US 20110153911A1

(19) **United States**

(12) **Patent Application Publication**

Sprouse et al.

(10) **Pub. No.: US 2011/0153911 A1**

(43) **Pub. Date: Jun. 23, 2011**

(54) **METHOD AND SYSTEM FOR ACHIEVING DIE PARALLELISM THROUGH BLOCK INTERLEAVING**

(76) Inventors: **Steven Sprouse**, San Jose, CA (US); **Chris Avila**, Sunnyvale, CA (US); **Jianmin Huang**, Sunnyvale, CA (US)

(21) Appl. No.: **12/642,181**

(22) Filed: **Dec. 18, 2009**

Publication Classification

(51) **Int. Cl.**

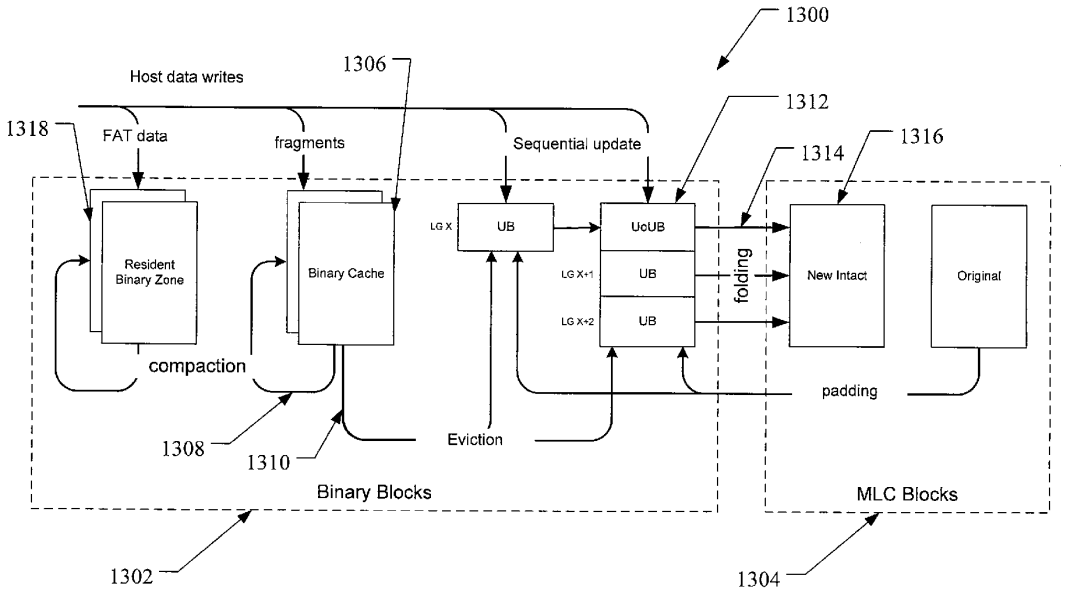
G06F 12/08 (2006.01)

G06F 12/02 (2006.01)

(52) **U.S. Cl.** **711/103**; 711/119; 711/E12.008; 711/E12.017

(57) **ABSTRACT**

A method and system for achieving die parallelism through block interleaving includes non-volatile memory having a multiple non-volatile memory dies, where each die has a cache storage area and a main storage area. A controller is configured to receive data and write sequentially addressed data to the cache storage area of a first die. The controller, after writing sequentially addressed data to the cache storage area of the first die equal to a block of the main storage area of the first die, writes additional data to a cache storage area of a next die until sequentially addressed data is written into the cache area of the next die equal to a block of the main storage area. The cache storage area may be copied to the main storage area on the first die while the cache storage area is written to on the next die.



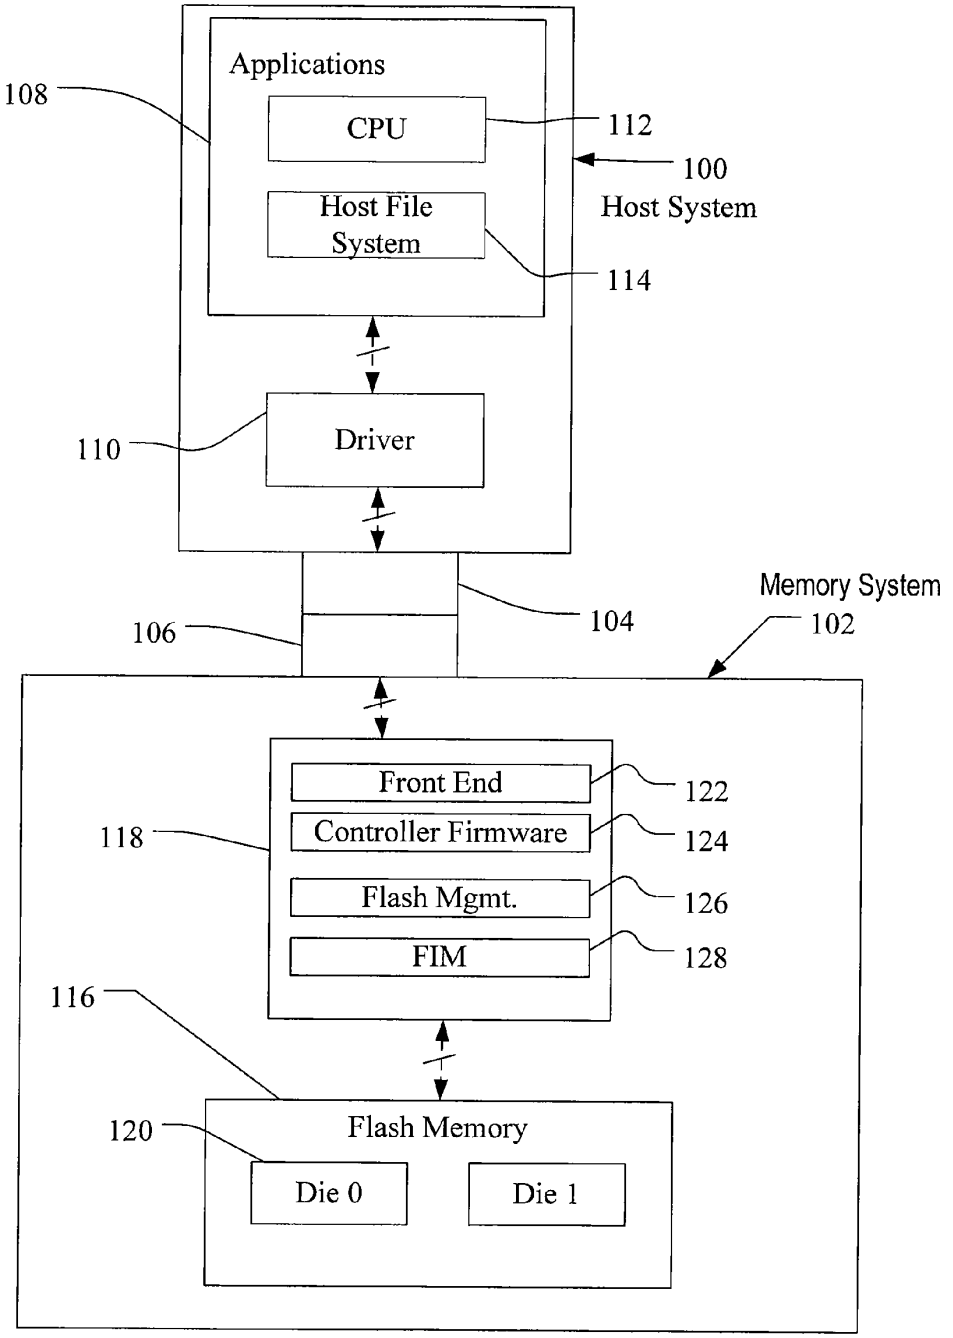

Applications
108
CPU
112
100
Host System
Host File System
114
Driver
110
Memory System
102
106
104
Front End
122
118
Controller Firmware
124
Flash Mgmt.
126
FIM
128
116
120
Flash Memory
Die 0
Die 1

FIG. 1

FIG. 2

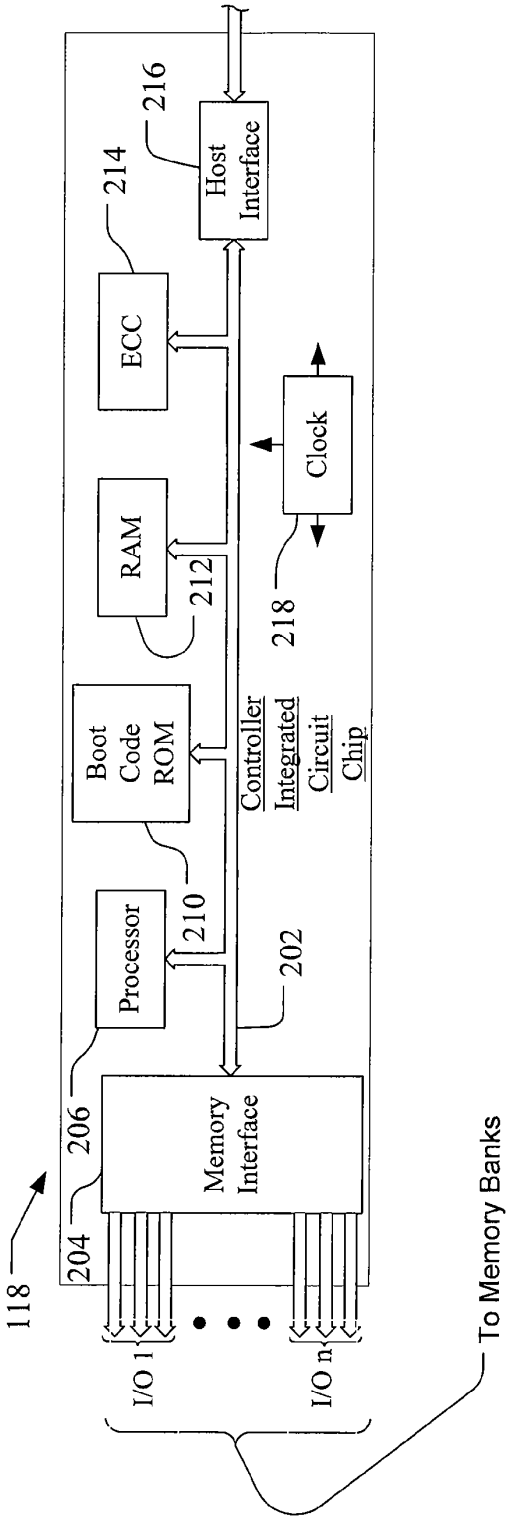
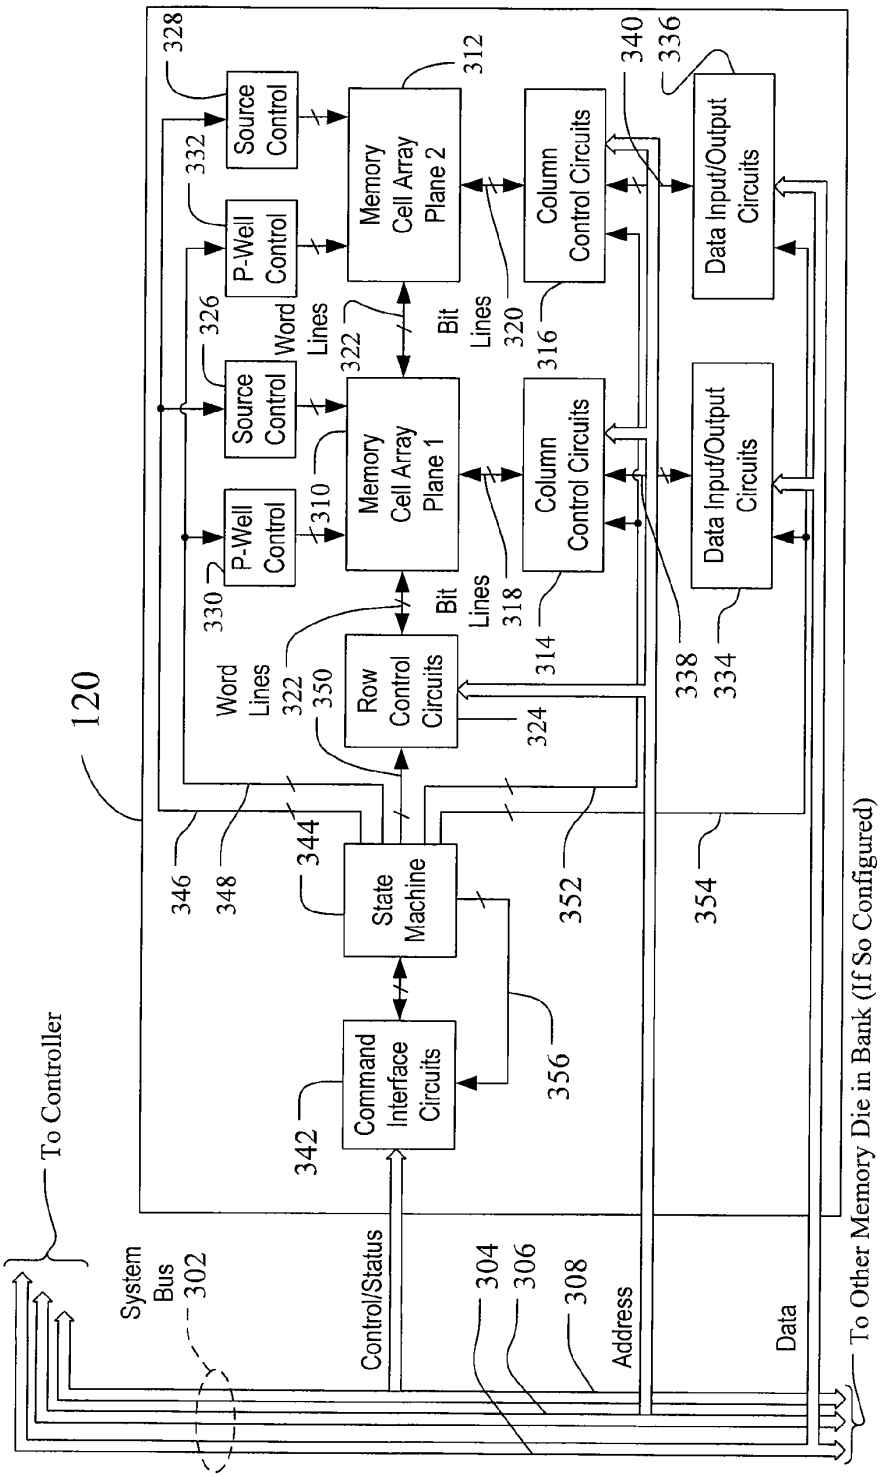

