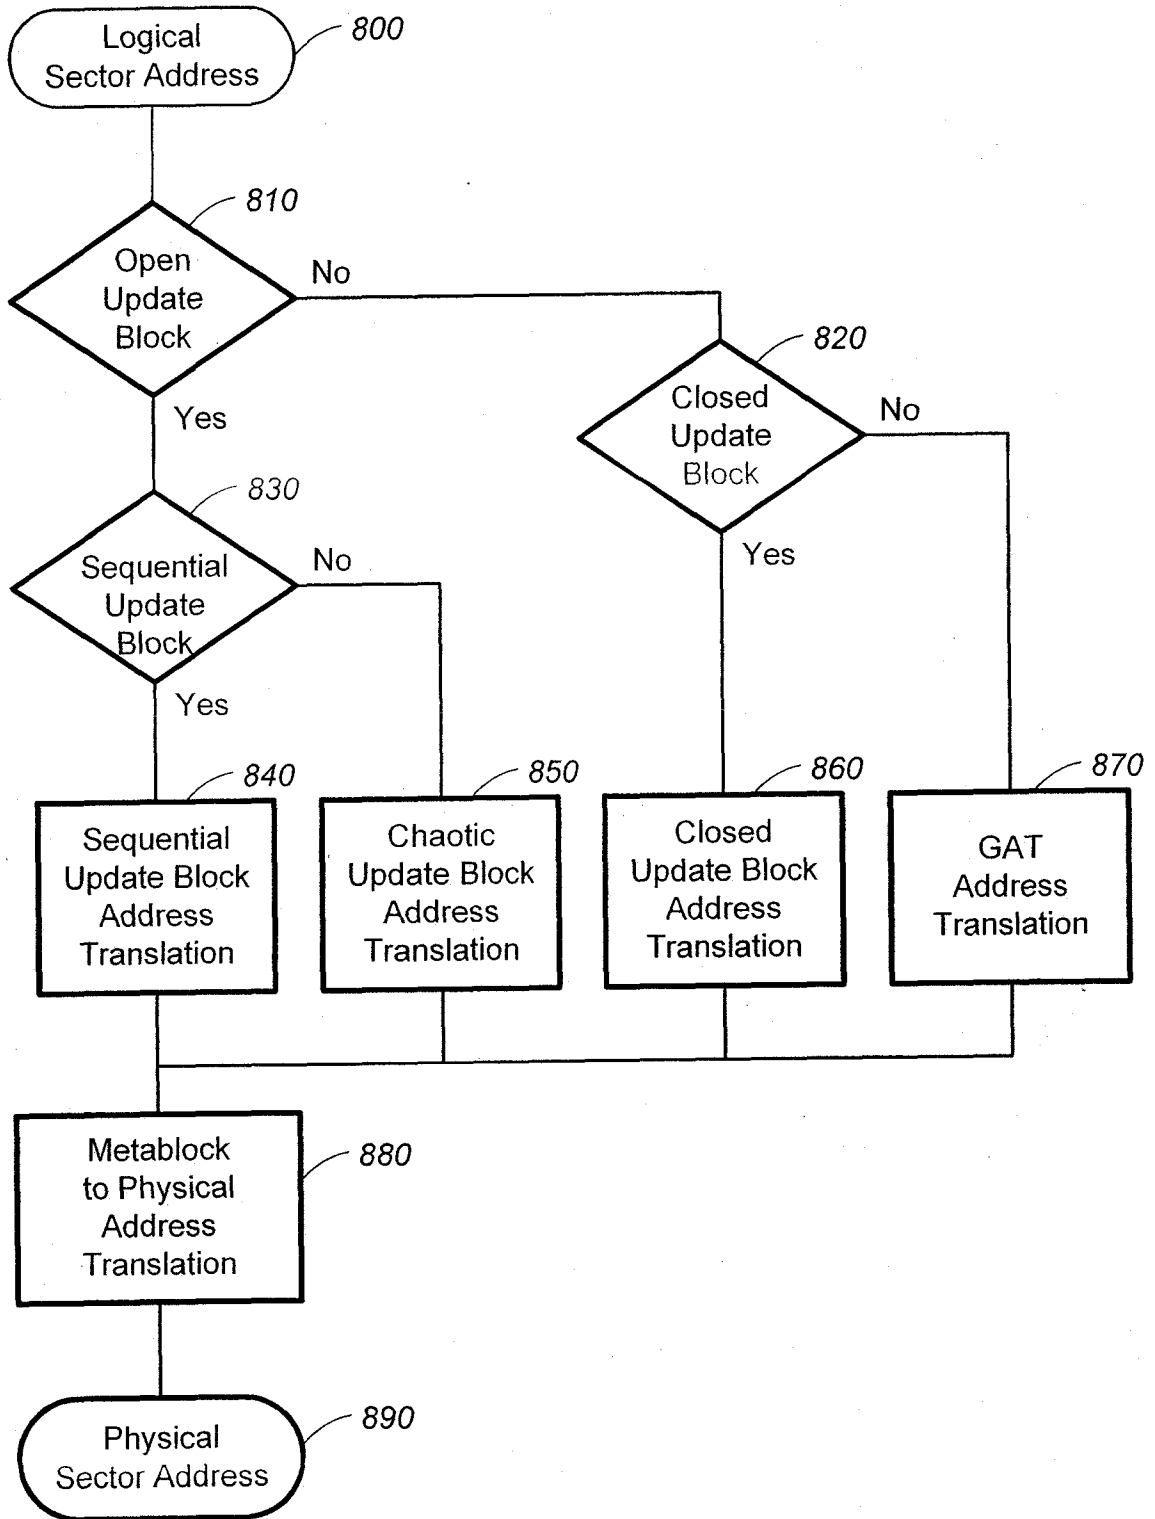
