

(12) UK Patent

(19) GB

(11) 2563713

(45) Date of B Publication

(13) B

15.01.2020

(54) Title of the Invention: NAND flash storage device with NAND buffer

(51) INT CL: **G06F 3/06** (2006.01) **G06F 12/02** (2006.01)

(21) Application No: 1804693.8

(22) Date of Filing: 23.03.2018

(30) Priority Data:
(31) 62524016 (32) 23.06.2017 (33) US
(31) 15816447 (32) 17.11.2017 (33) US

(43) Date of A Publication 26.12.2018

(56) Documents Cited:
US 8301832 B US 20170068456 A
US 20140304452 A US 20110258362 A

(58) Field of Search:
As for published application 2563713 A viz:
INT CL **G06F**
Other: **EPODOC, WPI**
updated as appropriate

Additional Fields
Other: **None**

(72) Inventor(s):
Monish Shah

(73) Proprietor(s):
Google LLC
1600 Amphitheatre Parkway, Mountain View 94043,
California, United States of America

(74) Agent and/or Address for Service:
Marks & Clerk LLP
62/68 Hills Road, CAMBRIDGE, CB2 1LA,
United Kingdom

GB

2563713

B

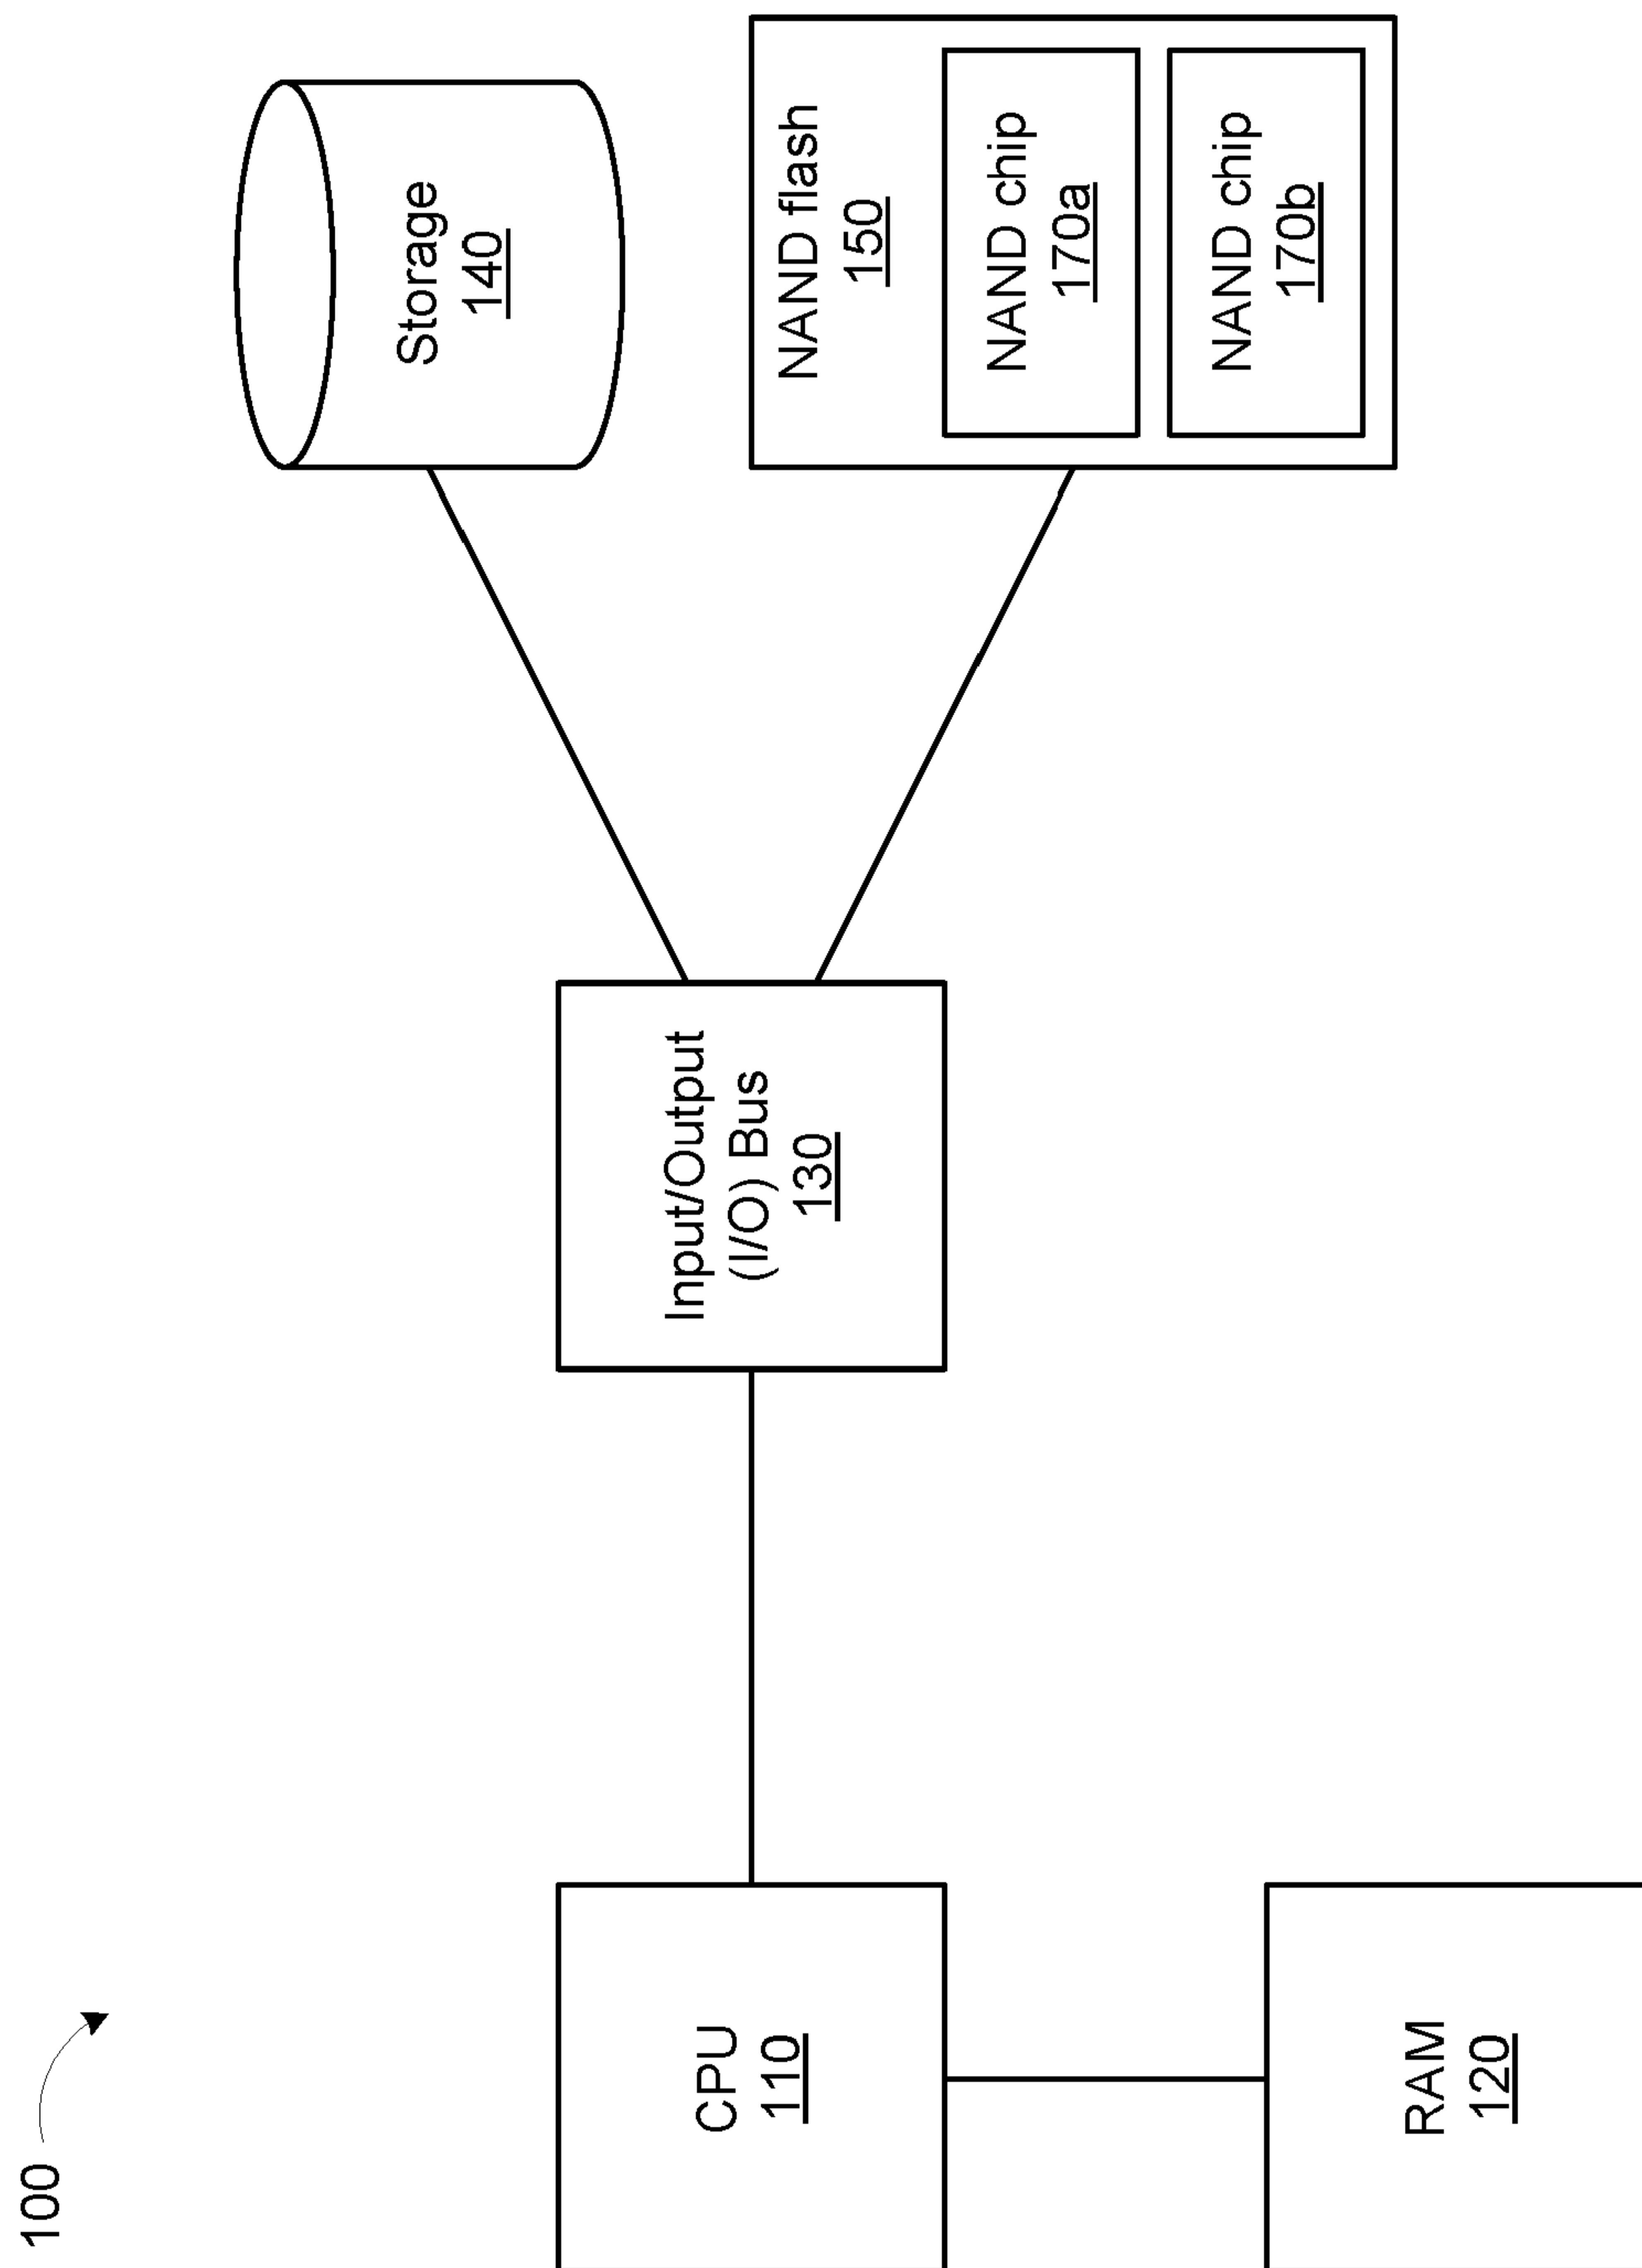


FIG. 1

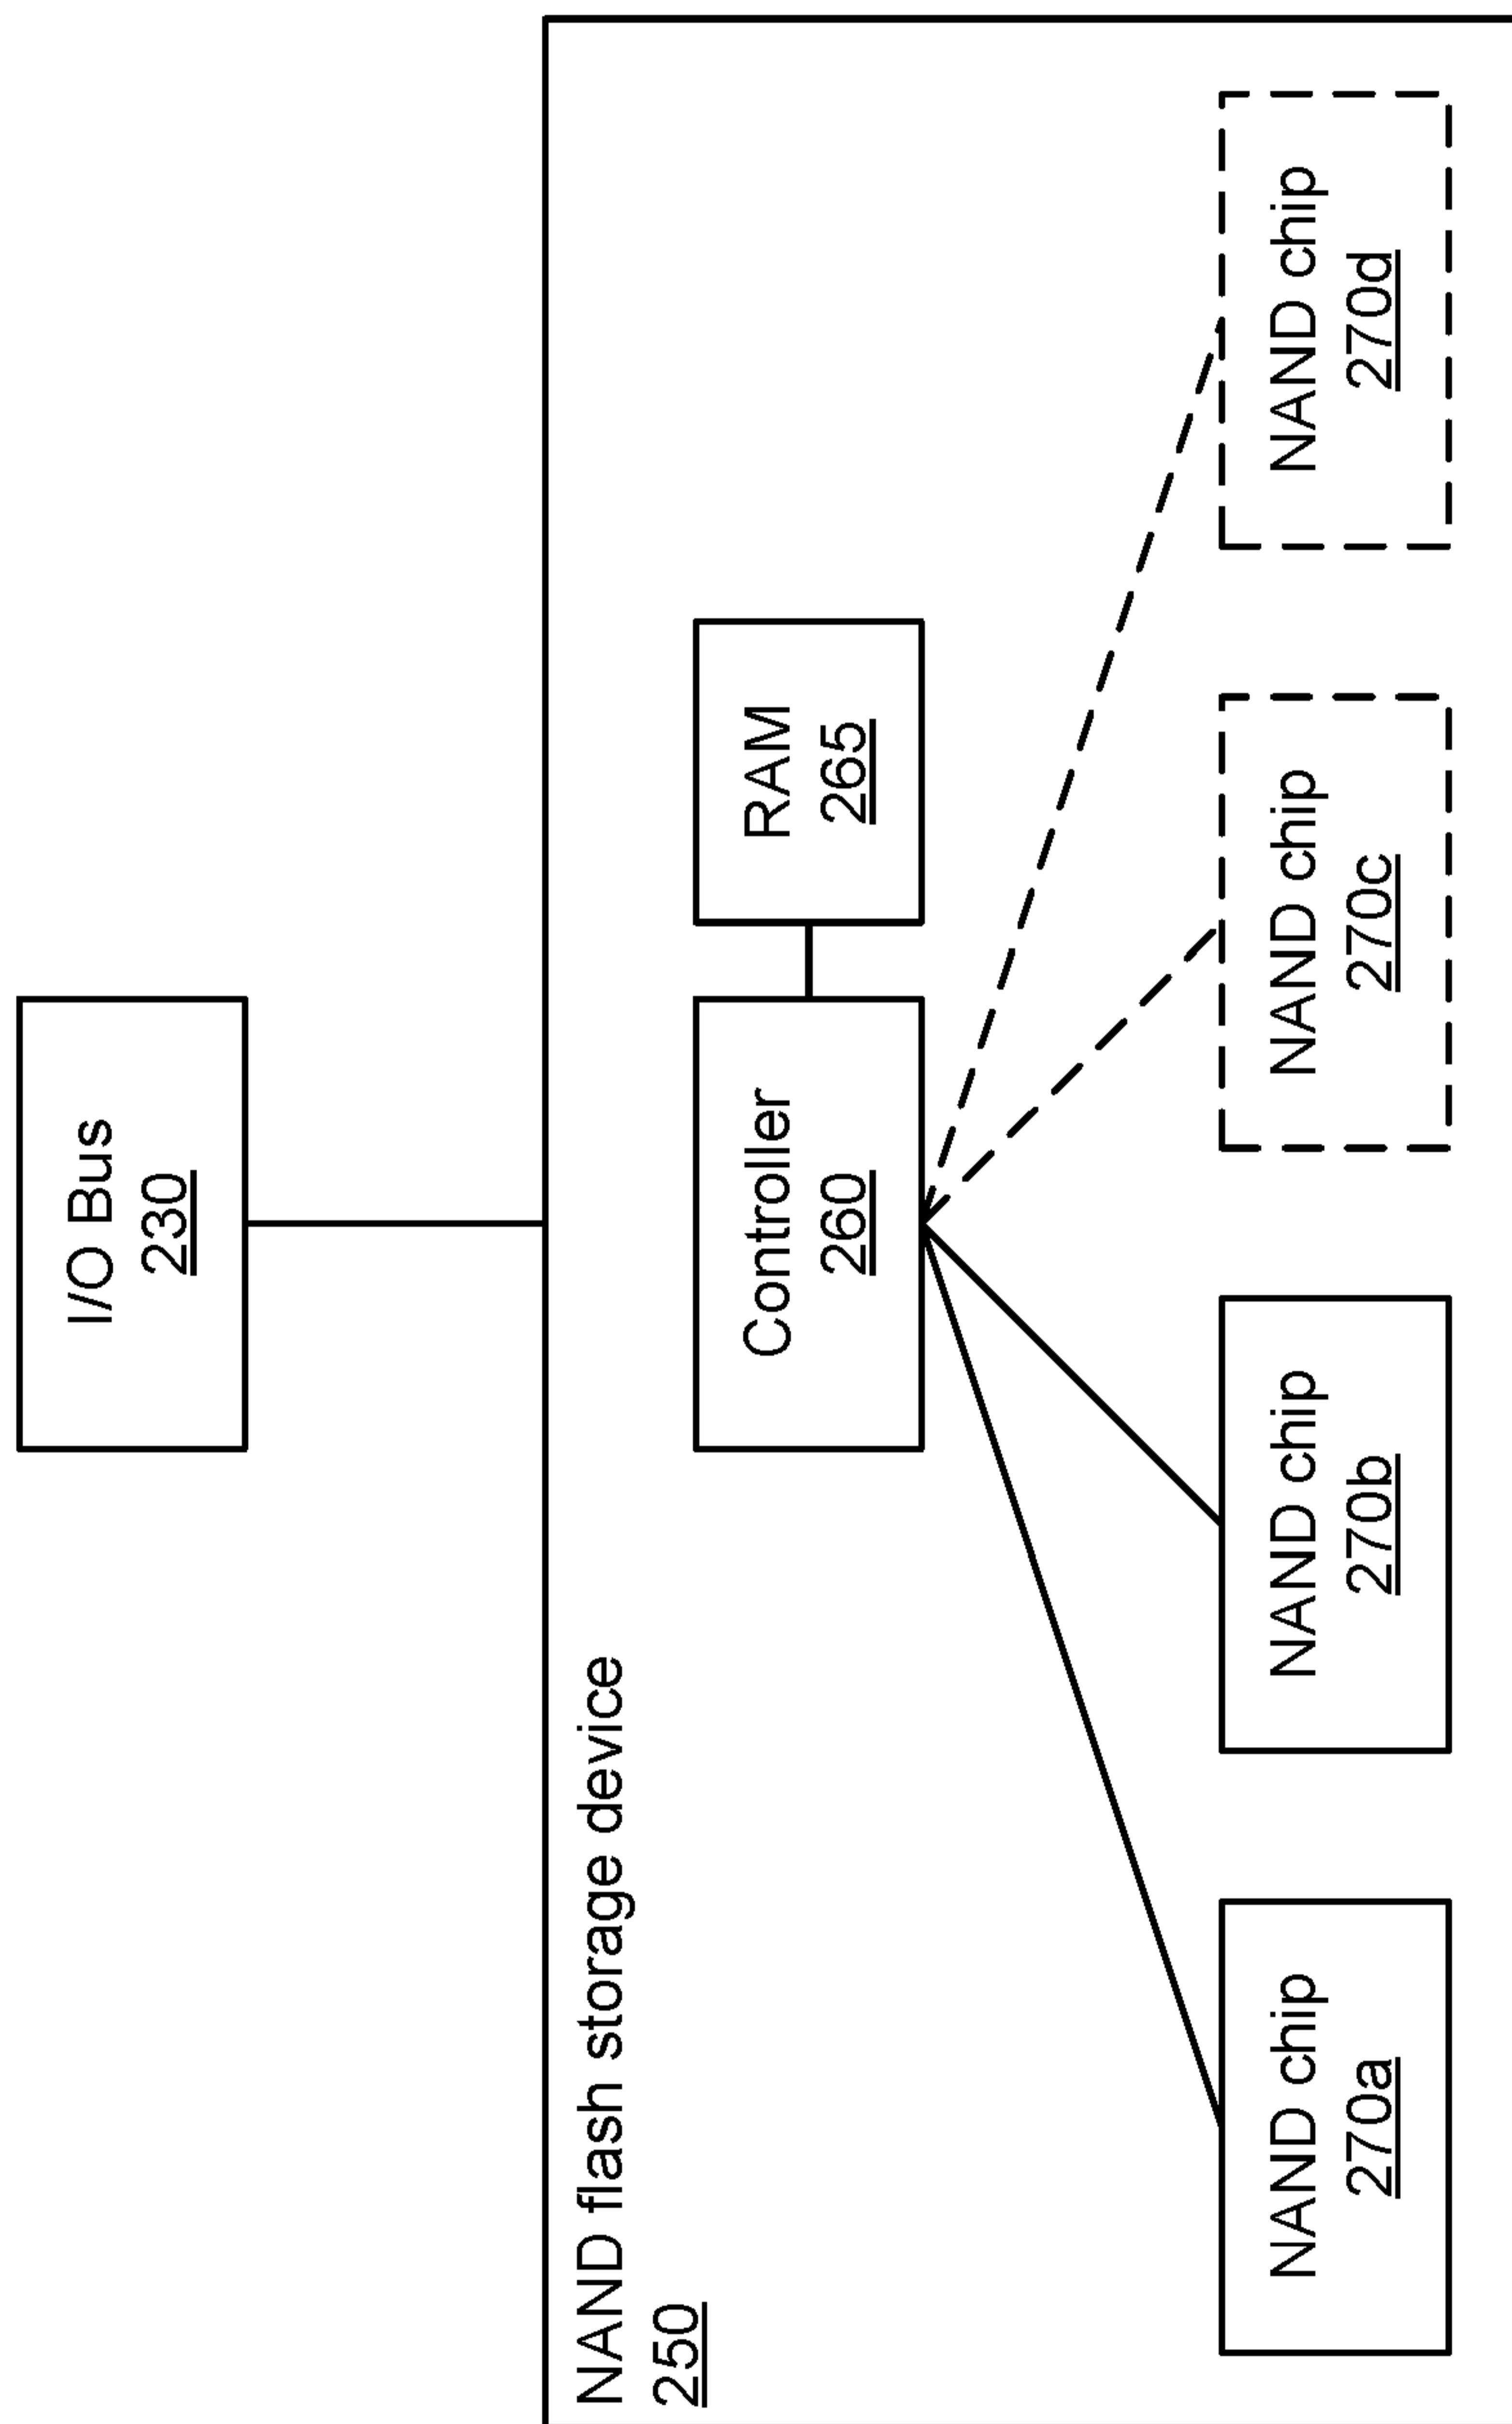


FIG. 2

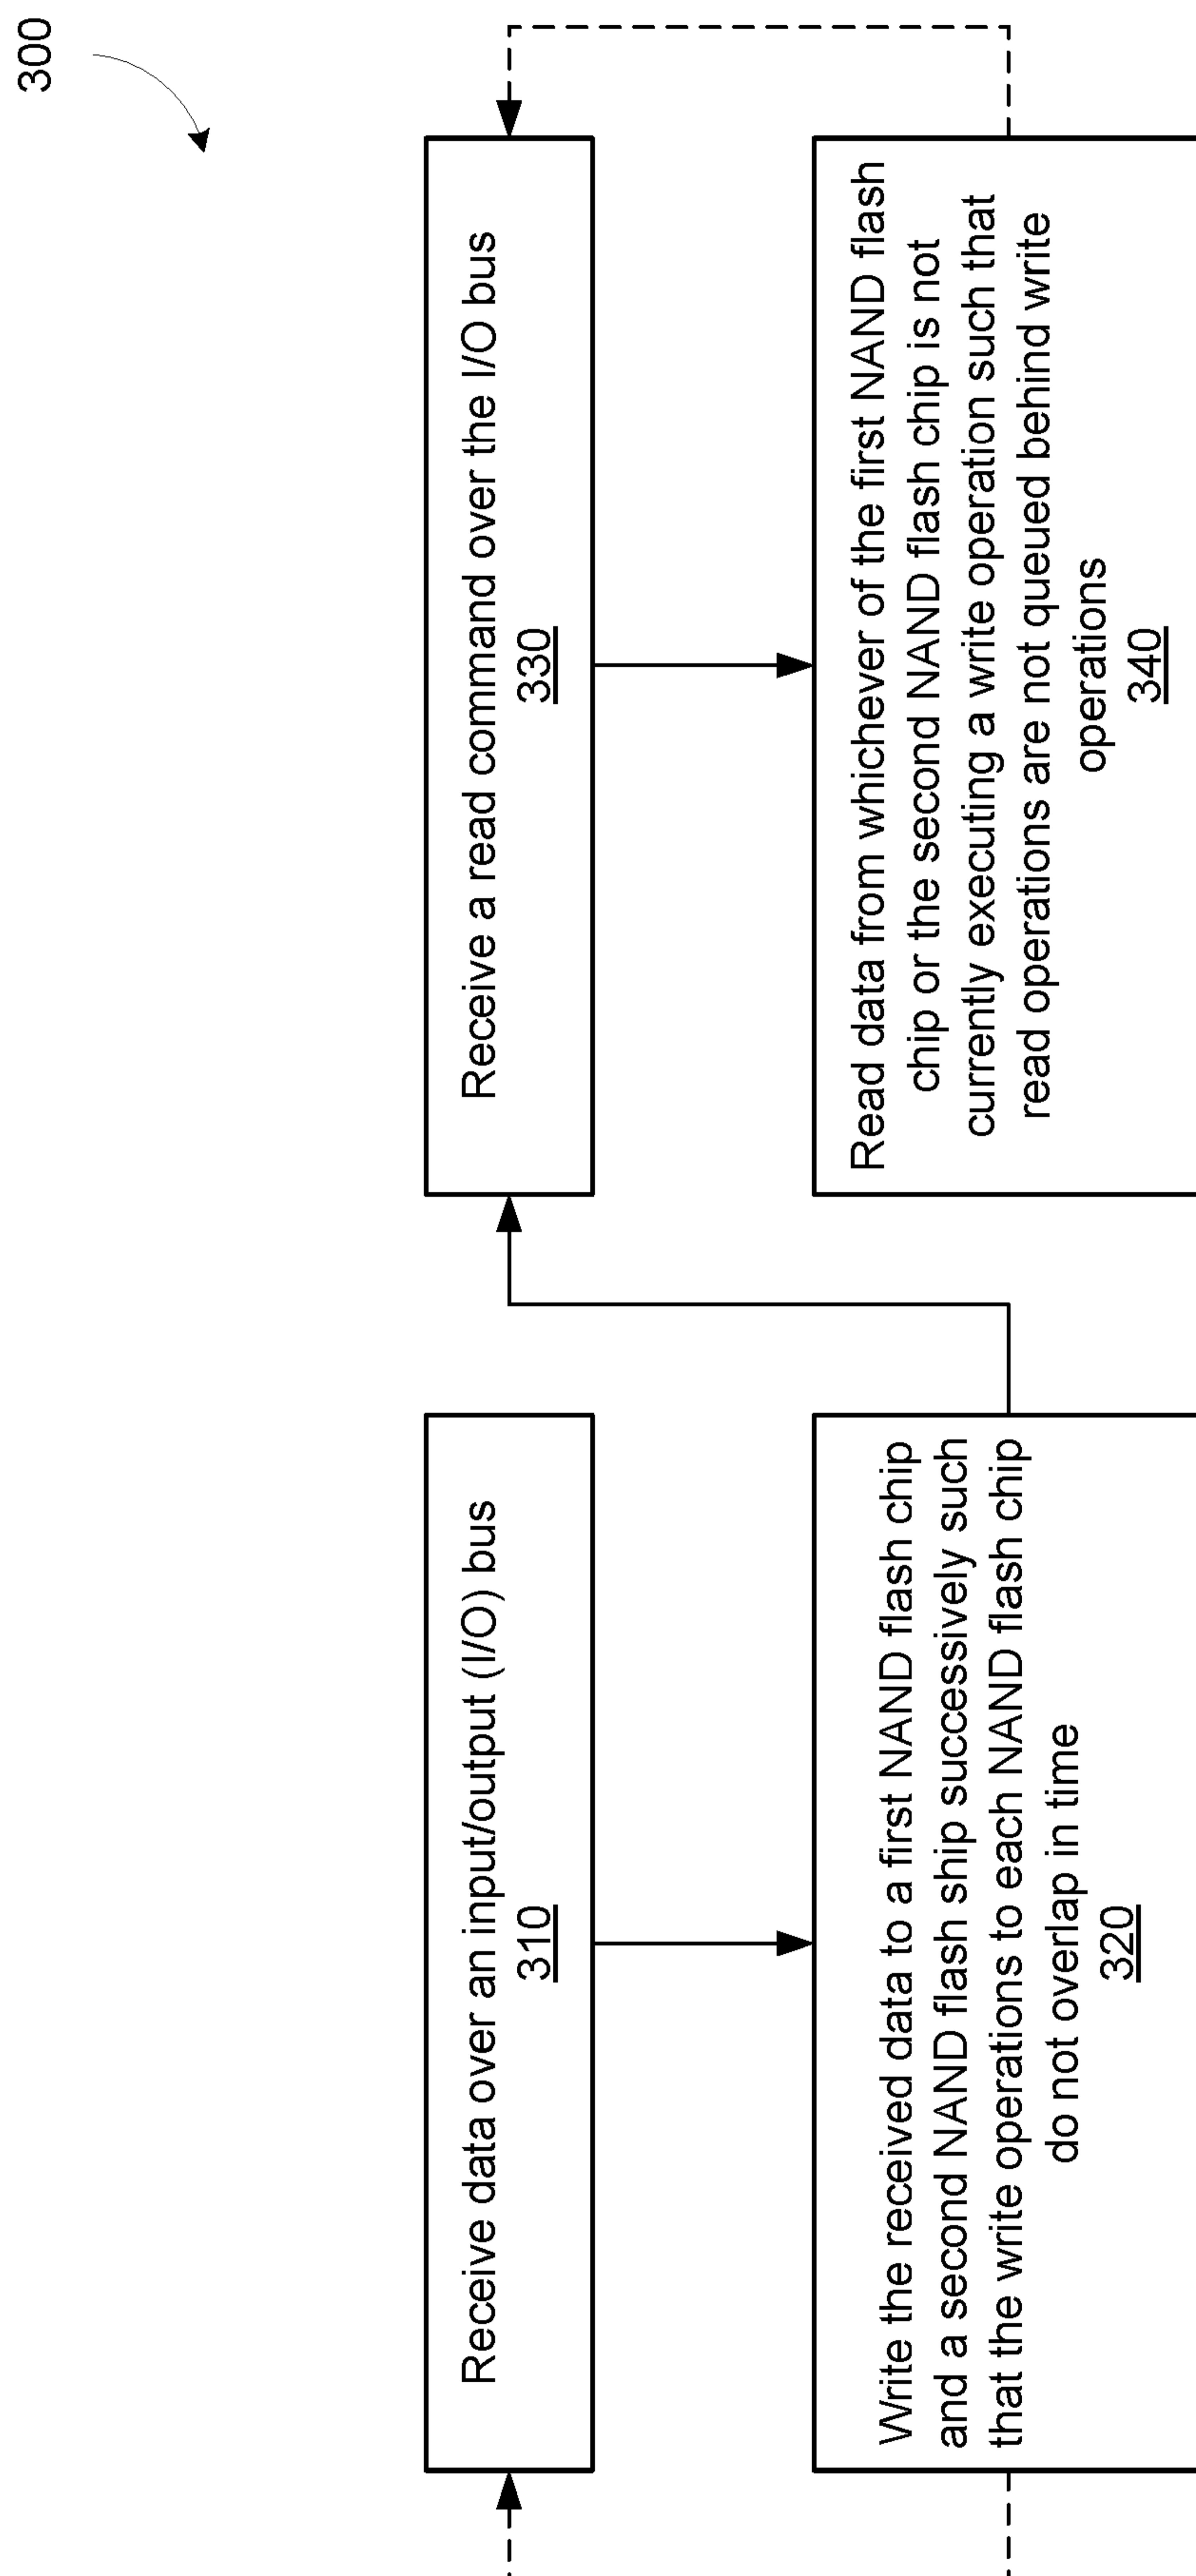


FIG. 3

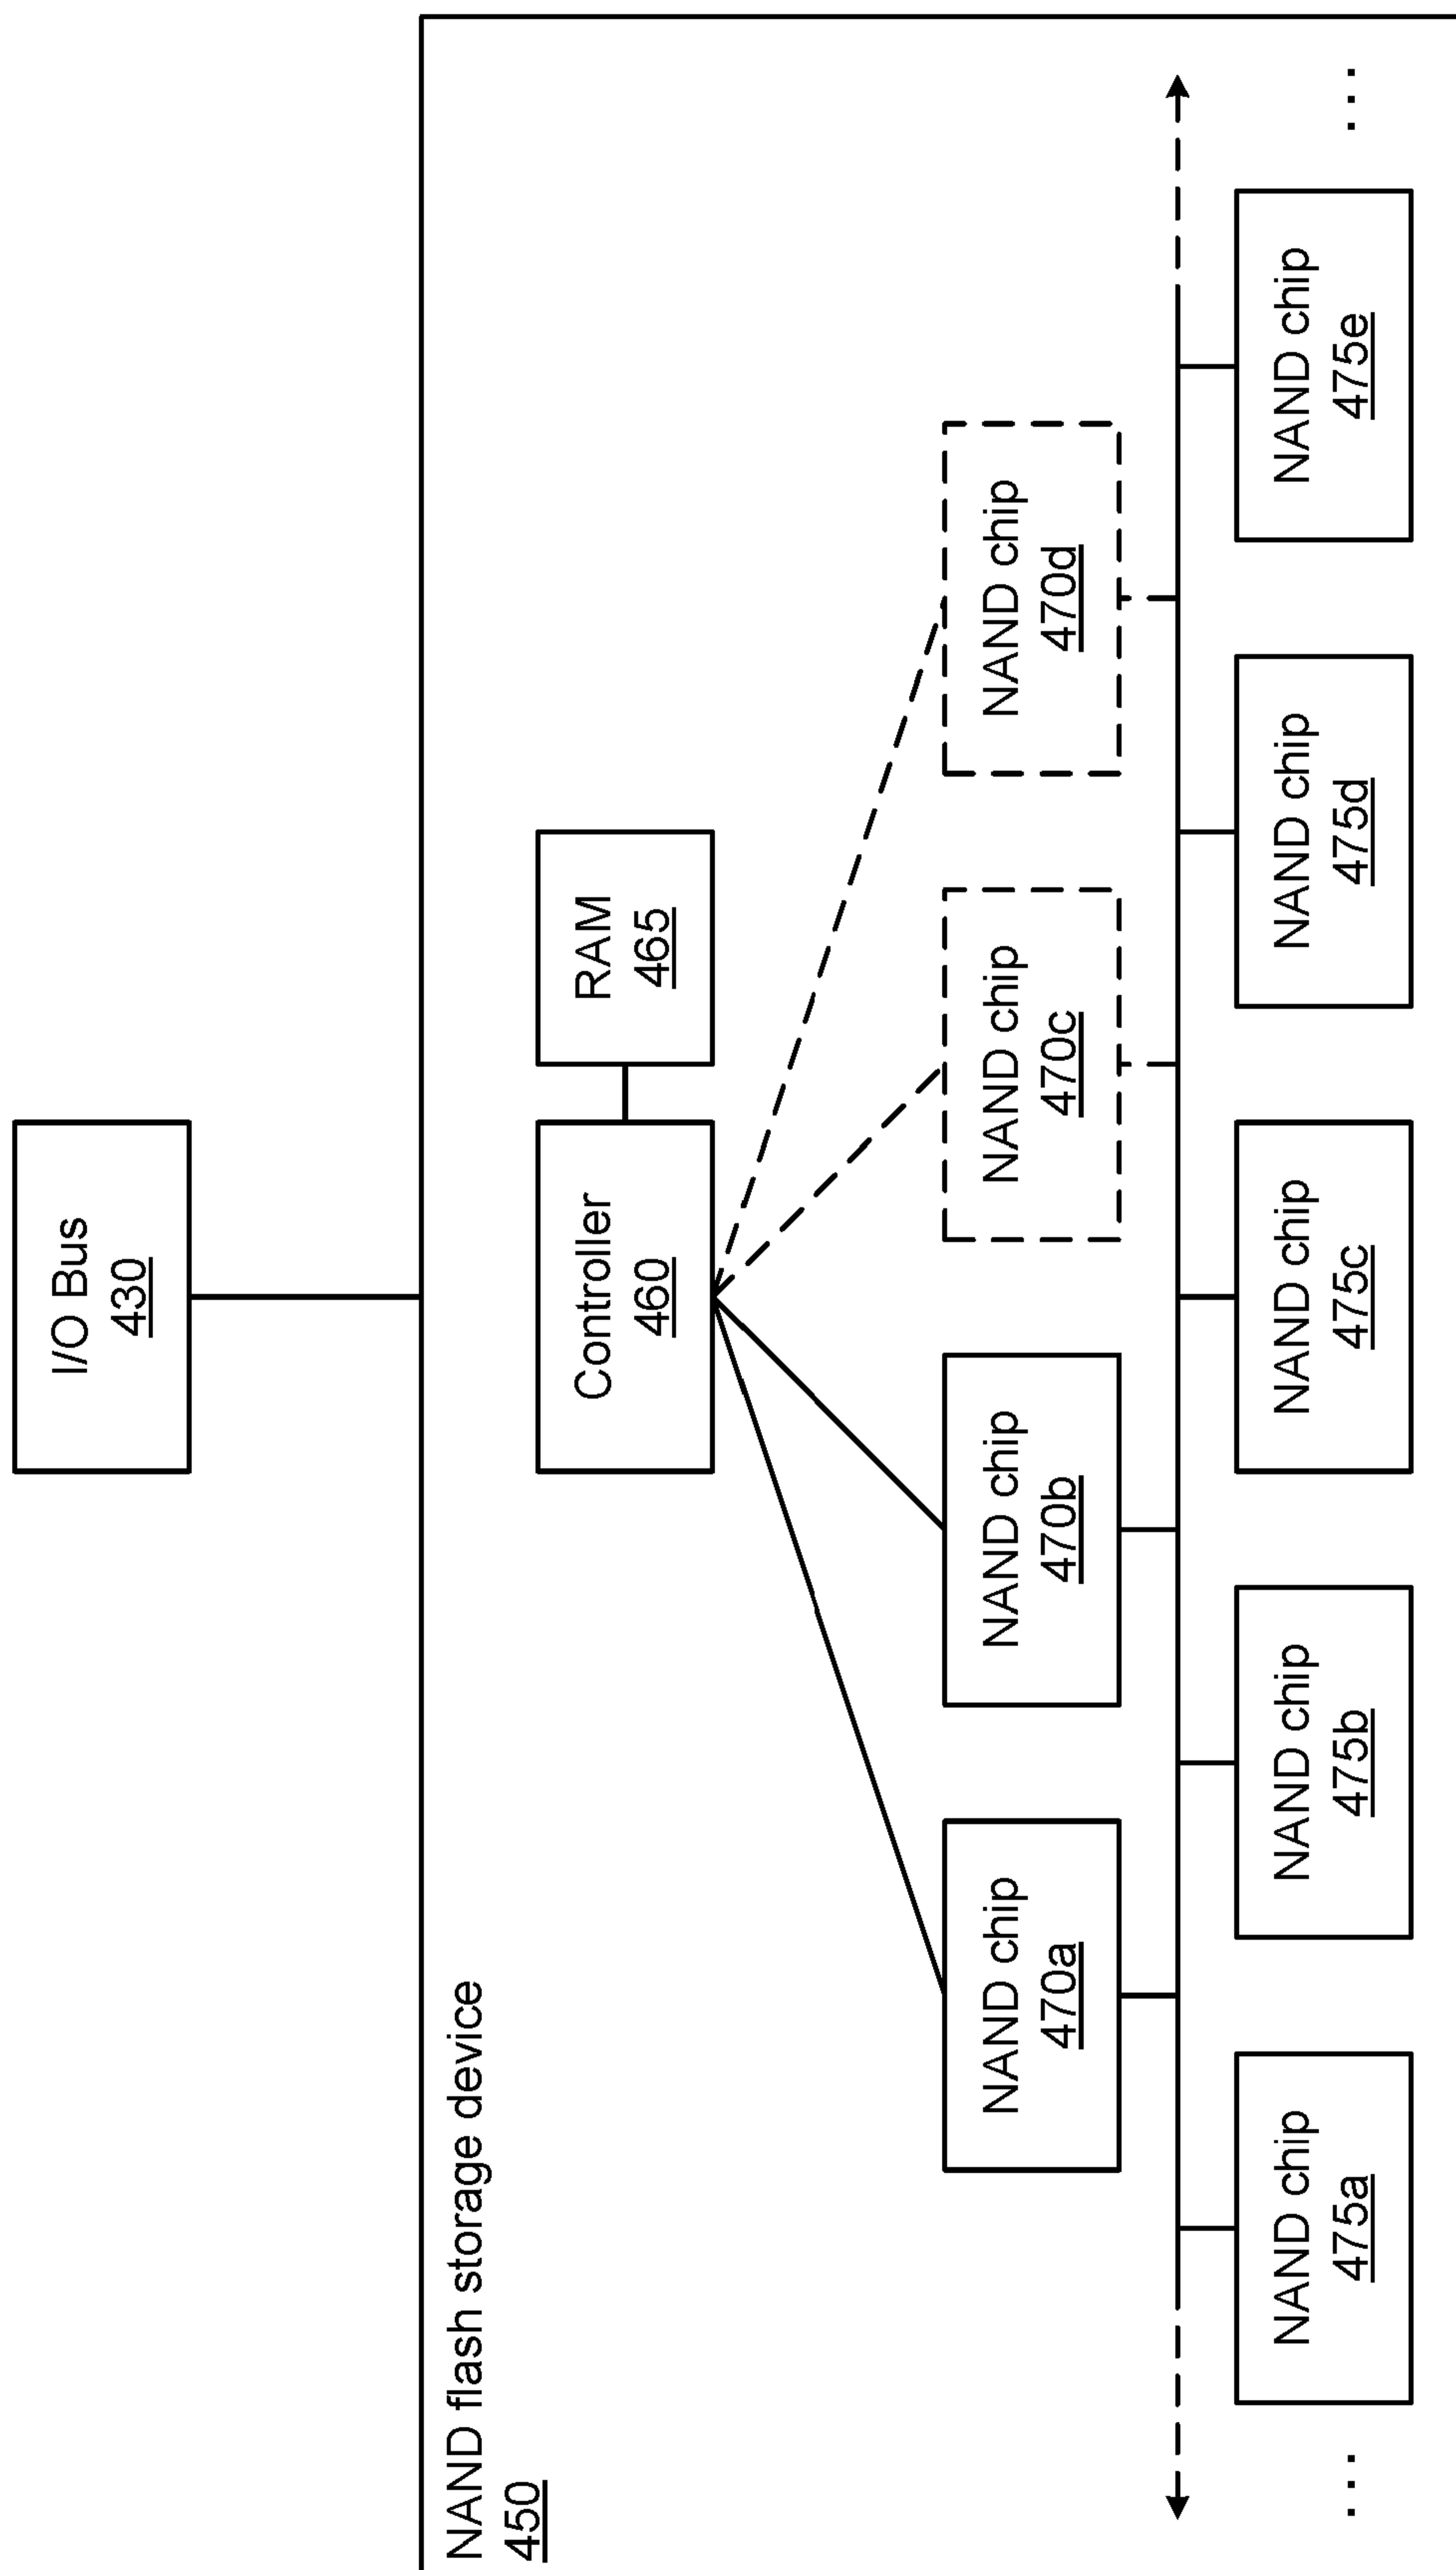


FIG. 4

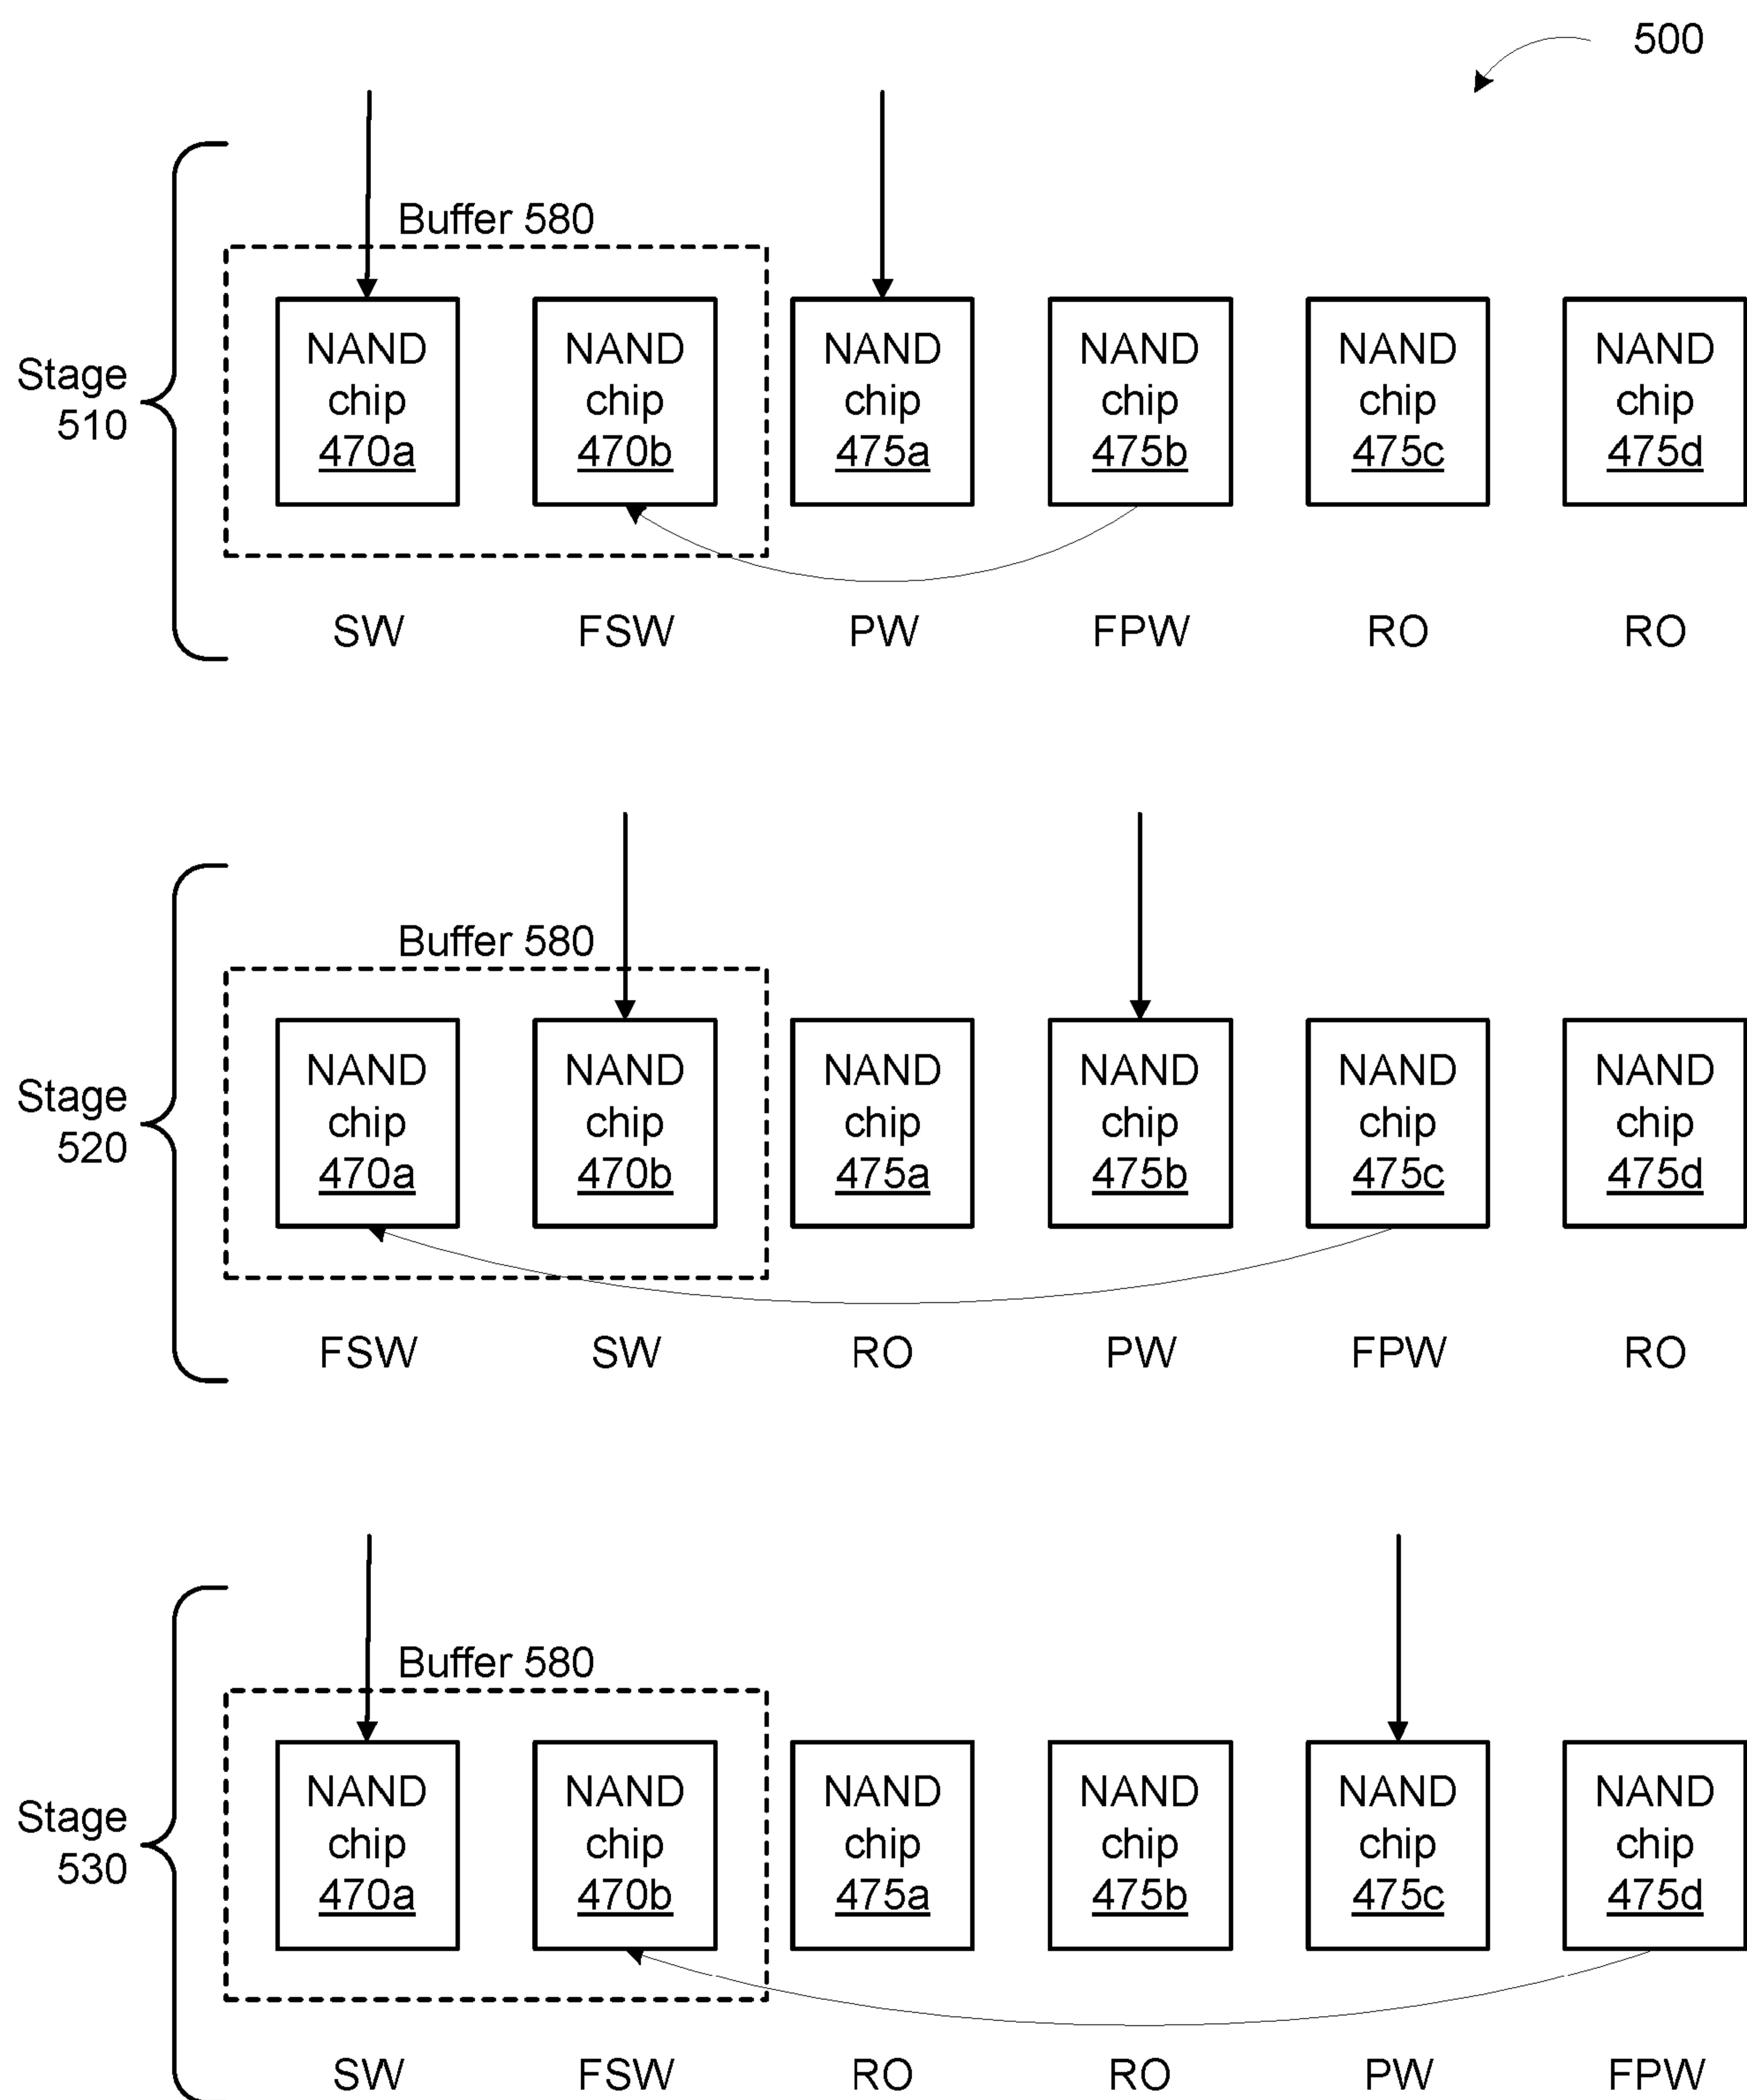


FIG. 5

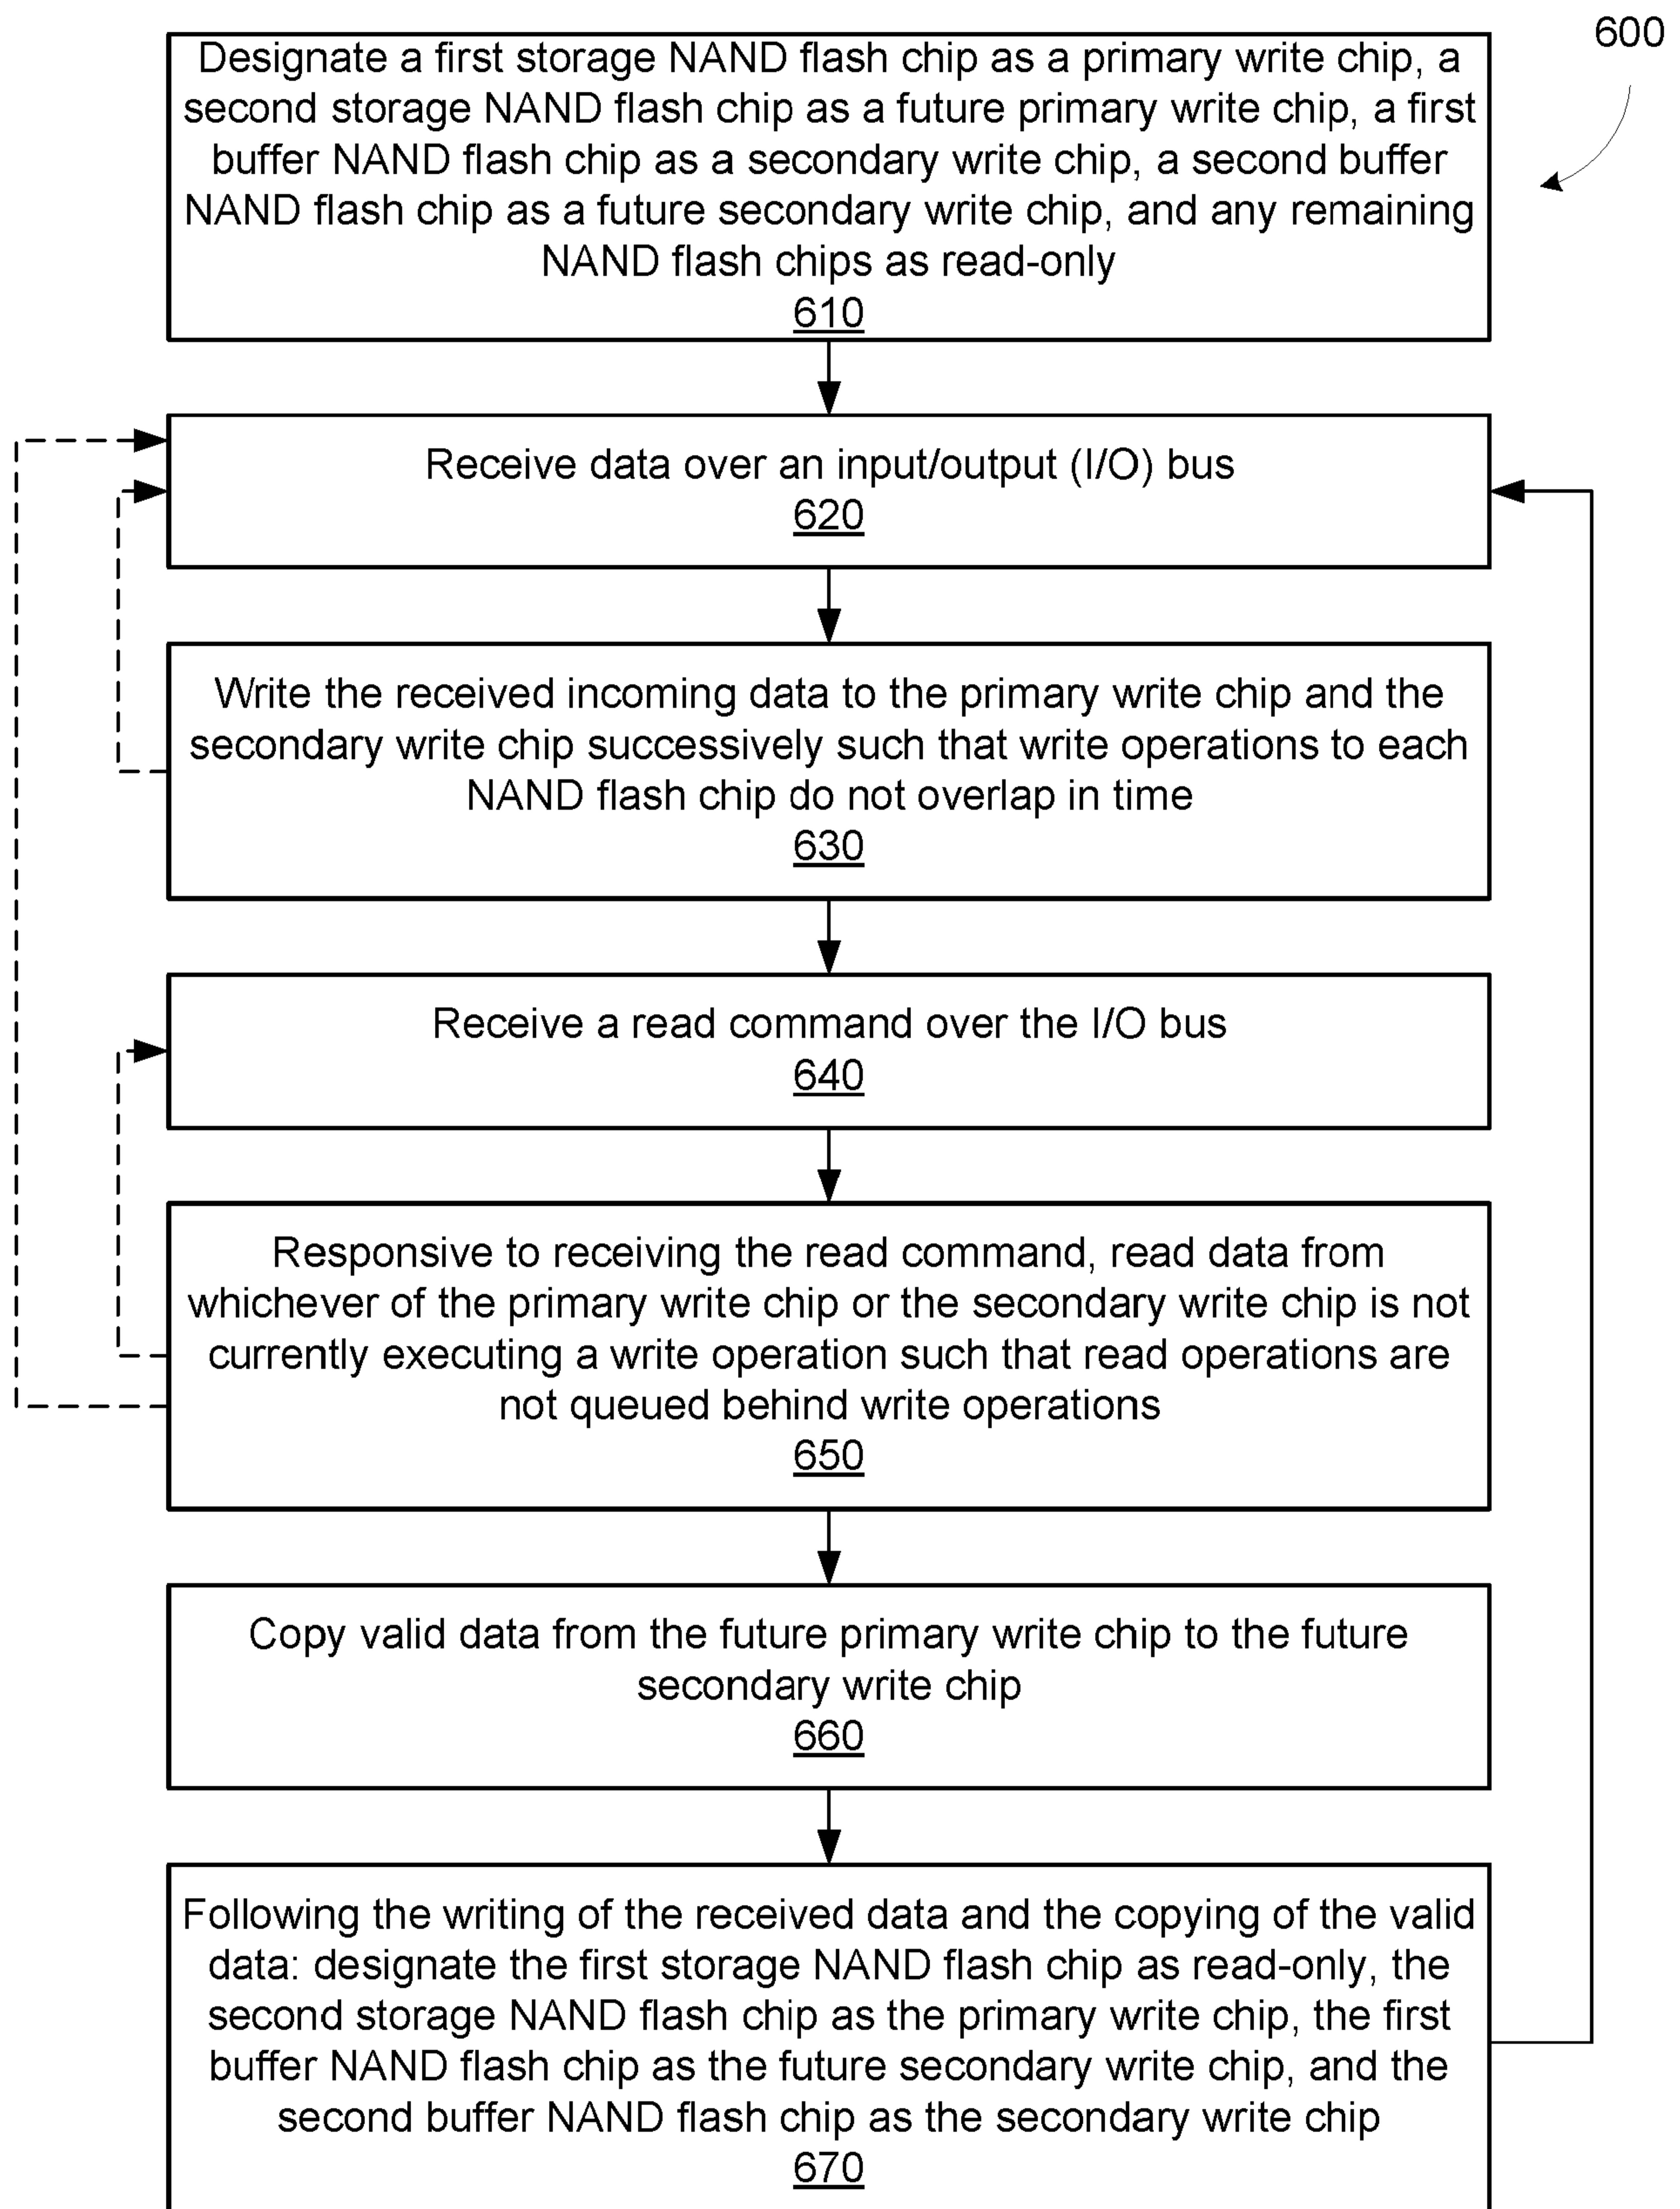


FIG. 6

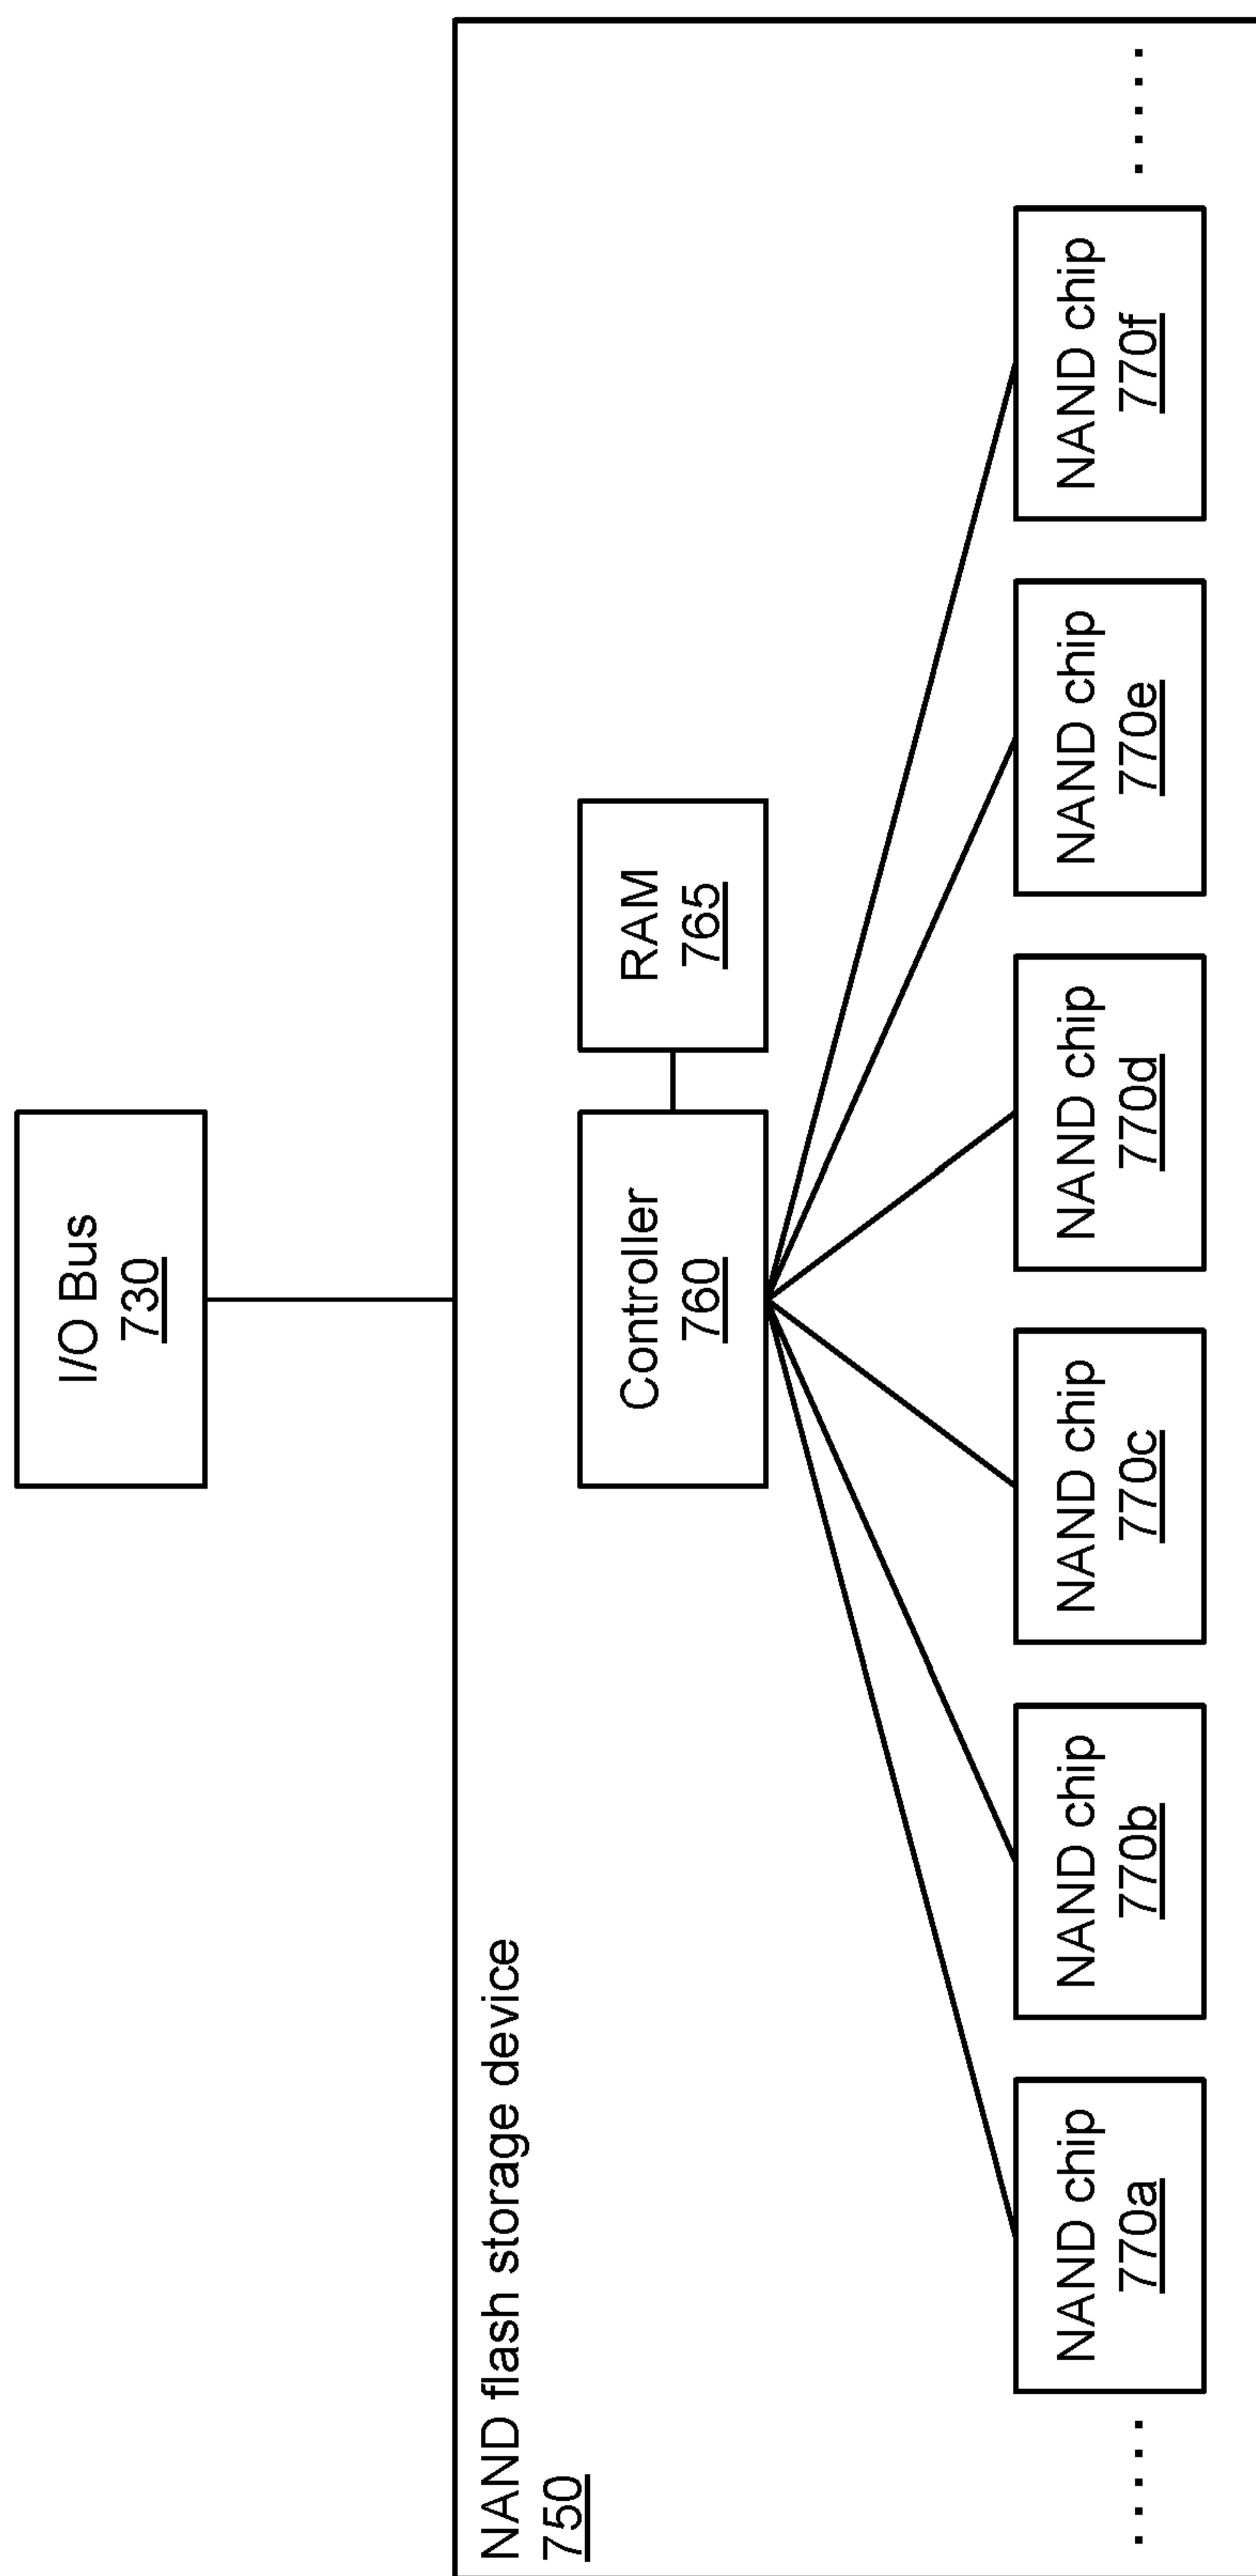


FIG. 7

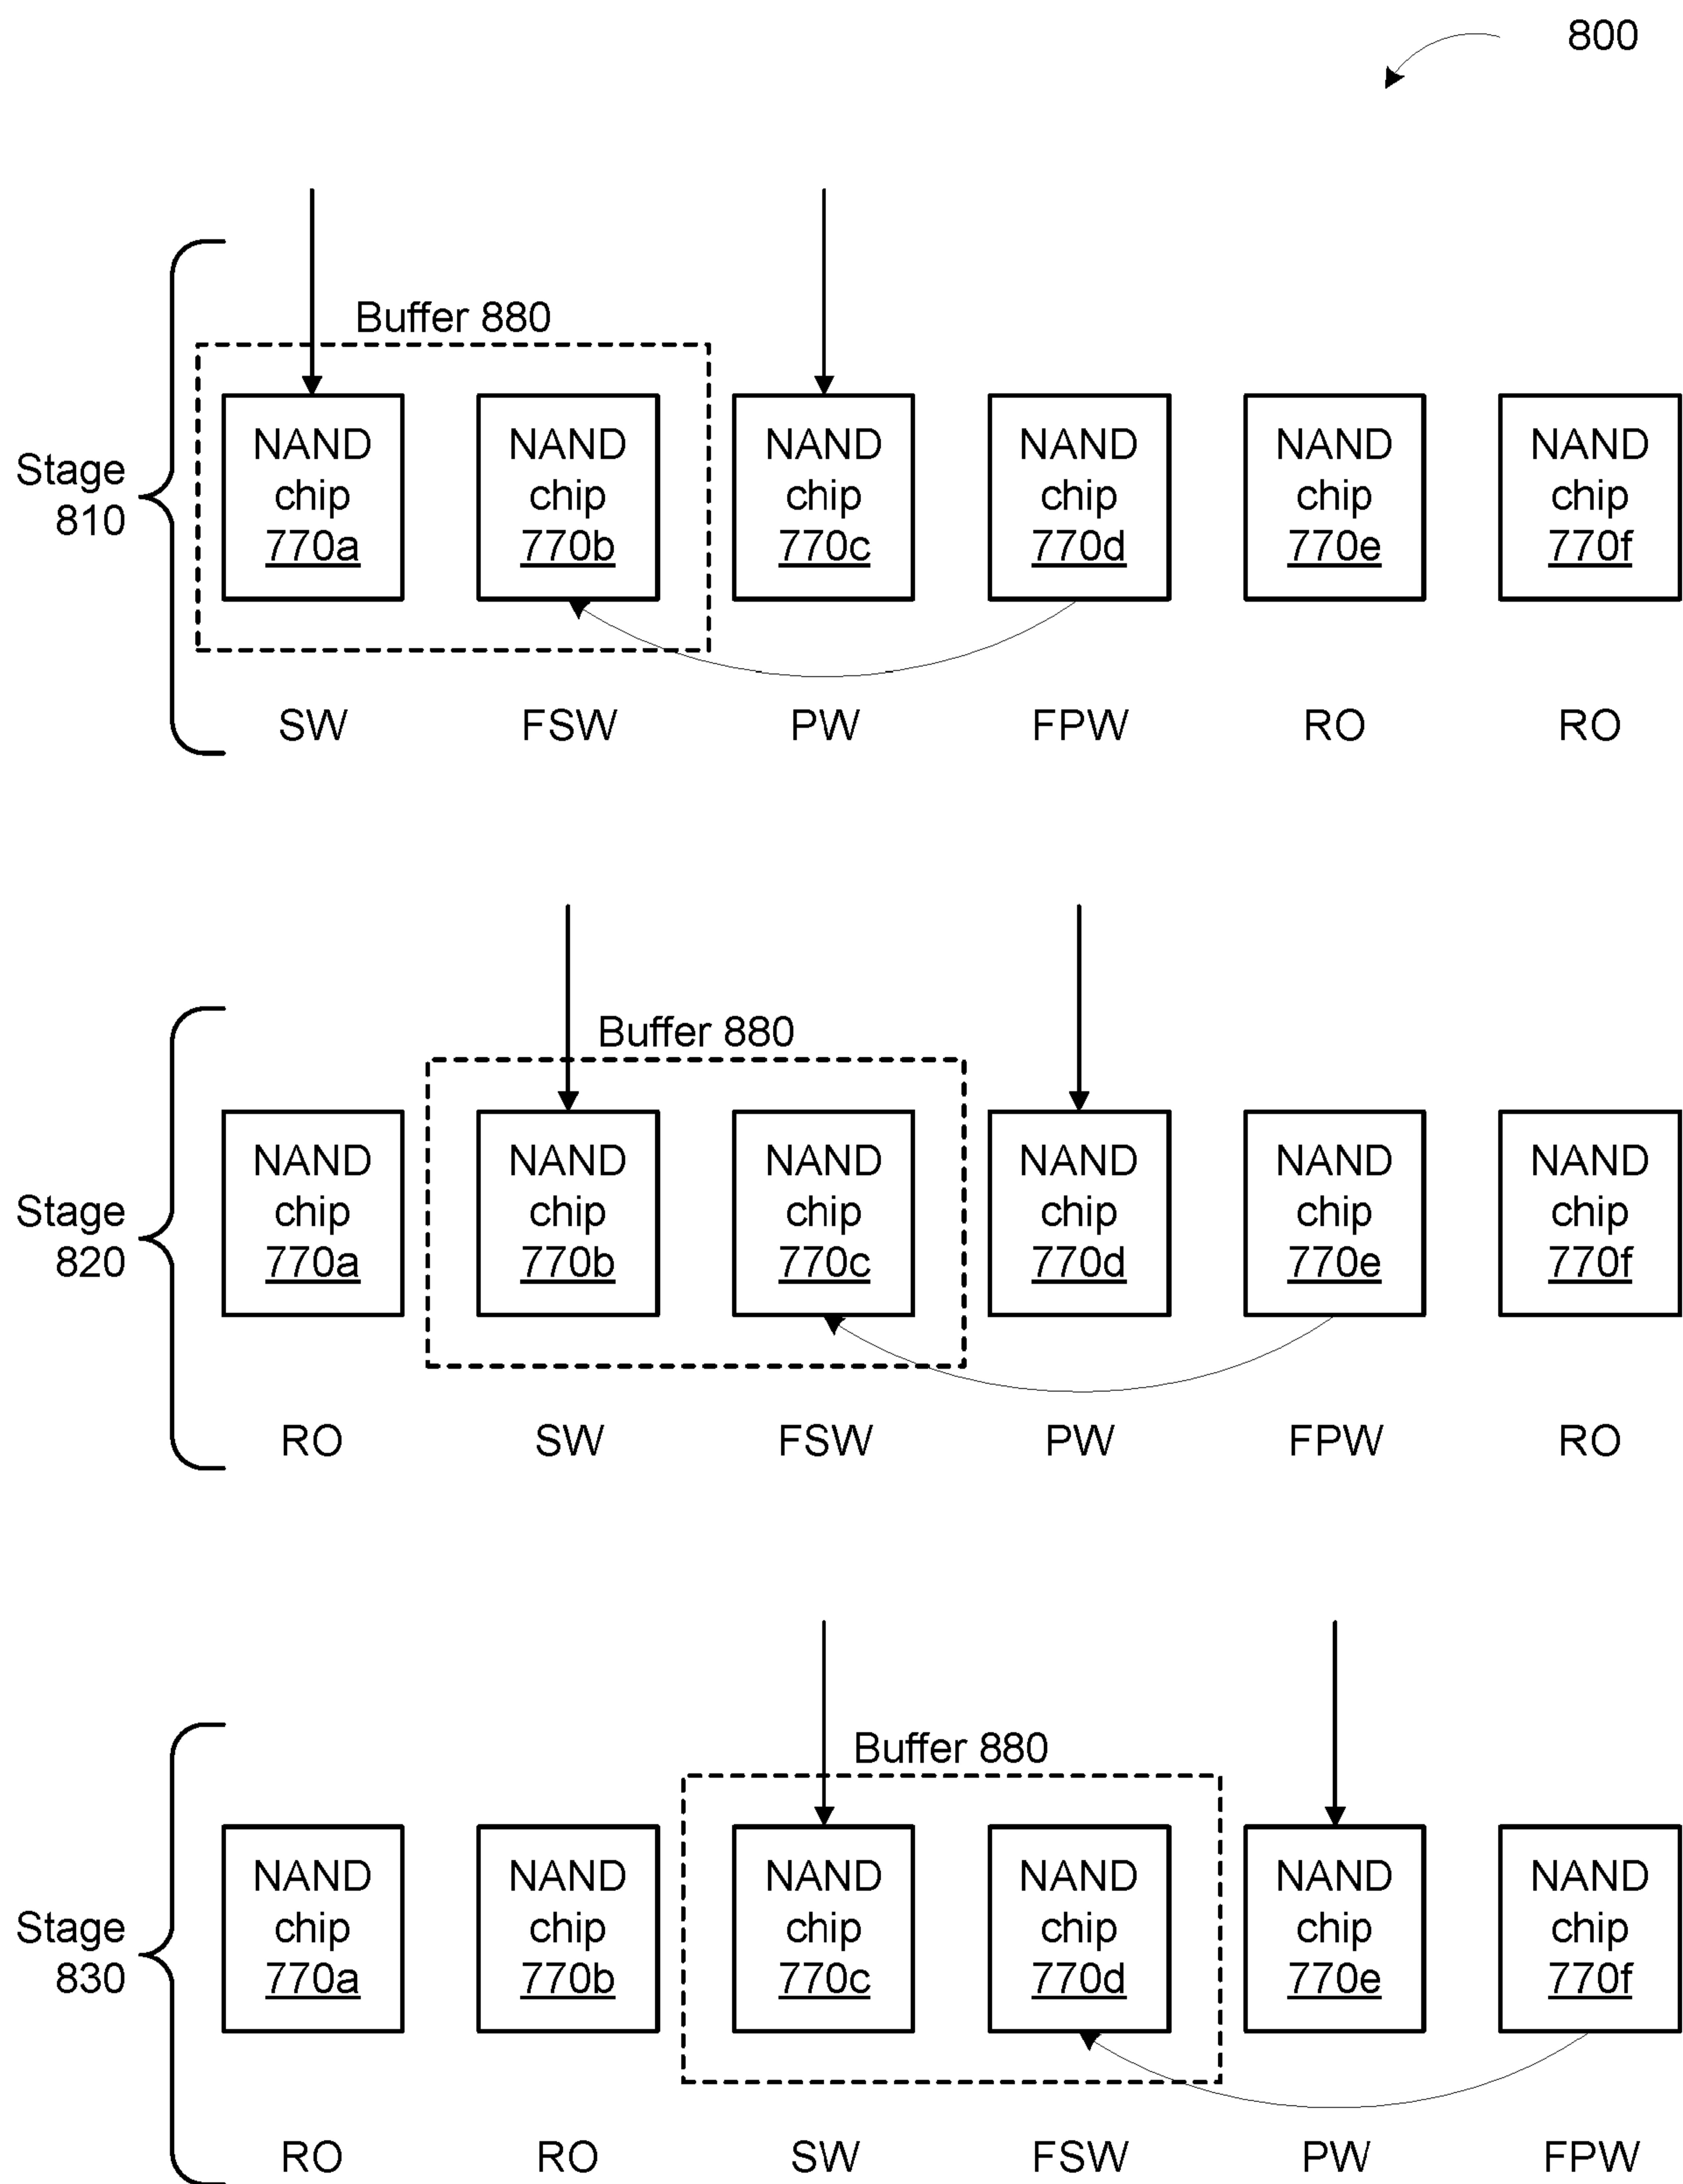


FIG. 8

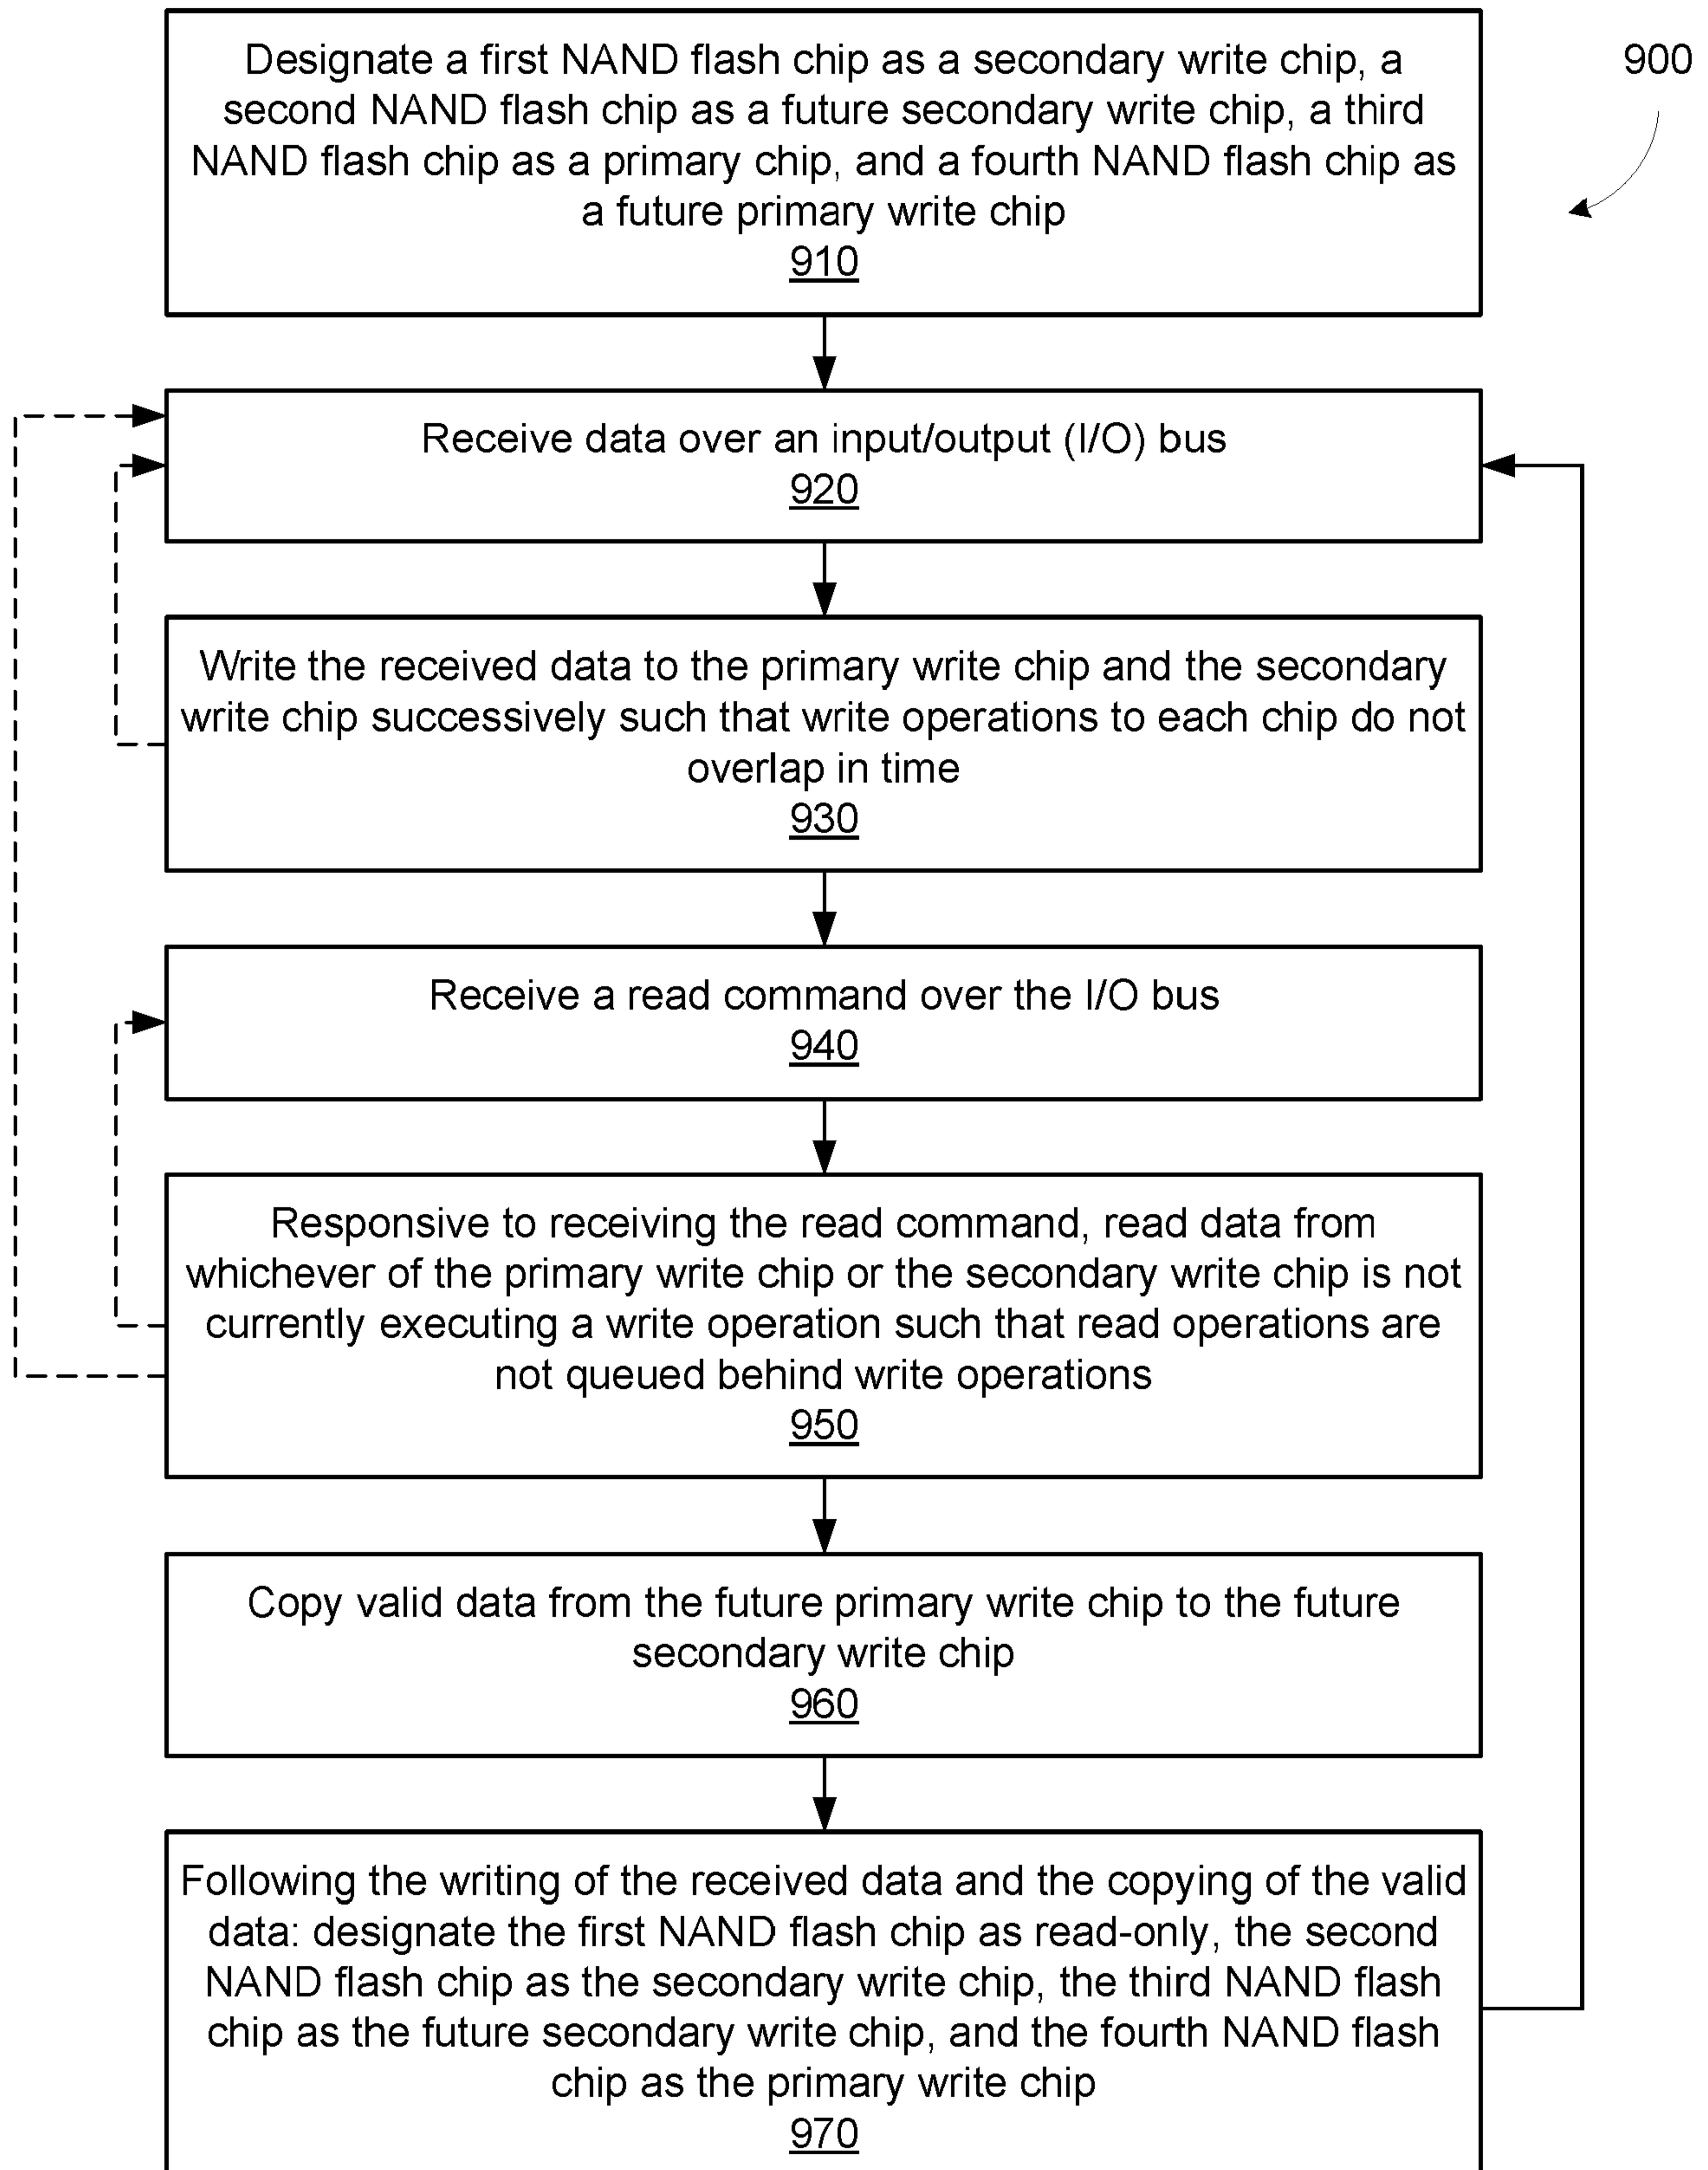


FIG. 9

NAND FLASH STORAGE DEVICE WITH NAND BUFFER

RELATED APPLICATIONS

[0001] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 62/524,016, titled “NAND FLASH STORAGE DEVICE WITH NAND BUFFER,” and filed on June 23, 2017, the entire contents of which are hereby incorporated by reference.

BACKGROUND