118
204
206
Memory Interface
Processor
210
Boot Code ROM
RAM
212
ECC
214
Host Interface
216
Clock
218
202
Controller Integrated Circuit Chip
I/O 1
I/O n
To Memory Banks

FIG. 3


120
To Controller
System Bus
302
Control/Status
304
306
308
Address
Data
To Other Memory Die in Bank (If So Configured)
Command Interface Circuits
342
State Machine
344
346
348
Row Control Circuits
350
Word Lines
322
324
352
354
356
P-Well Control
330
Source Control
328
P-Well Control
326
Source Control
332
Word Lines
322
Memory Cell Array Plane 1
310
Memory Cell Array Plane 2
312
Bit Lines
318
Bit Lines
320
Column Control Circuits
314
Column Control Circuits
316
Data Input/Output Circuits
334
Data Input/Output Circuits
336
338
340

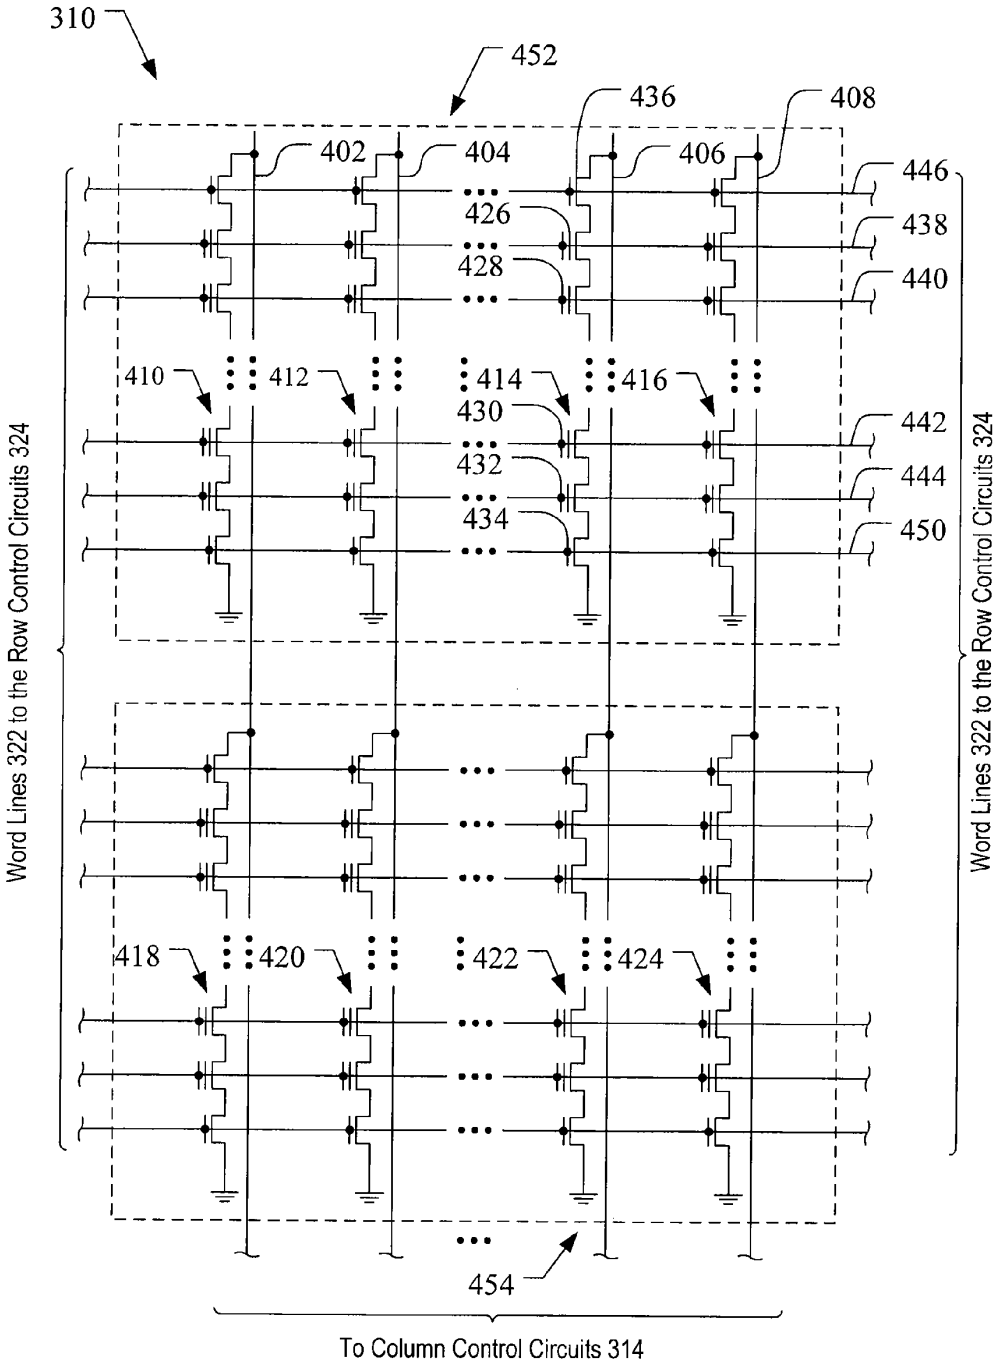

310
452
436
408
402
404
406
446
426
438
428
440
410
412
414
416
430
442
432
444
434
450
Word Lines 322 to the Row Control Circuits 324
Word Lines 322 to the Row Control Circuits 324
418
420
422
424
454
To Column Control Circuits 314

FIG. 4

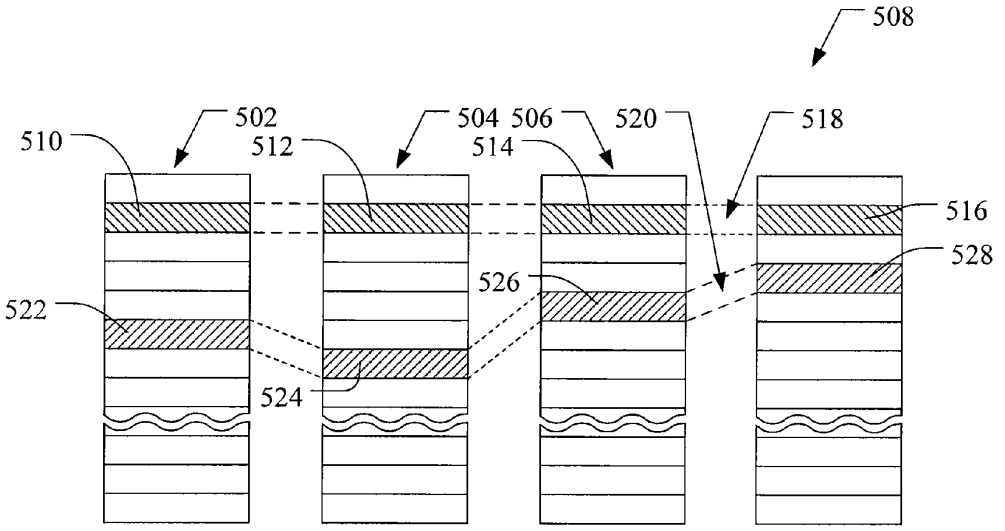

508
502
504
506
520
518
510
512
514
516
528
526
522
524

FIG. 5

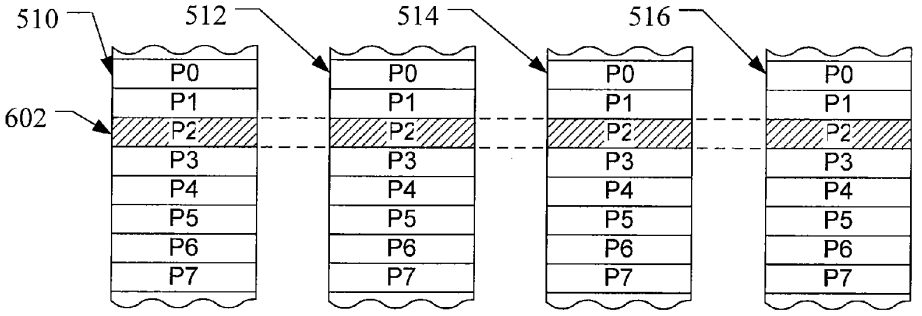

510
512
514
516
602
P0
P1
P2
P3
P4
P5
P6
P7

FIG. 6

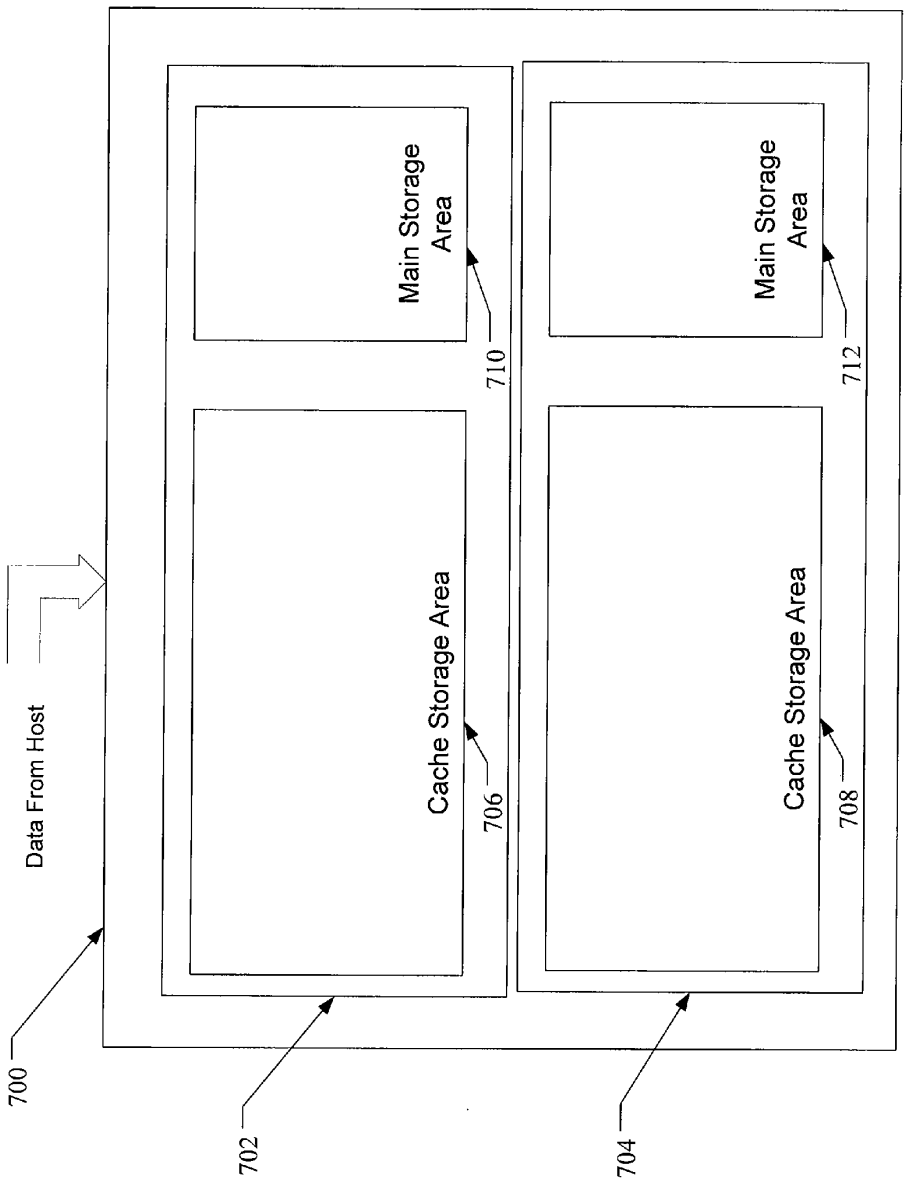

Data From Host
Main Storage Area
Cache Storage Area
Main Storage Area
Cache Storage Area
700
702
704
706
708
710
712