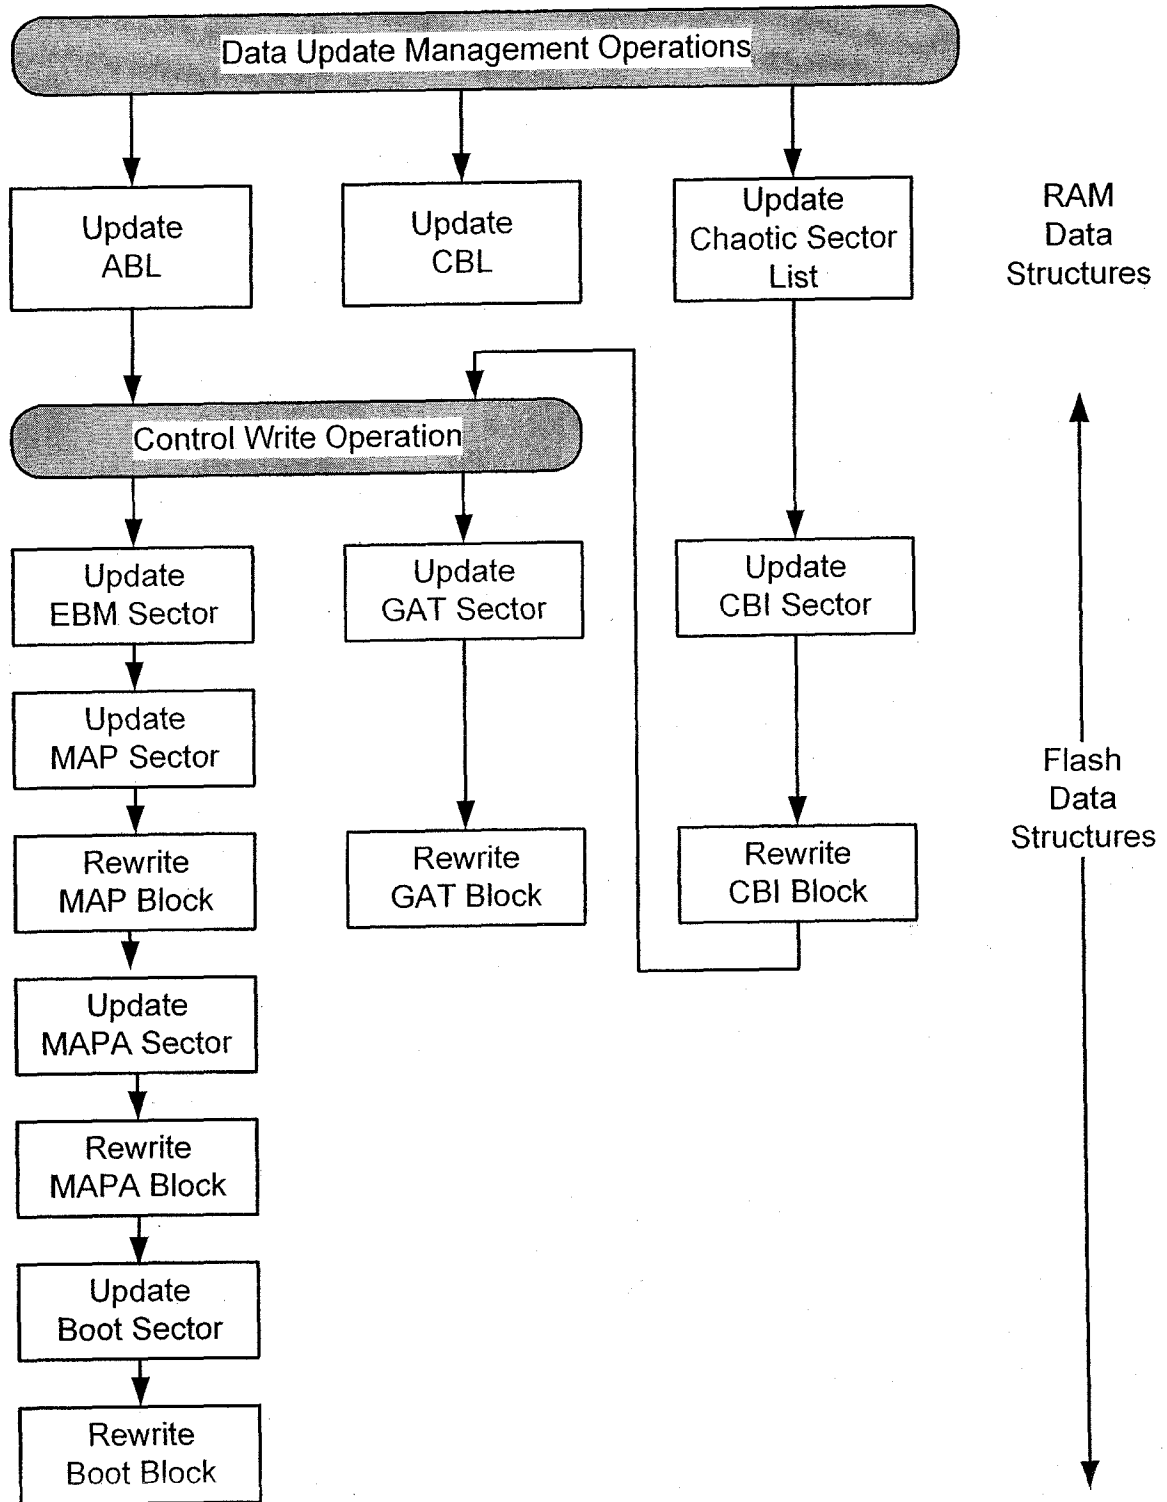