[0002] Storage devices based on NAND flash have good average latency; however, the worst-case latency, referred to as “tail latency” can be much longer. Providing more predictable latency in a NAND flash storage device may be beneficial in some applications of NAND flash technology; for example, in NAND flash-based storage devices and swap devices. A swap device can serve as extended memory space, sometimes referred to as swap space, in a computing system. The swap device can represent a tradeoff between the faster speed of RAM and the lower cost of alternative memory/storage mediums such as a hard disk drive (HDD) or a solid state drive (SSD). An SSD based on NAND flash, while faster than an HDD, may suffer from unpredictable tail latency due to the way NAND flash chips operate.

SUMMARY

[0003] At least one aspect is directed to a NAND flash storage device. The NAND flash storage device includes a plurality of NAND flash chips and a controller. The controller is configured to designate a first NAND flash chip as a secondary write chip, a second NAND flash chip as a future secondary write chip, a third NAND flash chip as a primary write chip, a fourth NAND flash chip as a future primary write chip, and any remaining NAND flash chips as read-only. The controller is configured to receive data over an input/output (I/O) bus. The controller is configured to write the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The controller is configured to receive a read command over the I/O bus. The controller is configured to, responsive to receiving the read command, read data from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read

07 10 19

operations are not queued behind write operations. The controller is further configured to copy valid data from the future primary write chip to the future secondary write chip; and following the writing of the received data and the copying of the valid data, designate the first NAND flash chip as read-only, the second NAND flash chip as the secondary write chip, the third NAND flash chip as the future secondary write chip, and the fourth NAND flash chip as the primary write chip.

[0004] In some implementations, the controller can be further configured to, following the writing of the received data and the copying of the valid data, designate a fifth NAND flash chip as the future primary write chip. The controller can be configured to write second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The controller can be configured to copy second valid data from the future primary write chip to the future secondary write chip. The controller can be configured to, following the writing of the second received data and the copying of the second valid data, designate the second NAND flash chip as read-only, the third NAND flash chip as the secondary write chip, the fourth NAND flash chip as the future secondary write chip, and the fifth NAND flash chip as the primary write chip.

[0005] In some implementations, the copy operation can include copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip. The copying operation can include copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip. The copying operation can include erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

[0006] In some implementations, the write operation and the copy operation are performed during overlapping time periods.