FIG. 7

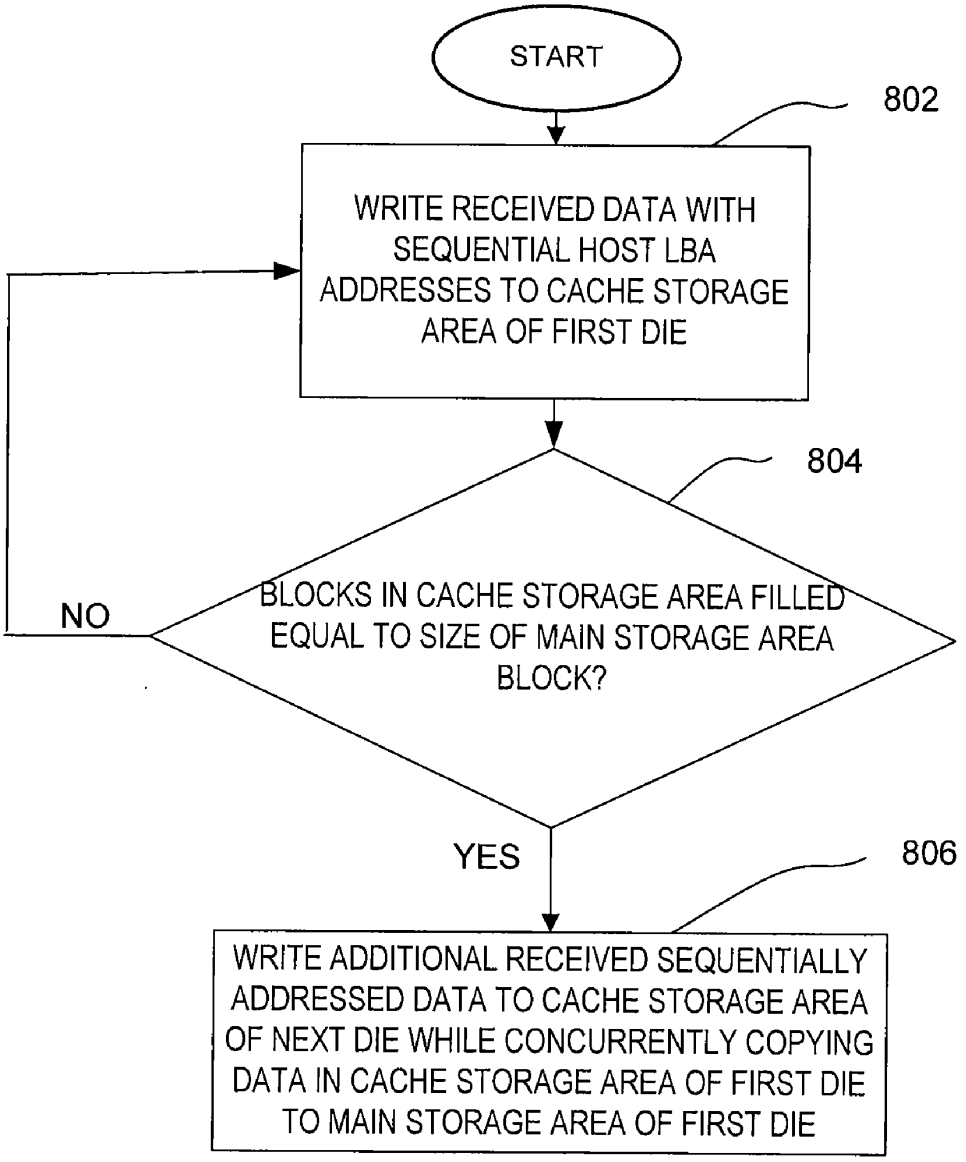

START
802
WRITE RECEIVED DATA WITH SEQUENTIAL HOST LBA ADDRESSES TO CACHE STORAGE AREA OF FIRST DIE
804
BLOCKS IN CACHE STORAGE AREA FILLED EQUAL TO SIZE OF MAIN STORAGE AREA BLOCK?
NO
YES
806
WRITE ADDITIONAL RECEIVED SEQUENTIALLY ADDRESSED DATA TO CACHE STORAGE AREA OF NEXT DIE WHILE CONCURRENTLY COPYING DATA IN CACHE STORAGE AREA OF FIRST DIE TO MAIN STORAGE AREA OF FIRST DIE

FIG. 8

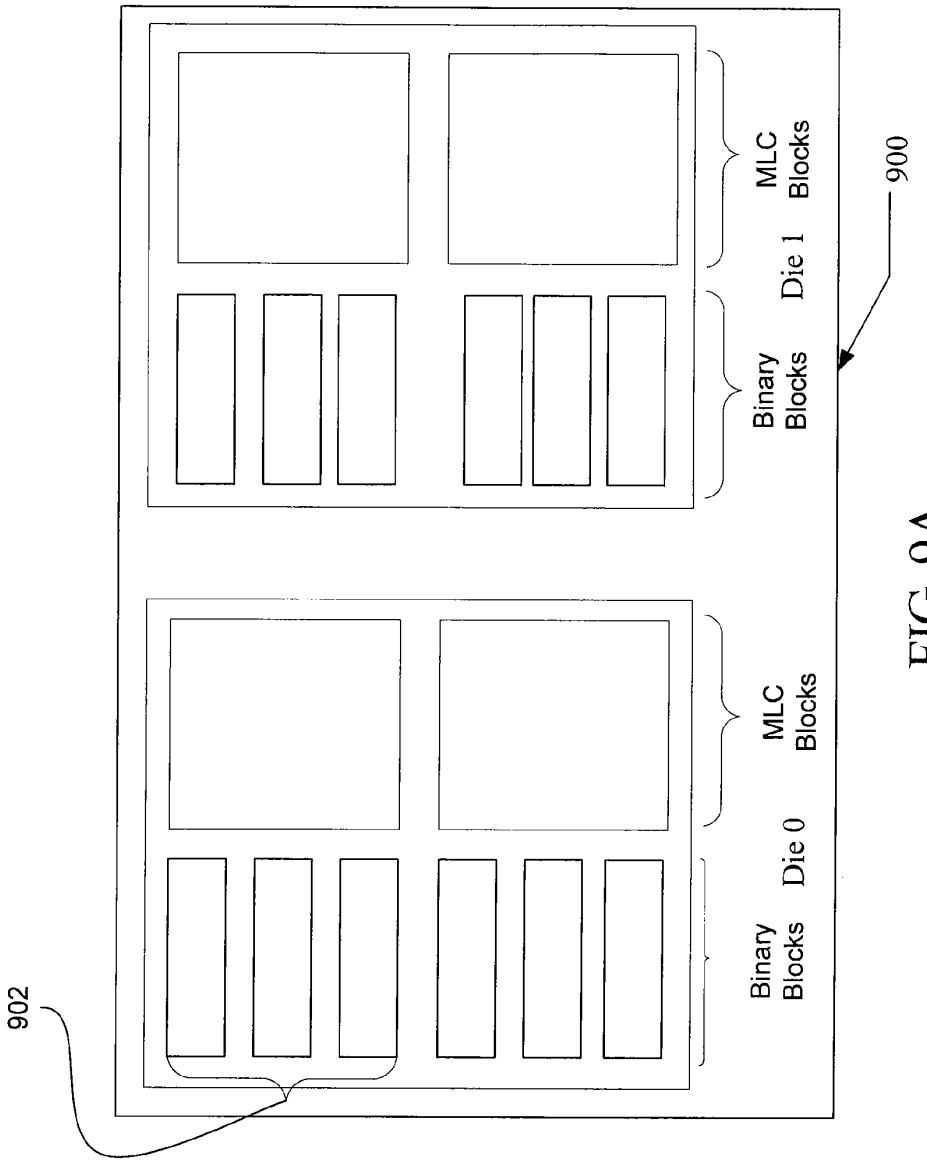

MLC Blocks
Die 1
Binary Blocks
900
MLC Blocks
Die 0
Binary Blocks
902

FIG. 9A

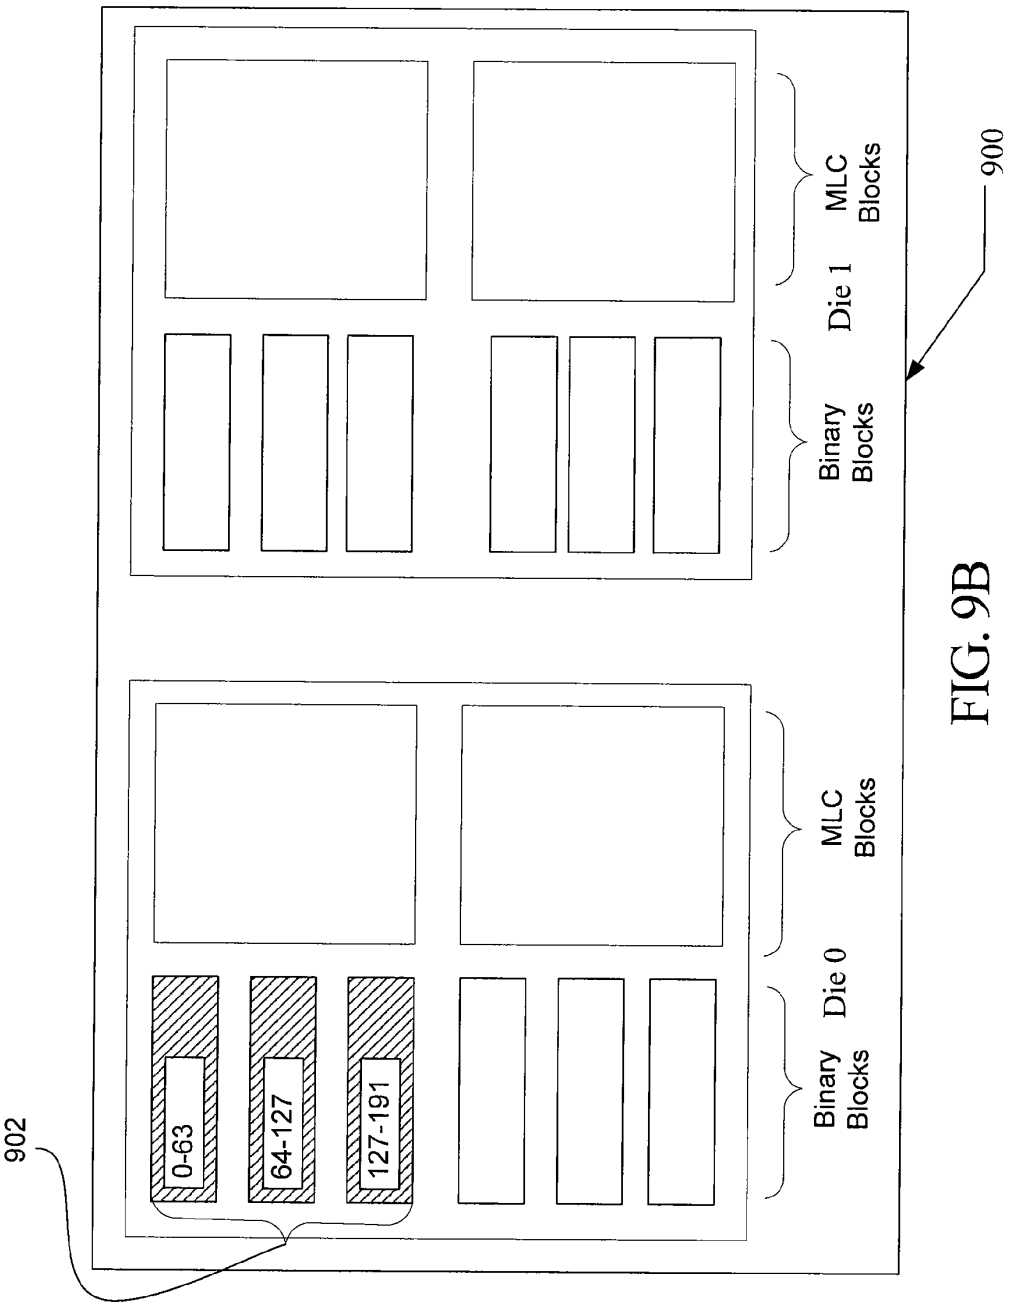

900
902
Die 1
MLC Blocks
Binary Blocks
Die 0
MLC Blocks
Binary Blocks
0-63
64-127
127-191