FIG. 18



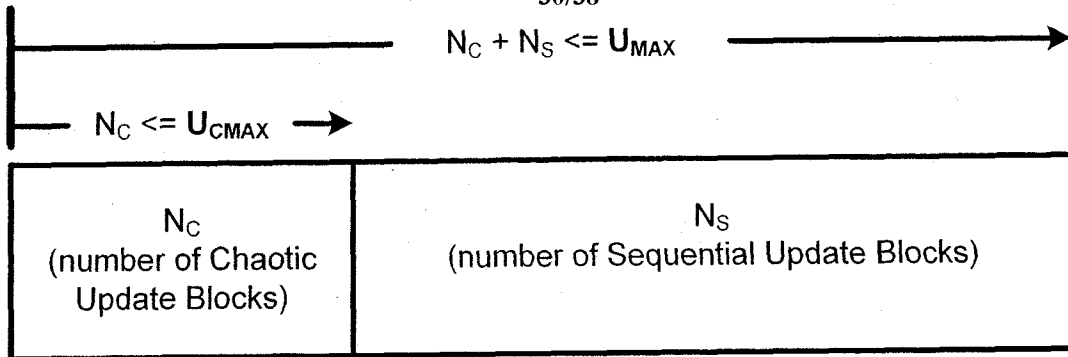
Address Translation

FIG. 19



Operations on Control Data Structures

FIG. 20



[U_{MAX} , U_{CMAX}] = [Max no. of all Update Blocks,
Max no. of all Chaotic or Non-sequential Update Blocks]


Limited Pool of Update Blocks

FIG. 21

Update Pool Restriction ($U_{MAX} - U_{CMAX}$)	U_{MAX} (Maximum No. of All Update Blocks)	U_{CMAX} (Maximum No. of Chaotic Update Blocks)
3-1	3	1
3-3	3	3
4-1	4	1
4-2	4	2
5-1	5	1
5-2	5	2
7-3	7	3


Examples of Update Blocks Limits

FIG. 22




(Update on a new logical group requiring a new update block allocation)

FIG. 23A




(In order U_{max} not be exceeded, the least active update block is closed out to make room)

FIG. 23B (Prior Art)




(A newly allocated update block S5 is introduced to the update pool for the new logical group)

FIG. 23C




(Update on a new logical group requiring a new update block allocation)

FIG. 24A



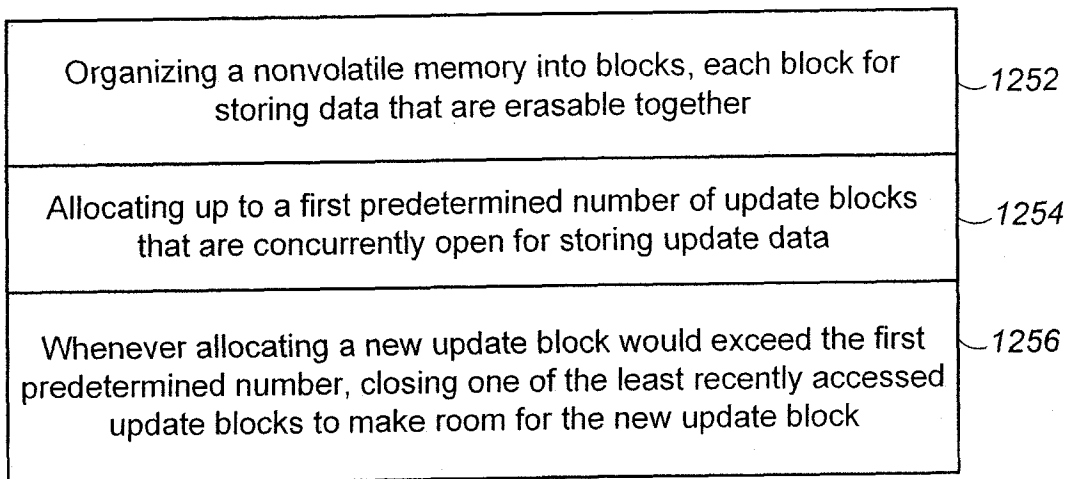
(In order U_{max} not be exceeded, the least active update block is closed out to make room)

FIG. 24B (Prior Art)