[0007] At least one aspect is directed to a NAND flash storage device. The NAND flash storage device includes a plurality of buffer NAND flash chips, a plurality of storage NAND flash chips, and a controller. The controller is configured to designate a first storage NAND flash chip as a primary write chip, a second storage NAND flash chip as a future primary write

chip, a first buffer NAND flash chip as a secondary write chip, a second buffer NAND flash chip as a future secondary write chip, and any remaining NAND flash chips as read-only. The controller is further configured to receive incoming data over an input/output (I/O) bus. The controller is configured to write the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The controller is configured to receive a read command over the I/O bus. The controller is configured to, responsive to receiving the read command, read data from whichever of the first buffer NAND flash chip or the first storage NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations. The controller is configured to copy valid data from the future primary write chip to the future secondary write chip; and following the writing of the received data and the copying of the valid data, designate the first storage NAND flash chip as read-only, the second storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the future secondary write chip, and the second buffer NAND flash chip as the secondary write chip.

[0008] In some implementations, the controller can be configured to, following the writing of the received data and the copying of the valid data, designate a third storage NAND flash chip as the future primary write chip. The controller can be configured to write second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The controller can be configured to copy second valid data from the future primary write chip to the future secondary write chip. The controller can be configured to, following the writing of the second received data and the copying of the second valid data, designate the second storage NAND flash chip as read-only, the third storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the secondary write chip, and the second buffer NAND flash chip as the future secondary write chip.

[0009] In some implementations, the copy operation can include copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip. The copy operation can include copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future

primary write chip; and erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

[0010] In some implementations, the buffer NAND flash chips can be single-level cell (SLC) NAND flash chips, and the storage NAND flash chips can be triple-level cell (TLC) NAND flash chips.

[0011] At least one aspect is directed to method of buffering a NAND flash storage device. The method includes designating a first NAND flash chip as a secondary write chip, a second NAND flash chip as a future secondary write chip, a third NAND flash chip as a primary write chip, a fourth NAND flash chip as a future primary write chip, and any remaining NAND flash chips as read-only. The method includes receiving data over an input/output (I/O) bus. The method includes writing the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time.. The method includes receiving a read command over the I/O bus. The method includes, responsive to receiving the read command, reading data from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations. The method includes copying valid data from the future primary write chip to the future secondary write chip; and following the