FIG. 9B

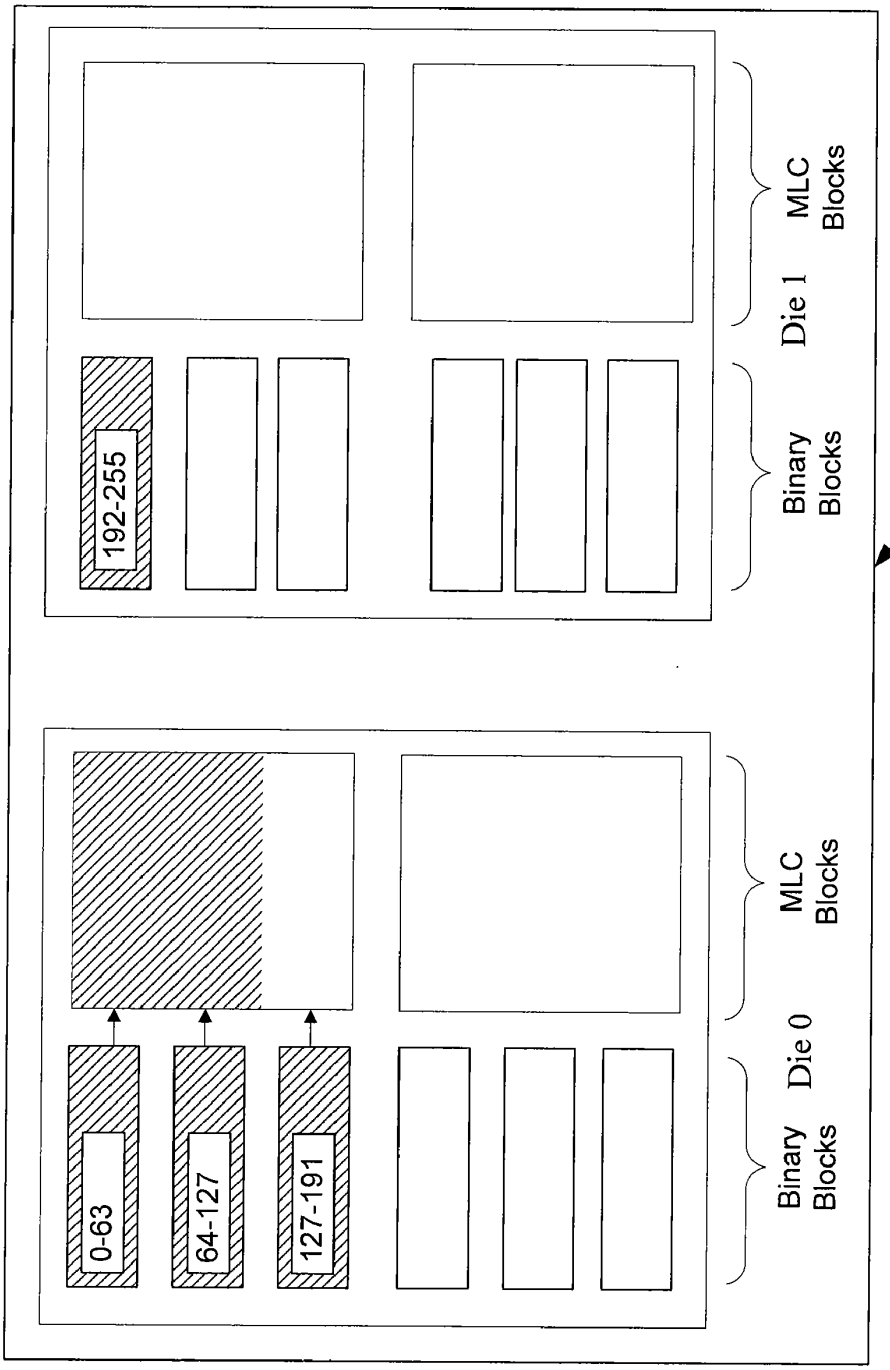

900
Die 1
MLC Blocks
Binary Blocks
192-255
Die 0
MLC Blocks
Binary Blocks
0-63
64-127
127-191

FIG. 9C

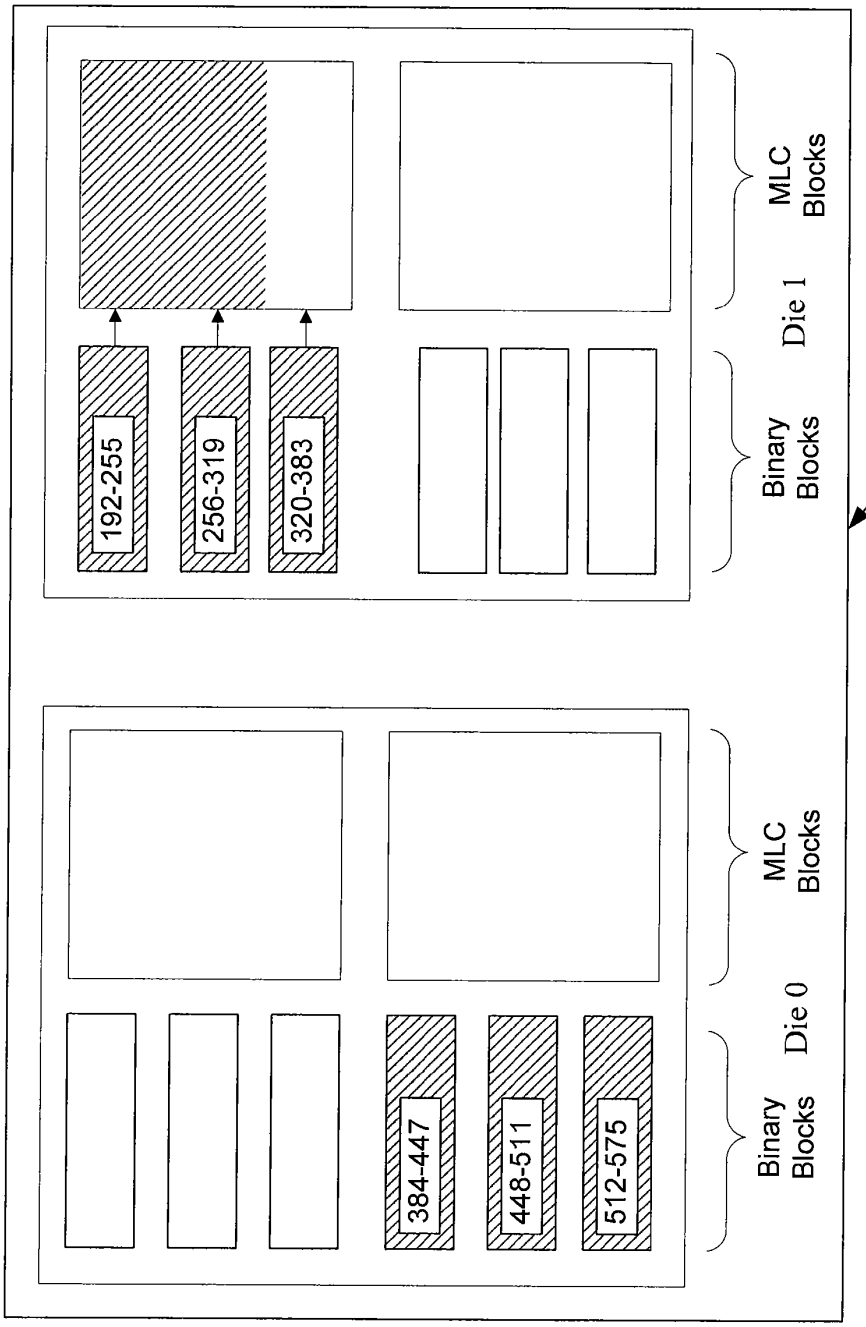

900
Die 1
MLC Blocks
Binary Blocks
192-255
256-319
320-383
Die 0
MLC Blocks
Binary Blocks
384-447
448-511
512-575

FIG. 9D

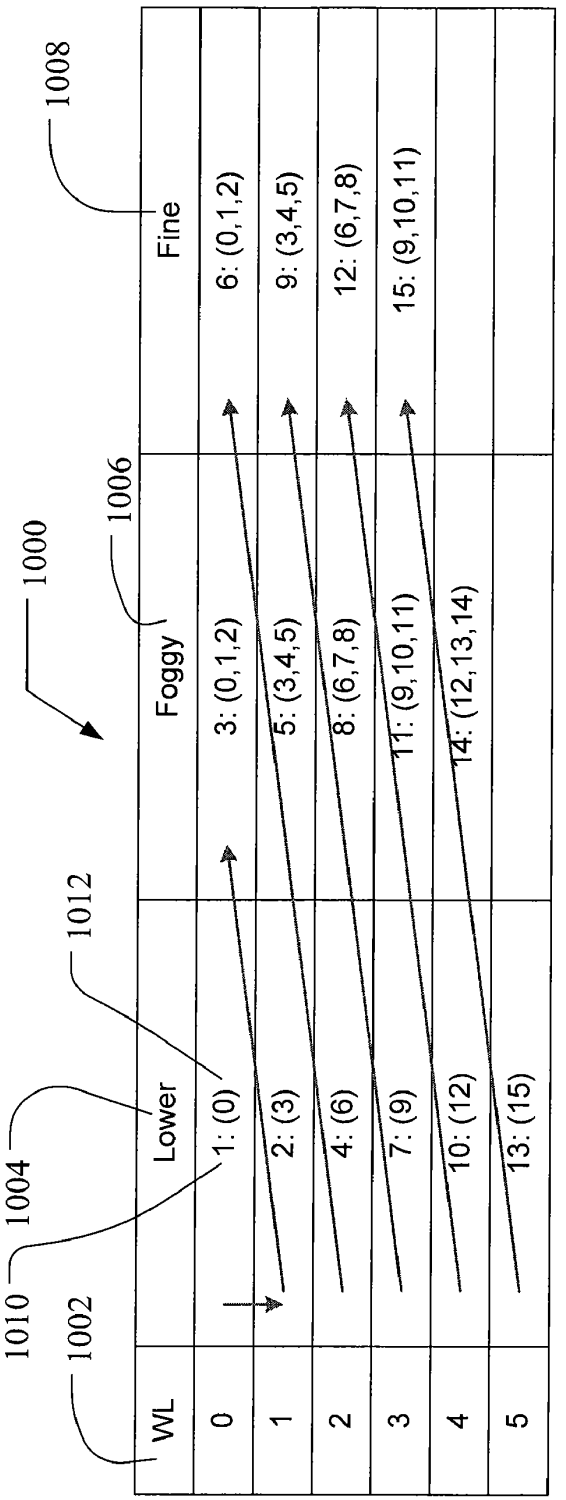

1000
1002
1004
1006
1008
1010
1012
WL
Lower
Foggy
Fine
0
1
2
3
4
5
1: (0)
2: (3)
4: (6)
7: (9)
10: (12)
13: (15)
3: (0,1,2)
5: (3,4,5)
8: (6,7,8)
11: (9,10,11)
14: (12,13,14)
6: (0,1,2)
9: (3,4,5)
12: (6,7,8)
15: (9,10,11)