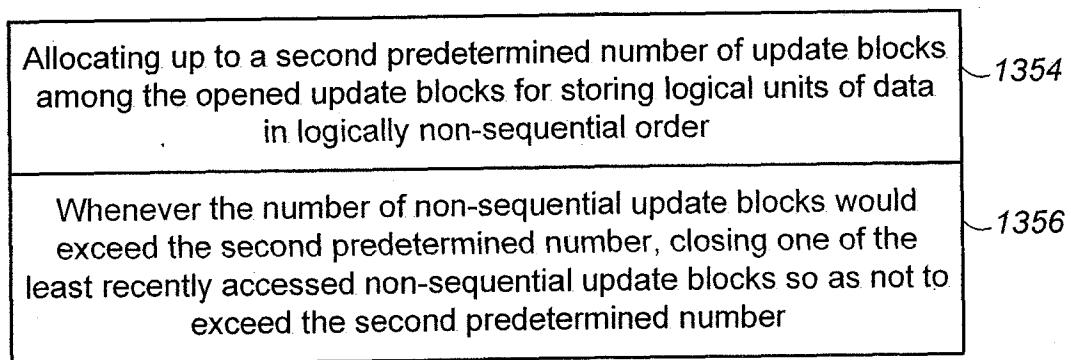
(A newly allocated update block S5 is introduced to the update pool for the new logical group)

FIG. 24C



Management of a Limited Set of Update Blocks

FIG. 25A (Prior Art)



Management of a Limited Set of Update Blocks

FIG. 25B (Prior Art)




FIG. 26A




FIG. 26B






FIG. 26C




(Updating pattern just turns an update block from sequential to chaotic)

FIG. 27A



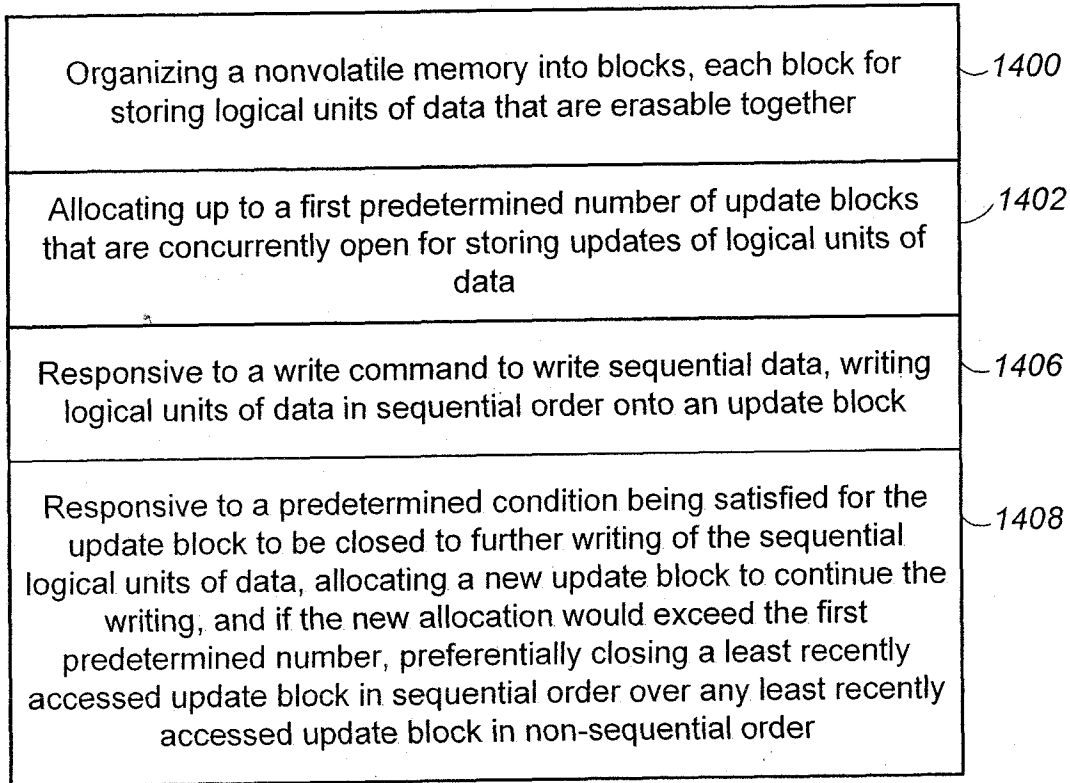
(In order U_{max} not be exceeded, the least active sequential update block is closed out to make room)