writing of the received data and the copying of the valid data: designating the first NAND flash chip as read-only, the second NAND flash chip as the secondary write chip, the third NAND flash chip as the future secondary write chip, and the fourth NAND flash chip as the primary write chip

[0012] In some implementations, the method can include, following the writing of the received data and the copying of the valid data, designating a fifth NAND flash chip as the future primary write chip. The method can include writing second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The method can include copying second valid data from the future primary write chip to the future secondary write chip. The method can include, following the writing of the second received data and the copying of the second valid data, designating the second NAND flash chip as read-only, the third NAND flash chip as the secondary write chip, the fourth NAND flash chip as the future secondary write chip, and the fifth NAND flash chip as the primary write chip.

[0013] In some implementations, the copy operation can include copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip. The copy operation can include copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip. The copy operation can include erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

[0014] In some implementations, the writing operation and the copying operation are performed during overlapping time periods.

[0015] At least one aspect is directed to a method of buffering a NAND flash storage device. The method includes designating a first storage NAND flash chip as a primary write chip, a second storage NAND flash chip as a future primary write chip, a first buffer NAND flash chip as a secondary write chip, a second buffer NAND flash chip as a future secondary write chip, and any remaining NAND flash chips as read-only. The method includes receiving data over an input/output (I/O) bus. The method writing the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The method includes receiving a read command over the I/O bus. The method includes, responsive to receiving the read command, reading data from whichever of the first buffer NAND flash chip or the first storage NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations. The method includes copying valid data from the future primary write chip to the future secondary write chip; and following the writing of the received data and the copying of the valid data: designating the first storage NAND flash chip as read-only, the second storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the future secondary write chip, and the second buffer NAND flash chip as the secondary write chip.

[0016] In some implementations, the method includes, following the writing of the received data and the copying of the valid data, designating a third storage NAND flash chip as the future primary write chip. The method includes writing second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time. The method includes copying second valid data from the future

primary write chip to the future secondary write chip. The method includes, following the writing of the second received data and the copying of the second valid data, designating the second storage NAND flash chip as read-only, the third storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the secondary write chip, and the second buffer NAND flash chip as the future secondary write chip.