FIG. 10

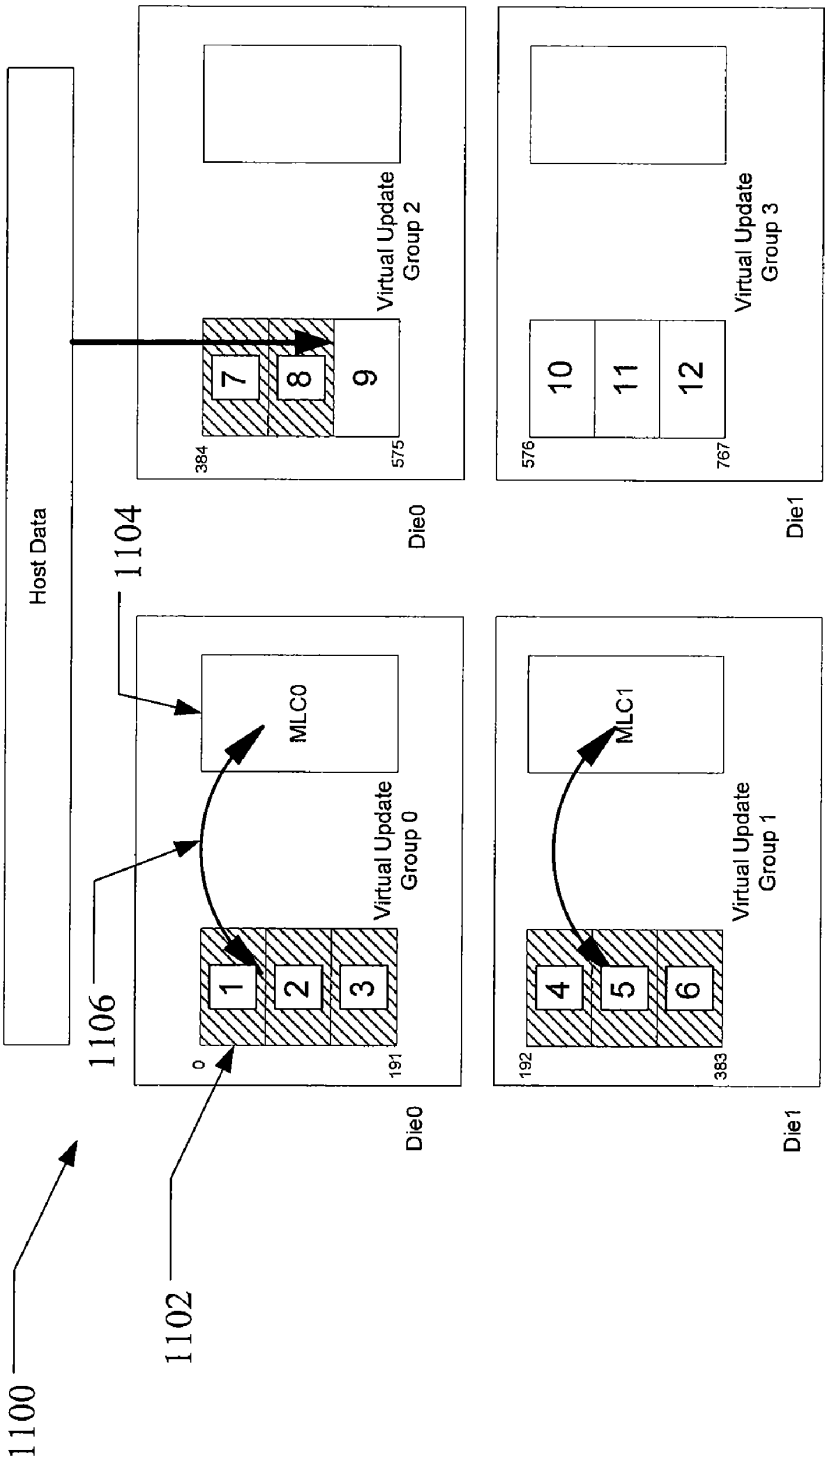

1100
1102
1104
1106
Host Data
MLC0
MLC1
Virtual Update Group 0
Virtual Update Group 1
Virtual Update Group 2
Virtual Update Group 3
Die0
Die1
Die0
Die1
0
191
192
383
384
575
576
767
1
2
3
4
5
6
7
8
9
10
11
12

FIG. 11

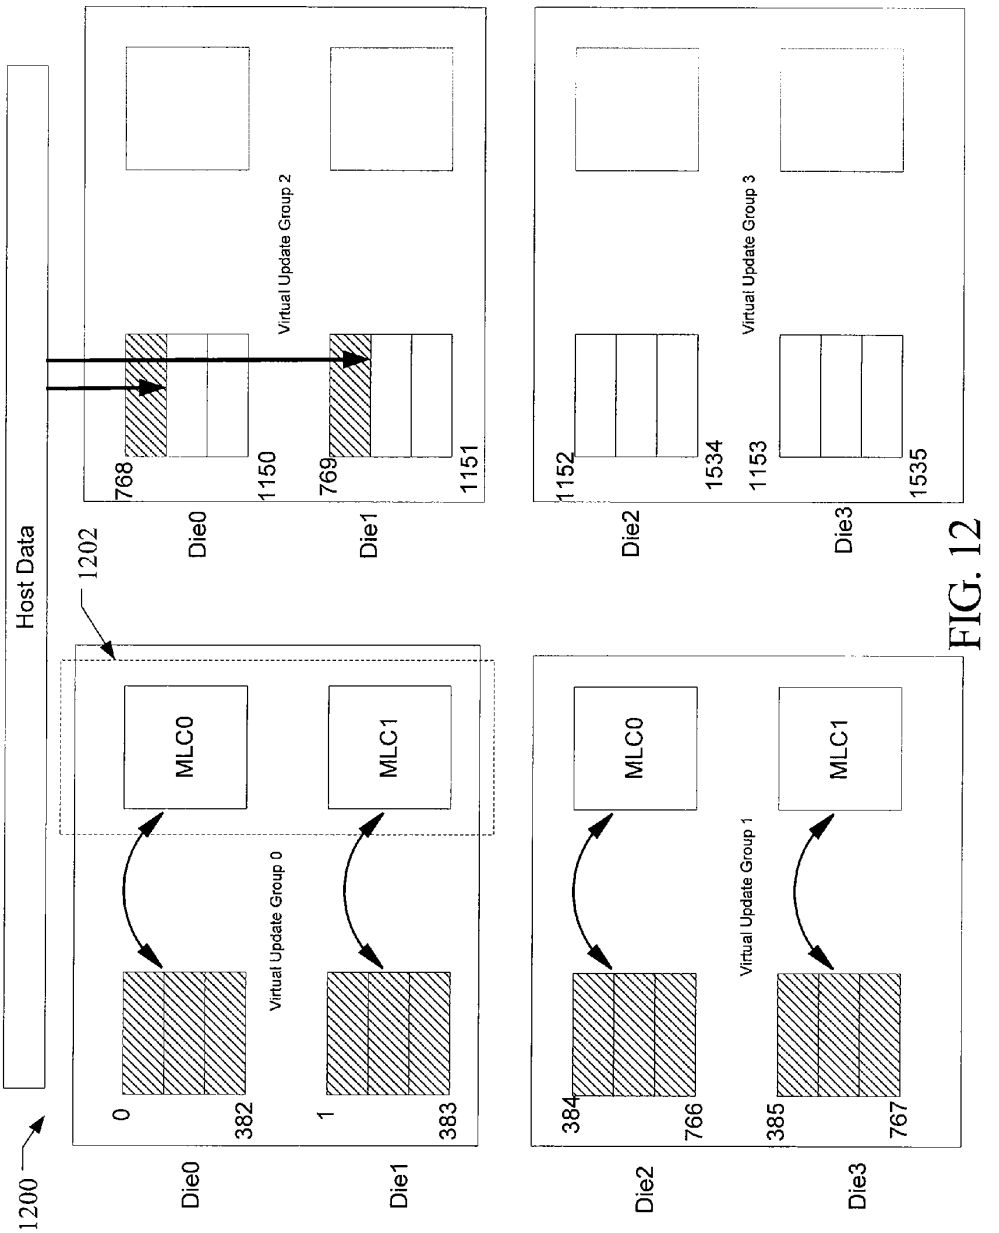

1200
Host Data
1202
MLC0
MLC1
Virtual Update Group 0
Virtual Update Group 1
Virtual Update Group 2
Virtual Update Group 3
Die0
Die1
Die2
Die3
0
382
1
383
384
766
385
767
768
1150
769
1151
1152
1534
1153
1535

FIG. 12

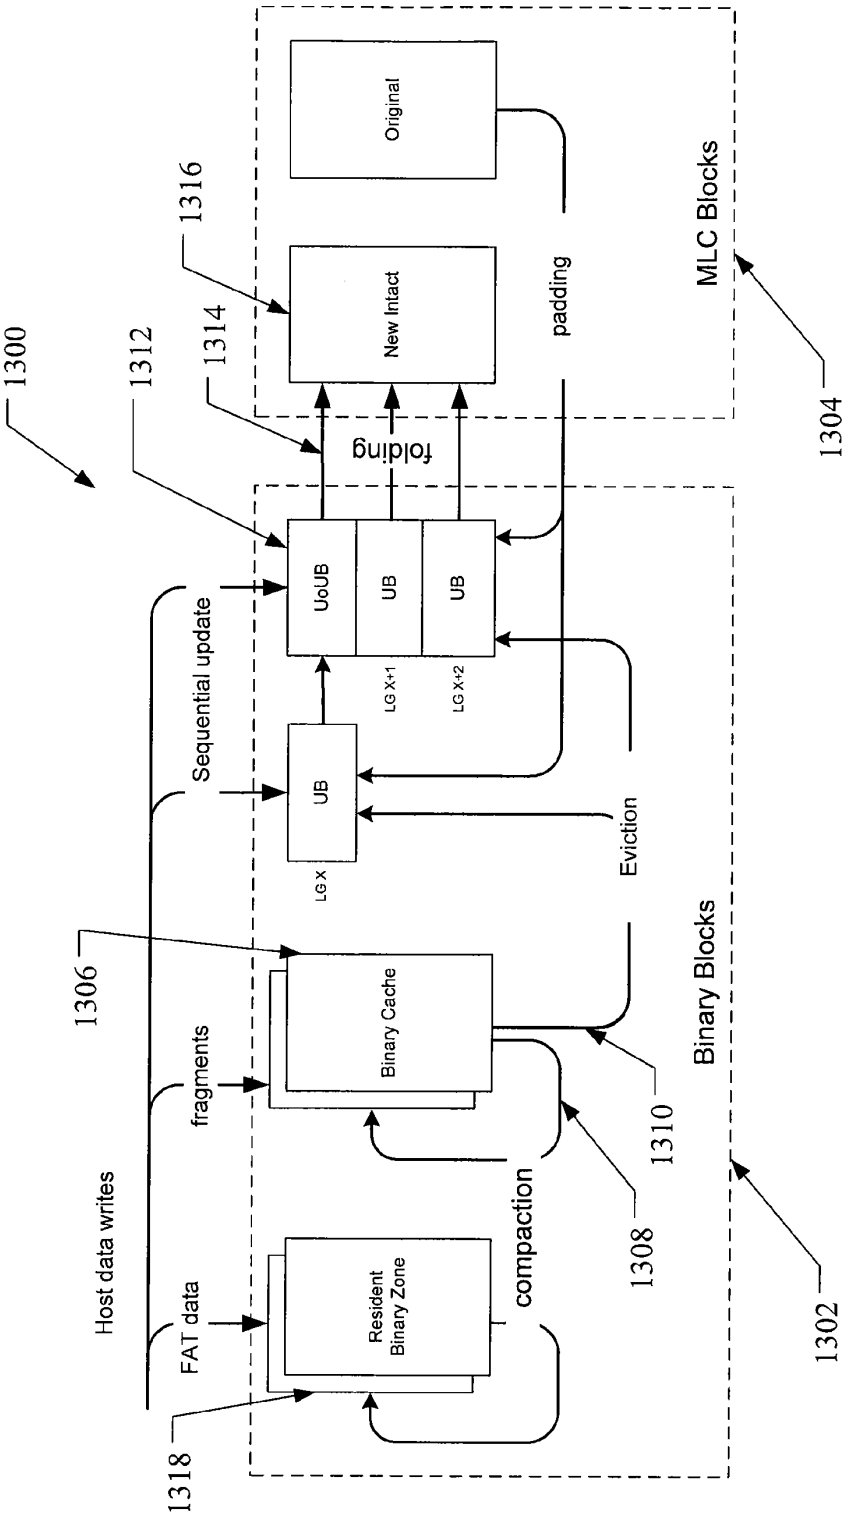

1300
1312
1314
1316
1304
1306
1308
1310
1302
1318
Original
New Intact
padding
MLC Blocks
folding
UoUB
UB
UB
LG X+1
LG X+2
Sequential update
UB
LG X
Eviction
Binary Cache
fragments
Host data writes
FAT data
Resident Binary Zone
compaction
Binary Blocks

FIG. 13

METHOD AND SYSTEM FOR ACHIEVING DIE PARALLELISM THROUGH BLOCK INTERLEAVING

TECHNICAL FIELD

[0001] This application relates generally to managing data in a memory system. More specifically, this application relates to the operation of a memory system to improve die parallelism in re-programmable non-volatile semiconductor flash memory having multiple die.