FIG. 27B



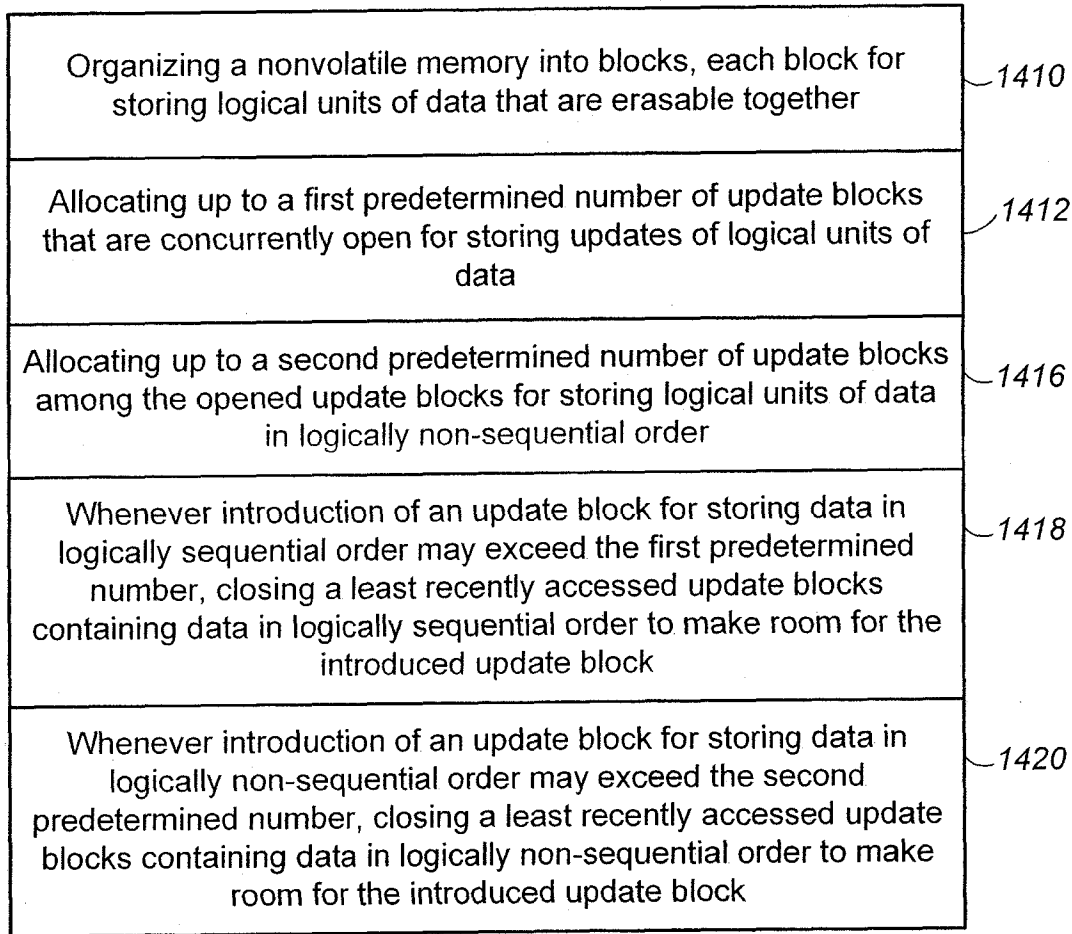
(A newly allocated update block S5 is introduced to the update pool for the new logical group)