BACKGROUND

[0002] When writing data to a conventional flash data memory system, a host typically assigns unique logical addresses to sectors, clusters or other units of data within a continuous virtual address space of the memory system. The host writes data to, and reads data from, addresses within the logical address space of the memory system. The memory system then commonly maps data between the logical address space and the physical blocks of the memory. Each block consists of a number of pages.

[0003] A common practice with flash memory is to operate multiple memory die in parallel to increase write performance. One method for achieving the multiple die parallelism is known as page level die interleaving, where one page is written to each of the different die in the memory system before writing a next page to each die. In this technique, sequentially addressed data that is written from the host to the memory system is striped across multiple die a page at a time. The write performance of a multiple die flash memory may be improved using this technique, but the size of the resulting metablock tends to be large. Larger metablock sizes can degrade the performance of a memory system when the memory system receives random writes, rather than long sequential writes, from a host.

BRIEF SUMMARY

[0004] In order to address the need for improved memory management in a multi-bank memory system, methods and systems are disclosed herein for achieving memory die parallelism at block level.

[0005] According to one aspect, a memory system includes a non-volatile memory having a plurality of non-volatile memory dies, wherein each of the plurality of non-volatile memory dies has a cache storage area and a main storage area. The memory includes a controller configured to receive data, where the received data is associated with host logical block addresses. The controller is further configured to write sequentially addressed data to the cache storage area of a first die of the plurality of non-volatile memory dies and, after writing an amount of sequentially addressed data to the cache storage area of the first die equal to a block of the main storage area of the first die, to write additional received sequentially addressed data to a cache storage area of a next die of the plurality of dies until a second amount of sequentially addressed data is written into the cache area of the next die equal to a block of the main storage area of the next die. The data received may be data from an external host as a result of a host write operation or from an internal source in the memory system as a result of a housekeeping operation in the memory system.

[0006] In another aspect of the invention, a method of operating a non-volatile memory system is disclosed, where the non-volatile memory system includes a non-volatile memory with a controller and a plurality of non-volatile memory dies, each of the plurality of non-volatile memory dies having a cache storage area and a main storage area. In the method, the controller receives data associated with host logical block addresses. The controller writes sequentially addressed data to the cache storage area of a first die of the plurality of non-volatile memory dies. After writing an amount of sequentially addressed data to the cache storage area of the first die equal to a block of the main storage area of the first die, the controller writes additional received sequentially addressed data to a cache storage area of a next die of the plurality of dies until a second amount of sequentially addressed data is written into the cache area of the next die equal to a block of the main storage area of the next die. The controller may write data from the cache storage area to the main storage area of the first die while writing additional received data to the cache area of the next die. Alternatively, the controller may write data from the cache storage area to the main storage area of the first die while writing data from the cache storage area to the main storage area of the next die.

[0007] Other features and advantages will become apparent upon review of the following drawings, detailed description and claims. Additionally, other embodiments are disclosed, and each of the embodiments can be used alone or together in combination. The embodiments will now be described with reference to the attached drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] FIG. **1** illustrates a host connected with a memory system having a multi-bank non-volatile memory containing multiple die.

[0009] FIG. **2** is an example block diagram of an example flash memory system controller for use in the multiple die non-volatile memory of FIG. **1**.

[0010] FIG. **3** is an example one flash memory bank suitable as one of the non-volatile memory banks illustrated in FIG. **1**.

[0011] FIG. **4** is a representative circuit diagram of a memory cell array that may be used in the memory bank of FIG. **3**.

[0012] FIG. **5** illustrates an example physical memory organization of the memory bank of FIG. **3**.

[0013] FIG. **6** shows an expanded view of a portion of the physical memory of FIG. **5**.

[0014] FIG. **7** illustrates a two die memory system.

[0015] FIG. **8** is a flow diagram of a method of achieving die parallelism in a memory system such as shown in FIG. **7**.

[0016] FIGS. **9**A-**9**D illustrates an embodiment of sequential write and parallel fold operations in a multiple die memory.

[0017] FIG. **10** illustrates a foggy-fine write process suitable for use in on-chip copy operations from binary metablocks to MLC metablocks.

[0018] FIG. **11** is an alternative embodiment of the sequential write and parallel folding operations of FIGS. **9**A-**9**D.

[0019] FIG. **12** is a second alternative embodiment of the sequential write and parallel folding operations of FIGS. **9**A-**9**D.

[0020] FIG. **13** is a memory system with a data flow suitable for incorporating the sequential write and parallel fold operations disclosed in FIGS. **9**A-**9**D.

DETAILED DESCRIPTION

[0021] A flash memory system suitable for use in implementing aspects of the invention is shown in FIGS. **1**-**7**. A host system **100** of FIG. **1** stores data into and retrieves data from a memory system **102**. The memory system may be flash memory embedded within the host, such as in the form of a solid state disk (SSD) drive installed in a personal computer. Alternatively, the memory system **102** may be in the form of a card that is removably connected to the host through mating parts **104** and **106** of a mechanical and electrical connector as illustrated in FIG. **1**. A flash memory configured for use as an internal or embedded SSD drive may look similar to the schematic of FIG. **1**, with the primary difference being the location of the memory system **102** internal to the host. SSD drives may be in the form of discrete modules that are drop-in replacements for rotating magnetic disk drives.

[0022] The host system **100** of FIG. **1** may be viewed as having two major parts, insofar as the memory system **102** is concerned, made up of a combination of circuitry and software. They are an applications portion **108** and a driver portion **110** that interfaces with the memory system **102**. In a PC, for example, the applications portion **110** can include a processor **112** running word processing, graphics, control or other popular application software, as well as the file system **114** for managing data on the host **100**. In a camera, cellular telephone or other host system that is primarily dedicated to performing a single set of functions, the applications portion **108** includes the software that operates the camera to take and store pictures, the cellular telephone to make and receive calls, and the like.

[0023] The memory system **102** of FIG. **1** may include non-volatile memory, such as flash memory **116**, and a system controller **118** that both interfaces with the host **100** to which the memory system **102** is connected for passing data back and forth and controls the memory **116**. The system controller **118** may convert between logical addresses of data used by the host **100** and physical addresses of the flash memory **116** during data programming and reading. The flash memory **116** may include any number of memory die **120** and two memory die are shown in FIG. **1** simply by way of illustration. Functionally, the system controller **118** may include a front end **122** that interfaces with the host system, controller logic **124** for coordinating operation of the memory **116**, flash management logic **126** for internal memory management operations such as garbage collection, and one or more flash interface modules (FIMs) **128** to provide a communication interface between the controller with the flash memory **116**.

[0024] The system controller **118** and may be implemented on a single integrated circuit chip, such as an application specific integrated circuit (ASIC) such as shown in FIG. **2**. The processor **206** of the system controller **118** may be configured as a multi-thread processor capable of communicating separately with each of the respective memory banks **120** via a memory interface **204** having I/O ports for each of the respective banks **120** in the flash memory **116**. The system controller **118** may include an internal clock **218**. The processor **206** communicates with an error correction code (ECC) module **214**, a RAM buffer **212**, a host interface **216**, and boot code ROM **210** via an internal data bus **202**.

[0025] Each die **120** in the flash memory **116** may contain an array of memory cells organized into multiple planes. One of FIG. **3** shows such planes **310** and **312** for simplicity but a greater number of planes, such as four or eight planes, may instead be used. Alternatively, the memory cell array of a memory bank may not be divided into planes. When so divided, however, each plane has its own column control circuits **314** and **316** that are operable independently of each other. The circuits **314** and **316** receive addresses of their respective memory cell array from the address portion **306** of the system bus **302**, and decode them to address a specific one or more of respective bit lines **318** and **320**. The word lines **322** are addressed through row control circuits **324** in response to addresses received on the address bus **19**. Source voltage control circuits **326** and **328** are also connected with the respective planes, as are p-well voltage control circuits **330** and **332**. If the bank **300** is in the form of a memory chip with a single array of memory cells, and if two or more such chips exist in the system, data are transferred into and out of the planes **310** and **312** through respective data input/output circuits **334** and **336** that are connected with the data portion **304** of the system bus **302**. The circuits **334** and **336** provide for both programming data into the memory cells and for reading data from the memory cells of their respective planes, through lines **338** and **340** connected to the planes through respective column control circuits **314** and **316**.

[0026] Although the processor **206** in the system controller **118** controls the operation of the memory chips in each bank **120** to program data, read data, erase and attend to various housekeeping matters, each memory chip also contains some controlling circuitry that executes commands from the controller **118** to perform such functions. Interface circuits **342** are connected to the control and status portion **308** of the system bus **302**. Commands from the controller **118** are provided to a state machine **344** that then provides specific control of other circuits in order to execute these commands. Control lines **346**-**354** connect the state machine **344** with these other circuits as shown in FIG. **3**. Status information from the state machine **344** is communicated over lines **356** to the interface **342** for transmission to the controller **118** over the bus portion **308**.

[0027] A NAND architecture of the memory cell arrays **310** and **312** is discussed below, although other architectures, such as NOR, can be used instead. An example NAND array is illustrated by the circuit diagram of FIG. **4**, which is a portion of the memory cell array **310** of the memory bank **300** of FIG. **3**. A large number of global bit lines are provided, only four such lines **402**-**408** being shown in FIG. **4** for simplicity of explanation. A number of series connected memory cell strings **410**-**424** are connected between one of these bit lines and a reference potential. Using the memory cell string **414** as representative, a plurality of charge storage memory cells **426**-**432** are connected in series with select transistors **434** and **436** at either end of the string. When the select transistors of a string are rendered conductive, the string is connected between its bit line and the reference potential. One memory cell within that string is then programmed or read at a time.

[0028] Word lines **438**-**444** of FIG. **4** individually extend across the charge storage element of one memory cell in each of a number of strings of memory cells, and gates **446** and **450** control the states of the select transistors at each end of the strings. The memory cell strings that share common word and control gate lines **438**-**450** are made to form a block **452** of memory cells that are erased together. This block of cells

contains the minimum number of cells that are physically erasable at one time. One row of memory cells, those along one of the word lines **438-444**, are programmed at a time. Typically, the rows of a NAND array are programmed in a prescribed order, in this case beginning with the row along the word line **444** closest to the end of the strings connected to ground or another common potential. The row of memory cells along the word line **442** is programmed next, and so on, throughout the block **452**. The row along the word line **438** is programmed last.

[0029] A second block **454** is similar, its strings of memory cells being connected to the same global bit lines as the strings in the first block **452** but having a different set of word and control gate lines. The word and control gate lines are driven to their proper operating voltages by the row control circuits **324**. If there is more than one plane in the system, such as planes **1** and **2** of FIG. **3**, one memory architecture uses common word lines extending between them. There can alternatively be more than two planes that share common word lines. In other memory architectures, the word lines of individual planes are separately driven.

[0030] The memory cells may be operated to store two levels of charge so that a single bit of data is stored in each cell. This is typically referred to as a binary or single level cell (SLC) memory. Alternatively, the memory cells may be operated to store more than two detectable levels of charge in each charge storage element or region, thereby to store more than one bit of data in each. This latter configuration is referred to as multi level cell (MLC) memory. Both types of memory cells may be used in a memory, for example binary flash memory may be used for caching data and MLC memory may be used for longer term storage. The charge storage elements of the memory cells are most commonly conductive floating gates but may alternatively be non-conductive dielectric charge trapping material.

[0031] FIG. **5** conceptually illustrates a multiple plane arrangement showing four planes **502-508** of memory cells. These planes **502-508** may be on a single die, on two die (two of the planes on each die) or on four separate die. Of course, other numbers of planes, such as 1, 2, 8, 16 or more may exist in each die of a system. The planes are individually divided into blocks of memory cells shown in FIG. **5** by rectangles, such as blocks **510**, **512**, **514** and **516**, located in respective planes **502-508**. There can be dozens or hundreds of blocks in each plane.

[0032] As mentioned above, a block of memory cells is the unit of erase, the smallest number of memory cells that are physically erasable together. For increased parallelism, however, the blocks are operated in larger metablock units. One block from each plane is logically linked together to form a metablock. The four blocks **510-516** are shown to form one metablock **518**. All of the cells within a metablock are typically erased together. The blocks used to form a metablock need not be restricted to the same relative locations within their respective planes, as is shown in a second metablock **520** made up of blocks **522-528**. Although it is usually preferable to extend the metablocks across all of the planes, for high system performance, the memory system can be operated with the ability to dynamically form metablocks of any or all of one, two or three blocks in different planes. This allows the size of the metablock to be more closely matched with the amount of data available for storage in one programming operation.

[0033] The individual blocks are in turn divided for operational purposes into pages of memory cells, as illustrated in FIG. **6**. The memory cells of each of the blocks **510-516**, for example, are each divided into eight pages P**0**-P**7**. Alternatively, there may be 32, 64 or more pages of memory cells within each block. The page is the unit of data programming and reading within a block, containing the minimum amount of data that are programmed or read at one time. In the NAND architecture of FIG. **3**, a page is formed of memory cells along a word line within a block. However, in order to increase the memory system operational parallelism, such pages within two or more blocks may be logically linked into metapages. A metapage **602** is illustrated in FIG. **6**, being formed of one physical page from each of the four blocks **510-516**. The metapage **602**, for example, includes the page P**2** in each of the four blocks but the pages of a metapage need not necessarily have the same relative position within each of the blocks. Within a die, a metapage is the maximum unit of programming.

[0034] FIG. **7** illustrates a memory **700** with two memory die **702**, **704** each having a cache storage area **706**, **708** and a main storage area **710**, **712** arrangement that may be used to implement block interleaving. As shown in FIG. **8**, block interleaving between multiple die such as die **702** and **704** in a flash memory **700** may be accomplished by writing sequentially addressed data received from a host to the cache storage area **706** of a first die **702** in flash memory (at **802**). When an amount of sequentially addressed data has been written to the cache storage area **706** of the first die that is equal to a size of a block of memory in a main storage area **710** of the first die, the controller of the flash memory writes a next group of received data from the host that is sequentially addressed to cache storage blocks in the next memory die **704** of the flash memory (at **804**, **806**). The writing of sequentially addressed data into the cache storage area **702** of the second die **704** continues until, as with the first die **702**, an amount of the cache storage area **708** of the second die **704** equal to a block size of a block in the main storage area **712** of the second die **704** is filled. In one embodiment, the sequentially addressed data in the cache storage area **706** of the first die **702** is written to main storage area **710** of the first die **702** concurrently with the writing of sequentially addressed data from the host to the cache storage area **708** of the second die **704** (at **806**). Accordingly, an entire block in the main storage area of each die is populated with sequentially addressed data each time data is copied from the cache storage area **706**, **708** to the main storage area **710**, **712**.

[0035] For simplicity of illustration, a controller is not shown in FIG. **7**, however the memory **700** of FIG. **7** may associated with a discrete controller such as controller **118** shown in FIG. **2**, may have one or more controllers integrated with the memory die on a single chip or may be configured to work with a combination of the two.

[0036] The steps illustrated in FIG. **8** assume a size of blocks in the cache storage area that is smaller than a size of blocks in the main storage area for each die of the flash memory. Also, the process described in FIG. **8** may be applied to any number of planes per die, where each plane in a die has cache and main storage areas, by writing to a number of blocks in the cache storage area of each respective plane in a die equal to the respective block size in the main storage area of each plane. Thus each of the cache storage areas for each plane of a die would be first written to prior to writing to the cache storage area of each plane in a next die. Also, the cache

storage area for each plane on the first die may be concurrently copied to the respective main storage area for the planes on the first die while host data is being written to cache storage area of each plane on the second die. Accordingly, the approach to block interleaving is scalable to any number of planes per die.

[0037] A more detailed example of one implementation of the process illustrated in FIG. **8** is shown in FIGS. **9**A-**9**D. In the example of FIGS. **9**A-**9**D a flash memory **900** is shown having two die, each with a single plane. In this example, the cache storage area is made up of binary (SLC) memory cells organized in blocks that contain 64 pages totaling 1 Megabyte of data and the main storage area in each die is made up of MLC memory cells (3 bits per cell in this example) organized in blocks that contain 3 Megabytes of data. Therefore, three binary blocks form an update group for each MLC block. Other size combinations of binary and MLC blocks may be implemented in different embodiments. Additionally, memory die with more than one plane per die may also be adapted in other embodiments.

[0038] FIG. **9**A illustrates the start of a sequential write operation in a flash memory **900** made up of two die, each die having a single plane arrangement of memory cells. For purposes of illustration, the sequence of FIGS. **9**A-**9**D is discussed with respect to data received from an external host, however the same process of sequentially writing to binary and then folding into MLC applies to data received from a source internal to the memory system (such as the binary cache). At the beginning of the sequential write operation, two empty MLC update blocks and 12 empty binary (SLC) update blocks are available, where half of each type of block are on each die. As data is written from the host, it is first written to the binary update blocks in Die **0**. Sequentially addressed host data is written to the binary update blocks in Die **0** until the entire update group of three binary blocks, also referred to as a virtual update block **902**, is completely filled. This state is shown in FIG. **9**B, where data for consecutive host LBA addresses 0-63, 64-127 and 127-191 have filled up three binary blocks. At this point, the controller of the flash memory **900** directs the next group of sequentially addressed host data to the first of the binary blocks in an update group in the binary cache of Die **1**. As illustrated in FIG. **9**C, concurrently with writing the next group of sequentially addressed host data to the binary cache in Die **1**, data in the completed binary block update group in Die **0** is copied to the associated MLC block in Die **0**. As described in greater detail below, the copy step from binary to MLC within a die is an on-chip copy process referred to herein as folding. More detail on folding is described below.

[0039] Referring to FIG. **9**D, after completing the internal folding operation to copy the data from the first binary update group to the first MLC update block in Die **0**, the controller will direct sequentially addressed data from the host to the second update group of binary blocks in Die **0** while completing the folding process from binary to MLC in the first update group in Die **1**. Also, after the folding process is completed on the first update group in Die **0**, the binary blocks in the first update group may be erased and returned to a pool of unallocated blocks. In one implementation, the number of pages written to binary blocks should equal the number of pages folded into MLC blocks, therefore the controller should write **3** metapages to binary update blocks for every folding step of copying each completed binary update block from an update group into a corresponding MLC block.

[0040] When gathering the necessary sequentially addressed host data in a virtual update block **902** of three binary blocks for folding into one 3 bit-per-cell MLC block, the logical group of addresses may be predetermined. Thus, the logical group to be associated with each virtual update block **902** grouped with an MLC block has a known starting address such as address **0** for the first group of update blocks of Die **0** in FIG. **9**B. Alternatively, rather than a fixed mapping where each virtual update block is assigned to a predetermined address range, a non-fixed mapping of addresses to the binary blocks may be implemented. In this non-fixed mapping implementation, the three binary blocks making up a virtual update block may be assigned sequentially addressed data, regardless of the starting address of the data, as long as sequentially addressed data is contained in the three binary blocks.

[0041] In another embodiment, the method may include filling up two or more virtual update blocks in each die before proceeding to write data to virtual update blocks of a next die. Thus, referring to the binary block triplets that define virtual update blocks **902** in FIGS. **9**A-**9**D, more than one virtual update block **902** may be filled up before writing to virtual update blocks **902** in the next die.

[0042] The folding process of performing an on-chip copy of data from three binary blocks to one MLC block in the example above may be accomplished using a "foggy-fine" technique to transfer the data from the 1 bit (2 states) per cell memory cells of the binary blocks to the higher density 3 bit (8 states) per cell arrangement of the MLC block. As is generally understood, binary flash memory cells can have two possible states (0 or 1) that represent 1 bit of information. For a given total operating voltage range available in a cell, the states of a binary cell can be represented by voltage levels above or below a reference voltage set halfway between the maximum and minimum voltage of the range. MLC memory cells, which may have the same physical configuration as a binary cell, need to partition their operating voltage range into finer increments to allow for discernable peaks at each desired state. For a three bit per cell MLC arrangement, a distribution of detectable eight states (000, 001, 010, 011, 100, 101, 110, 111) must be fit into the same voltage range that the binary memory cell can use for just two states.

[0043] In order to accurately copy data from the three binary update blocks to the associated MLC block after the update blocks have been written to, a binary to MLC write process such as the foggy-fine process described below is used. The process of writing from the three binary blocks to the one denser MLC block requires repeated writes to each wordline of the MLC block in order to achieve the charge resolution in the MLC cells necessary to distinguish between the greater possible states in the MLC cells. Referring to the foggy-fine folding process **1000** illustrated in FIG. **10** for the binary to 3 bit MLC example that has been discussed above, the memory controller writes to each wordline (WL) **1002** of the MLC three times. The first programming pass for the wordline **1002** is referred to as a lower resolution programming pass **1004**. The second programming pass for the wordline is referred to as the foggy programming pass **1006** and the third and final programming pass is the fine programming pass **1008**. Although each wordline **1002** will be written to three times, the order of the writing may be interleaved among the wordlines.

[0044] The folding process begins with the first wordline (WL **0**). In the table of FIG. **10**, the write order **1010** precedes

the pages written to **1012** under the type of write resolution pass for each wordline. For example, the sixth write operation is on the first wordline (WL **0**) and is a fine write pass for pages **0-2**. The order of writing is illustrated by the diagonal arrows that show how a different one of a lower, foggy or fine programming pass may be made sequentially on three sequential wordlines such that up to seven pages may need to be cached at any one time. This can be seen, for example in the concurrent writing to wordlines **2**, **1** and **0**, where the lower resolution write is made to page **6** on WL **2**, a foggy resolution write is made to pages **3-5** on WL **1** and a fine resolution write is made to pages **0-2** on WL **0**. The fine resolution write to WL **0** is the third and final pass for WL **0** so that the next round of writes begins on WL **3** and ends at WL **1**, and so on as indicated in FIG. **10**.

[0045] Although many more wordlines are contemplated and the number of pages per wordline may be greater than 3, for ease of illustration only five wordlines and 16 pages are shown in FIG. **10**. In the example of FIG. **10**, a maximum of 7 pages is included in a write pass and is maintained in binary memory rather than controller buffer RAM so that the VRAM size for the controller may be kept smaller and so that power loss issues can be avoided. The two die example of FIGS. **9**A-**9**D is provided for simplicity of illustration. In other embodiments, different numbers of die and/or numbers of planes per die may be utilized.

[0046] In the embodiment of FIG. **11**, where a two die memory Die **0**, Die **1** is shown having virtual update groups of three binary update blocks **1102** associated with each MLC block **1104** as in the example of FIG. **9**, another method for metablock interleaving may include first filling up a virtual update group of three update blocks in each of the two die (Die **0** and Die **1**), then folding **1106** the resulting six update blocks into their respective two MLC blocks simultaneously. While folding the sequentially addressed data in each die from binary to MLC blocks, the next sequentially addressed data may be concurrently written into the binary update blocks in each of the Die **0** and Die **1**. Once binary blocks D**1**-D**6** are filled the folding process will begin folding data from D**1-3** to MLC**0** and D**4-6** to MLC**1**. The host continues sequential writes to D**1** blocks D**7-9** and then D**10-12** that are interspersed with the folding operations of Virtual Update Groups **0** and **1**. As in the example of FIGS. **9**A-**9**D, the number of pages written to the binary update blocks in the example of FIG. **11** should be equal to the number of pages folded. Every fold operation **1106** results in three metapages of data being transferred from binary to MLC. Therefore it is necessary for the host to write three metapages to binary for every fold operation. In this example, two folds are performed in parallel so a total of six metapages need to be written to the binary buffers for every pair of fold operations.