FIG. 27C



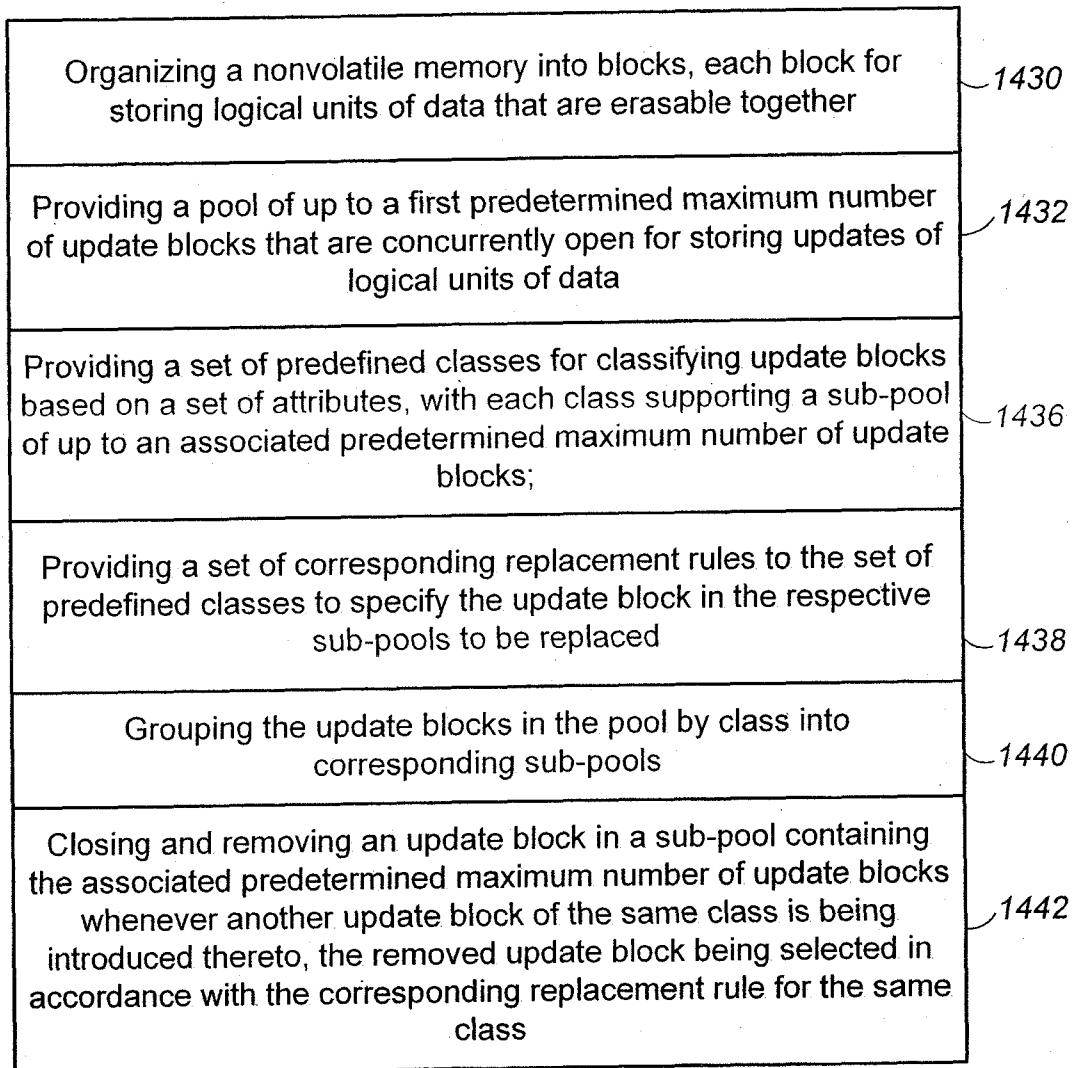
**Management of a Limited Set of Update Blocks
During a Sequential Update, according to a First Embodiment**

FIG. 28



**Management of a Limited Set of Update Blocks
Having Two Limits, according to a Second Embodiment**

FIG. 29



Class-based Replacement Rules for a Set of Update Blocks

FIG. 30