[0047] In an embodiment illustrated in FIG. **12**, a hybrid combination of page level interleaving and metablock level interleaving is also contemplated. For example, in a 4-die memory **1200** (Die **0**-Die **3**) data may be interleaved across 4 MLC planes in two die to create a 6 MB metablock **1202**, and then the two pairs of die (Die **0**, Die **1** and Die **2**, Die **3**) may be block interleaved. This approach achieves a mix of sequential performance through 2-die interleaving but limits the metablock size to 6 Megabytes (MB) instead of 12 MB that would result from a full 4-die page level interleave. The order and timing of the host data write sequence to the page interleaved binary blocks of Die **0** and Die **1** and the subsequent folding operation from binary blocks to MLC blocks may be managed such that all 12 MB of binary update blocks in Virtual Update Group **0** and Virtual Update Group **1** are completed before simultaneously folding the binary blocks into their respective MLC blocks.

[0048] The above-described techniques of writing sequentially addressed data to binary blocks and then folding multiple binary blocks into MLC blocks in parallel may be implemented as part of an overall data flow in a memory system. The process and system discussed above is suitable for sequential writes to binary blocks and parallel folds into MLC blocks for data coming from multiple sources. For example, the data that triggers a fold operation may be received at the memory system from an external host. Alternatively, the trigger for a folding operation may be data previously received from the host and already in the memory system. The fold operation would then be triggered by a housekeeping function. Examples of housekeeping functions might be the eviction of data from a binary cache in the memory when enough non-sequential data runs have been collected to assemble a complete sequentially addresses virtual update block. Other housekeeping functions that might trigger a fold operation may be wear leveling operations that would result in a MLC-to-MLC block copy that would go through a virtual update block. Another example of a housekeeping operation capable of triggering a fold operation may be the process of cleaning up and closing binary blocks in a way that leads to the creation of a virtual update block. Thus the creation of virtual update blocks may be through sequentially written data received directly from an external source (e.g. host) or sources internal to the memory where previously received host data is moved about by housekeeping functions.

[0049] One example of a memory system **1300** with a data flow suitable for incorporating the above-described techniques is illustrated in FIG. **13**. FIG. **13** also illustrates the various sources from which data may be received to form virtual update blocks **1312** in binary memory and trigger fold operations in MLC memory. The memory system **1300** includes binary blocks **1302** and MLC blocks **1304**. A memory controller will route small fragments less than a metapage to the binary cache **1306**. The binary cache **1306** will be periodically compacted **1308** to clean up stale data. When the binary cache **1306** fills up, a logical group will be selected for eviction **1310**. The internal housekeeping function of binary cache eviction copies the logical group to a virtual update block **1312** and schedules it for folding **1314** to an MLC block **1316**.

[0050] Host data writes with a length greater than a metapage will be written directly to virtual update blocks. Several virtual update blocks **1312** can be open at a time though only one virtual update block per die will be folding at a time. When the system **1300** has exhausted the virtual update block resources, a virtual update block **1312** will be selected and scheduled for folding. Once data is successfully folded from binary blocks **1302** to MLC blocks **1304**, the virtual update block **1312** is retired and the binary block resources previously tied up with the virtual update block can be recovered. It should be noted that data that is written to binary cache may experience at least two binary writes before being folded to MLC blocks. It is written once to the binary cache **1306**, compacted **1308** zero or more times within the binary cache, and finally evicted **1310** and written to a binary virtual update block to be scheduled for folding. Long

sequential, aligned writes may experience only one binary write directly to a virtual update block before being folded to an MLC block.

[0051] FIG. **13**. also illustrates the process of adding an update of update block (UoUB) to a virtual update block. An update of update block contains updated information that is already in a binary update block, but that has been partially updated again prior to the original update block being folded into an MLC block. Thus the valid sequentially addressed data for a particular host address range may be found in two separate update blocks and an MLC block that is then copied into a virtual update block for folding into a new MLC block.

[0052] A special host resident binary cache zone **1318** can be reserved for File Allocation Table (FAT) areas or other logical regions that require frequent updates or faster write performance. MLC blocks will only be written by the operation of folding, or copying of the entire logical group from single, fully written, closed virtual update block. In other words, a logical group has to be fully consolidated to virtual update blocks **1312** before folding to an MLC metablock. The logical group needs to be consolidated together into the last Virtual Update block with an error correction code (ECC) check upon read from flash sources and ECC correction if necessary. In addition, some data, like FAT table data, can be routed to the host resident binary cache zone **1318** and stay resident there, thus not triggering any logical group consolidations.

[0053] A system and method has been described for achieving memory die parallelism by interleaving sequentially written data at a block level. Binary (or SLC) blocks may be programmed serially in sequential page order; an integral number of binary blocks are then folded into an MLC block in parallel. By interleaving sequentially written data across blocks, the same page order may be preserved, for any integral interleave factor. Accordingly, the number of memory die that can be operated in parallel does not increase the metablock size. Also, a combination of page interleaving and metablock interleaving may be implemented in a multiple die memory to obtain improved sequential write performance and reduce metablock size. The entirety of the following concurrently filed (Dec. 18, 2009), commonly owned U.S. patent application is incorporated herein by reference: "DATA TRANSFER FLOWS FOR ON-CHIP FOLDING" (attorney reference number 0084567.669US0 (SDD-1569/SDA-1537-US)).

[0054] It is intended that the foregoing detailed description be understood as an illustration of selected forms that the invention can take and not as a definition of the invention. It is only the following claims, including all equivalents, that are intended to define the scope of this invention. Also, some of the following claims may state that a component is operative to perform a certain function or configured for a certain task. It should be noted that these are not restrictive limitations. It should also be noted that the acts recited in the claims can be performed in any order and not necessarily in the order in which they are recited.

What is claimed is:

1. A non-volatile memory system comprising:

a non-volatile memory comprising a plurality of non-volatile memory dies, wherein each of the plurality of non-volatile memory dies has a cache storage area and a main storage area; and

a controller configured to:

receive data, the received data associated with host logical block addresses;

write sequentially addressed data to the cache storage area of a first die of the plurality of non-volatile memory dies; and

after writing an amount of sequentially addressed data to the cache storage area of the first die equal to a block of the main storage area of the first die, write additional received sequentially addressed data to a cache storage area of a next die of the plurality of dies until a second amount of sequentially addressed data is written into the cache area of the next die equal to a block of the main storage area of the next die.

2. The non-volatile memory system of claim **1**, wherein the controller is further configured to write data from the cache storage area to the main storage area of the first die while writing additional received data to the cache area of the next die.

3. The non-volatile memory system of claim **1**, wherein the controller is further configured to write data from the cache storage area to the main storage area of the first die while writing data from the cache storage area to the main storage area of the next die.

4. The non-volatile memory system of claim **1**, wherein the cache storage area of each of the plurality of dies comprises single-level cell (SLC) flash memory and the main storage area of each of the plurality of dies comprises multi-level cell (MLC) flash memory.

5. The non-volatile memory system of claim **4**, wherein for each of the plurality of non-volatile memory dies, a size of a block of MLC flash memory in the main storage area is greater than a size of a block of SLC memory in the cache storage area.

6. The non-volatile memory system of claim **5**, wherein the size of the block of MLC memory in the memory storage area is an integer multiple of the size of the block of SLC memory in the cache storage area.

7. The non-volatile memory system of claim **2**, wherein the data written from the cache storage area to the main storage area of the first die is maintained in a same sequentially addressed order such that the block of the main storage area contains sequentially addressed pages of the data written from the cache storage area.

8. The non-volatile memory system of claim **1**, wherein the plurality of non-volatile memory dies each comprise at least two planes, wherein each of the plurality of non-volatile memory dies has a cache storage area and a main storage area for each of the at least two planes; and

wherein the controller is further configured to:

write sequentially addressed data to the cache storage area of a first plane of a first die of the plurality of non-volatile memory dies;

after writing an amount of sequentially addressed data to the cache storage area of the first plane of the first die equal to a block of the main storage area of the first plane, write additional received sequentially addressed data to a cache storage area of a next plane of the first die until a second amount of sequentially addressed data is written into the cache area of the next plane equal to a block of the main storage area of the next plane; and

after writing to each of the at least two planes of the first die, write respective additional sequentially addressed data received from the host to a cache stor-

age area, and then from the cache storage area to a main storage area, for each respective plane of the next die.

9. A method of operating a non-volatile memory system comprising:

in a non-volatile memory system having a controller and a non-volatile memory comprising a plurality of non-volatile memory dies, wherein each of the plurality of non-volatile memory dies has a cache storage area and a main storage area:

the controller receiving data associated with host logical block addresses;

the controller writing sequentially addressed data to the cache storage area of a first die of the plurality of non-volatile memory dies; and

after writing an amount of sequentially addressed data to the cache storage area of the first die equal to a block of the main storage area of the first die, the controller writing additional received sequentially addressed data to a cache storage area of a next die of the plurality of dies until a second amount of sequentially addressed data is written into the cache area of the next die equal to a block of the main storage area of the next die.

10. The method of claim **9**, further comprising the controller writing data from the cache storage area to the main storage area of the first die while writing additional received data to the cache area of the next die.

11. The method of claim **10**, wherein receiving data comprises receiving data from an external host.

12. The method of claim **10**, wherein receiving data comprises receiving data from a source internal to the memory system.

13. The method of claim **9**, further comprising the controller writing data from the cache storage area to the main storage area of the first die while writing data from the cache storage area to the main storage area of the next die.

14. The method of claim **9**, wherein writing sequentially addressed data to the cache storage area comprises writing sequentially addressed data to single-level cell (SLC) flash memory.

15. The method of claim **10**, wherein writing data from the cache storage area to the main storage area comprises writing data from single-level cell (SLC) flash memory to multi-level cell (MLC) flash memory.

16. The method of claim **15**, wherein for each of the plurality of non-volatile memory dies, a size of a block of MLC flash memory in the main storage area is greater than a size of a block of SLC memory in the cache storage area.

17. The method of claim **16**, wherein the size of the block of MLC memory in the memory storage area is an integer multiple of the size of the block of SLC memory in the cache storage area.

18. The method of claim **10**, further comprising writing data from the cache storage area to the main storage area of the first die in a same sequentially addressed order such that the block of the main storage area contains sequentially addressed pages of the data written from the cache storage area.

19. The method of claim **9**, wherein the plurality of non-volatile memory dies each comprise at least two planes, and each of the plurality of non-volatile memory dies has a cache storage area and a main storage area for each of the at least two planes; and

the method further comprising:

the controller writing sequentially addressed data to the cache storage area of a first plane of a first die of the plurality of non-volatile memory dies;

after writing an amount of sequentially addressed data to the cache storage area of the first plane of the first die equal to a block of the main storage area of the first plane, the controller writing additional received sequentially addressed data to a cache storage area of a next plane of the first die until a second amount of sequentially addressed data is written into the cache area of the next plane equal to a block of the main storage area of the next plane; and

after writing to each of the at least two planes of the first die, the controller, for each plane in a next die, writing respective additional sequentially addressed data received from the host to a cache storage area, and then from the cache storage area to a main storage area, for each respective plane of the next die.

20. A non-volatile memory system comprising:

a non-volatile memory comprising a plurality of non-volatile memory dies, wherein each of the plurality of non-volatile memory dies has a cache storage area and a main storage area; and

a controller configured to:

receive data, the received data associated with host logical block addresses;

write sequentially addressed data to the cache storage area of a first die of the plurality of non-volatile memory dies; and

after writing an amount of sequentially addressed data to the cache storage area of the first die equal to a plurality of blocks of the main storage area of the first die, write additional received sequentially addressed data to a cache storage area of a next die of the plurality of dies until a second amount of sequentially addressed data is written into the cache area of the next die equal to a plurality of blocks of the main storage area of the next die.

21. The non-volatile memory system of claim **20**, wherein the controller is further configured to write data from the cache storage area to the main storage area of the first die while writing additional received data to the cache area of the next die.

22. The non-volatile memory system of claim **20**, wherein the controller is further configured to write data from the cache storage area to the main storage area of the first die while writing data from the cache storage area to the main storage area of the next die.

23. The non-volatile memory system of claim **21**, wherein the controller is configured to receive data from a host external to the memory system, and wherein the controller is configured to write sequentially addressed data to the cache storage areas and write data from the cache storage area to the main storage area of the first die in response to receiving data from the host.

24. The non-volatile memory system of claim **21**, wherein the controller is configured to receive data from a source internal to the memory system as a result of a housekeeping operation, and wherein the controller is configured to write sequentially addressed data to the cache storage areas and write data from the cache storage area to the main storage area of the first die in response to receiving data from the source internal to the memory system.

* * * * *