[0017] In some implementations, the method includes copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip. The method includes copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip. The method includes erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

[0018] In some implementations, the buffer NAND flash chips can be single-level cell (SLC) NAND flash chips, and the storage NAND flash chips can be triple-level cell (TLC) NAND flash chips.

[0019] These and other aspects and implementations are discussed in detail below. The foregoing information and the following detailed description include illustrative examples of various aspects and implementations, and provide an overview or framework for understanding the nature and character of the claimed aspects and implementations. The drawings provide illustration and a further understanding of the various aspects and implementations, and are incorporated in and constitute a part of this specification.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] The accompanying drawings are not intended to be drawn to scale. Like reference numbers and designations in the various drawings indicate like elements. For purposes of clarity, not every component may be labeled in every drawing. In the drawings:

[0021] Figure 1 is block diagram of an example computer system including a NAND flash storage device with a NAND buffer;

[0022] Figure 2 is a block diagram of an example NAND flash storage device with a NAND buffer;

[0023] Figure 3 is a flowchart showing an example method of using a NAND flash storage device with a NAND buffer;

[0024] Figure 4 is a block diagram of a NAND flash storage device with a NAND buffer and additional NAND storage, according to an implementation consistent with the claims;

[0025] Figure 5 shows operations of a NAND flash storage device with a NAND buffer and additional NAND storage, according to an implementation consistent with the claims;

[0026] Figure 6 is a flowchart showing a method of using a NAND flash storage device with a NAND buffer and additional NAND storage, according to an implementation consistent with the claims;

[0027] Figure 7 is block diagram of a NAND flash storage device with a rotating NAND buffer for wear leveling, according to a implementation consistent with the claims;

[0028] Figure 8 shows operations of a NAND flash storage device with a rotating NAND buffer for wear leveling, according to an implementation consistent with the claims; and

[0029] Figure 9 is a flowchart showing a method of using a NAND flash storage device with a rotating NAND buffer for wear leveling, according to an implementation consistent with the claims.

DETAILED DESCRIPTION

[0030] Following below are descriptions of various concepts related to, and implementations of, a predictable-latency NAND flash storage device with NAND buffer and methods of using the same. The various concepts introduced above and discussed in greater detail below may be implemented in any of numerous ways, as the described concepts are not limited to any particular manner of implementation. Examples of specific implementations and applications are provided primarily for illustrative purposes.

[0031] This disclosure generally relates to a predictable-latency NAND flash storage device with a NAND buffer. The storage device can be a solid-state drive (SSD) or a swap device. A swap device is a secondary storage device that can function as an extension of Random Access Memory (RAM). When the amount of data to be retained in memory exceeds the capacity of RAM, some of the data can be moved to the swap device in an operation sometimes referred to as “paging.” Traditionally, the secondary device could be a hard disk drive (HDD). Reading and writing data from an HDD is very slow relative to RAM, however, making an HDD impractical as a swap device in settings where retrieval speed is important. More recent SSDs, particularly NAND flash-based SSDs, although still slower than RAM, exhibit faster retrieval times relative to HDDs. NAND SSDs also tend to be less expensive than RAM, thus reducing RAM capacity in favor of SSD capacity can lead to cost savings. More recently developed low-latency NAND (LLNAND), such as Samsung’s Z-NAND, are designed with special circuitry that improves its typical read latency. LLNAND-based SSDs are even faster, and thus practical for paging and other temporary data storage in some server applications.

[0032] NAND flash SSDs have other drawbacks, however. While NAND read time is very fast, on the order of 10s of microseconds, erase and write operations are much slower, on the order of milliseconds each. And, because a NAND flash chip cannot execute a read operation and write operation at the same time, a read operation queued behind an erase/write operation can exhibit a latency increase of two or three orders of magnitude from normal. This is called “tail latency.” While it is possible to buffer data in RAM for writing to a NAND flash chip, the overhead can be cost prohibitive. For example, a device having ten NAND flash chips will require RAM equivalent to one tenth of the total NAND capacity as a buffer. Because DRAM can be four times as expensive as NAND storage, the cost overhead may be 40%.

[0033] This disclosure therefore proposes employing a NAND buffer to improve tail latency at a lower cost overhead than that of a RAM buffer. The system works by writing each page of data into two chips in sequence: first into a first chip, and then into a second chip. The second chip may be considered a NAND buffer. The write operations do not overlap in time. Read operations can be executed by the chip that is presently not executing a write operation. Thus, the read operations need not be queued behind write operations, allowing the NAND flash storage device to have more predictable latency.

[0034] In some implementations, the first chip (and, by extension, the second chip) are filled, and two new chips are chosen as write buffers. The first chip can be switched to read only and, because the first chip can perform read operations without interruptions from write operations, the copy of the data stored on the second chip is no longer necessary. The second chip is now free for future primary write operations. Two new chips, a third chip and a fourth chip, can be selected for writing. In this case, the fourth chip may be considered a NAND buffer.

[0035] For such a system to work, it is helpful for the two new chips selected for writing to contain duplicate data. Accordingly, prior to receiving new data, existing data on the third chip can first be copied to the fourth chip. The copying operation involves “garbage collection,” an operation memory devices perform to free up space occupied by data that has been invalidated (is no longer in use). Garbage collection can be time consuming and, like erase/write operations, can prevent read operations from executing until it is complete. Therefore, the garbage collection and copying operations are executed such that initially, only the fourth chip is being written and third chip is available for reading. Later, the garbage collected version is copied back from the fourth chip to the third chip such that the fourth chip is available for reading. Further, the garbage collection and copy operations between third and fourth chip are overlapped in time with writing into the first and second chip. As long as the garbage collection and copy operation takes less time than the writing operations, the third and fourth chips will be ready to take over execution of write operations when the first and second chips are filled. A new pair of chips will then be selected for garbage collection and copy operation.

[0036] In such a system, at any given time, two buffer chips contain a duplicate of the data in two other chips. Accordingly, a device having ten NAND flash chips for storage can have two additional NAND flash buffer chips. The cost overhead will thus be 20%, which can be half the cost overhead of using a DRAM buffer.

[0037] Additional buffer NAND flash chips can be employed to improve the write bandwidth of the NAND flash storage device. The write bandwidth of a NAND flash storage device refers to the amount of data the NAND flash storage device can receive and store in a given amount of time. In some cases, the write bandwidth of a NAND flash storage device can be limited by the time it takes to write data to a particular NAND flash chip. Therefore, write bandwidth can be increased by adding more buffer NAND flash chips, and distributing the write

operations among them. Accordingly, in some implementations, the NAND flash storage device can include additional buffer chips (i.e., 4, 6, 8, or more) to improve the write bandwidth.

[0038] In some implementation, the NAND flash storage device can employ wear leveling. A NAND flash chip can only perform a finite number of write operations before exhibiting excessive bit errors. Eventually, after performing many write operations, the NAND flash chip will wear out and begin to fail. To prolong the life of the NAND flash storage device, the NAND flash storage device can rotate which NAND flash chips are designated as buffer chips for write operations. The designation can be rotated to equalize or substantially equalize wear among some or all of the NAND flash chips in the device, as opposed to concentrating buffering write operations on a small number of chips. In this manner, the wear can be spread, or “leveled,” across many NAND flash chips of the device.

[0039] Figure 1 is block diagram of an example computer system 100 including a NAND flash storage device with a NAND buffer. The system 100 can be a computing device; for example and without limitation, a personal computing device such as a laptop or desktop computer, a shared computing device such as a server or mainframe, or component of the aforementioned or other computing device. The system 100 includes a processor such as a central processing unit (CPU) 110, RAM 120 such as DRAM, an input/output (I/O) bus 130 such as a Peripheral Component Interconnect Express (PCIe), a data storage device 140 such as a hard disk drive (HDD) or a solid-state drive (SSD) for long-term data storage, and a NAND flash storage device 150. The RAM 120 may be a main memory and can be directly accessible to the CPU 110 over, for example, an address bus that carries a memory address of the desired location, and a data bus used to read or write data from/to the specified memory address.

[0040] The NAND flash storage device 150 can be a swap device providing a swap space to the system 100. In some computing environments, the kernel or operating system can employ a memory/storage abstraction such as virtual memory, which provides applications executing on the CPU 110 with transparent and consistent access to one or more memory devices that can vary in type and location. For example, virtual memory can allow applications executing on the CPU 110 to access an extended memory space. The extended memory space can exceed the capacity of the RAM 120. When the data sought to be stored in memory exceed the capacity of the RAM 120, the virtual memory function can perform paging, by which the system 100 can store data in a secondary memory device such as an HDD or a SSD. This extended memory space can be

referred to as swap space, and the storage device providing the swap space can be referred to as a swap device.

[0041] Different types of memory and storage devices can vary greatly in access speed, with RAM generally being much faster than HDDs, and NAND flash falling in the middle. The slow access times of HDDs make them impractical for applications requiring fast access to data. NAND flash storage devices are much faster than HDDs, and thus become useful for a wider range of applications. NAND flash storage devices have the additional advantage of being less expensive than some common types of RAM such as DRAM. Newer types of NAND such as low-latency NAND (LLNAND) can be even faster, exhibiting latency for read operations of approximately 10 μ s versus approximately 70 μ s regular NAND. LLNANDs also exhibit less latency for program (write) operations; i.e., 100–200 μ s versus a few milliseconds for regular NAND. A write operation, however, generally involves both a program operation and an erase operation. LLNANDs may still suffer from slow erase times, which can take on the order of 5 ms. As an additional challenge, a NAND flash chip or die can generally only perform one operation at a time. For example, if a NAND flash chip is executing an erase/program sequence of a write operation, any incoming read commands will be queued until completion. Thus, some read operations may require up to several milliseconds to complete despite a nominal latency of tens of microseconds. These outlier latencies are referred to as “tail latency.” Both the duration and the unpredictability of this tail latency can present difficulties when using NAND flash. Some systems can use a RAM to cache the contents of a NAND flash chip being written to such that the RAM can perform the read operation while the NAND flash chip executes write operations—see, for example, U.S. Patent Application Number 14/936,486. Because RAM can be more expensive than NAND flash, however, the RAM adds a cost overhead that can be unacceptable for some applications. Therefore, this disclosure proposes providing one or more additional NAND flash chips to provide a cache for data residing in NAND flash chips executing write operations. A NAND flash-based buffer is described in more detail below with regard to Figure 2.

[0042] Figure 2 is a block diagram of an example NAND flash storage device 250 with a NAND buffer. The NAND flash storage device 250 can serve as the NAND flash storage device 150. To avoid tail latency associated with read operations being queued behind write operations, the NAND flash storage device 250 can employ double buffering using two NAND flash chips 270a and 270b (collectively, “NAND flash chips 270”). A controller 260 can handle storage

access commands received over the I/O bus 230. The controller 260 can include programmable logic such as an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a microcontroller, or a microprocessor. The controller 260 can be integral with, or external to the NAND flash storage device 250. The controller 260 can orchestrate read and write operations of the NAND flash storage device 250 such that the CPU 110 need not be aware of, or itself execute, the buffering operations of the NAND flash storage device 250.

Alternatively, it is possible to design a simpler controller that relies on the CPU 110 to orchestrate the operations, while the controller performs the lower level work of completing each operation. Accordingly, the controller 260 is configured to receive data and/or commands over the I/O bus 230, and write the received data first to one of the NAND flash chips 270, and then to a second of the NAND flash chips 270. The controller 260 executes the write operations successively—either autonomously or under the direction of CPU 110—such that they do not overlap in time; thus, at any given time, the NAND flash chip 270 that is not presently performing a write operation can promptly perform read operations. Once the controller 260 executes each write operation, two of the NAND flash chips 270 can contain the same data. The controller 260 can execute read commands by reading from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations. In this manner, tail latency can be improved.

[0043] In some implementations, the NAND flash storage device 250 can include additional NAND flash chips 270c and 270d to further improve write bandwidth. For example, a NAND flash storage device 250 having four NAND flash chips 270a through 270d can split write operations between two NAND flash chips 270; e.g., NAND flash chips 270a and 270b. After the NAND flash chips 270a and 270b complete their respective write operations, those write operations can be repeated on the NAND flash chips 270c and 270d. In this manner, the write bandwidth of the NAND flash storage device 250 can be approximately doubled. The NAND flash storage device 250 can employ additional NAND flash chips 270 for further increases in write bandwidth.

[0044] In some implementations, the NAND device 250 can include some RAM 265 to serve as a cache or buffer to store data while it is written to each NAND flash chip 270. This RAM 265 need not have the capacity to buffer the entire contents of a NAND flash chip 270, but rather can be only large enough to buffer a single block or page of data. Accordingly, it can be much

smaller than a dedicated RAM buffer, and thus not contribute appreciably to the cost overhead of the NAND device 250.

[0045] Figure 3 is a flowchart showing an example method 300 of using a NAND flash storage device with a NAND buffer. The NAND flash storage device can include any of the NAND flash storage devices 150, 250, 450, or 750 described herein. The method 300 includes receiving data over an input/output (I/O) bus (stage 310). The method 300 includes writing the received data to a first NAND flash chip and a second NAND flash chip successively such that the write operations to each NAND flash chip do not overlap in time (stage 320). The method includes receiving a read command over the I/O bus (stage 330). The method 300 includes reading data from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations (stage 340).

[0046] The method 300 includes receiving data over an I/O bus (stage 310). The NAND flash storage device can receive a write command over the I/O bus, such as one of the I/O busses 130 or 230 previously described with respect to the NAND flash storage devices 150 or 250, respectively. The write command can include data to be stored by the NAND flash storage device.

[0047] The method 300 includes writing the received data to a first NAND flash chip and a second NAND flash chip successively such that the write operations to each NAND flash chip do not overlap in time (stage 320). The write operations can be executed by a controller of the NAND flash storage device, such as the CPU 110 or the controller 260 previously described. The controller can execute the write operation on one of two NAND flash chips, such as the NAND flash chip 270a or 270b. While the controller executes a write operation on one of the NAND flash chips 270, that particular NAND flash chip will not be able to process other operations. Accordingly, the controller will execute the write operations on the two NAND flash chips successively such that they do not overlap in time. In some implementations, the receiving stage 310 and the writing stage 320 can be repeated until the first and second NAND flash chips are full or substantially full. In some implementations, the NAND flash chips are low-latency NAND (LLNAND) chips.

[0048] The method includes receiving a read command over the I/O bus (stage 330). The method 300 includes reading data from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations (stage 340). During execution of the write operations, the NAND flash storage device may receive a read command over the I/O bus. The controller can direct the read command to whichever of the NAND flash chips is not currently executing a write operation. In this manner, stages 330 and 340 can proceed concurrently with stages 310 and 320, and read operations can be processed promptly without being queued behind slower write operations. The stages of the method 300 can repeat as additional data and read commands are received over the I/O bus. In some implementations, the role played by each particular NAND flash chip can cycle, with different NAND flash chips performing different buffer and storage functions at different times. Examples of different manners of NAND flash chip buffering and storage are described below with reference to Figures 4 through 9.

[0049] Figure 4 is a block diagram of an example NAND flash storage device 450 with a NAND buffer and additional NAND storage, according to an implementation consistent with the claims. The NAND flash storage device 450 includes a controller 460 for executing read and write commands received over the I/O bus 430. The controller 460 can include RAM 465, or be in communication with the RAM 465 via a memory bus. The controller 460 can include programmable logic such as an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a microcontroller, or a microprocessor. The controller 460 can be integral with, or external to the NAND flash storage device 450. The NAND flash storage device 450 can include two buffer NAND flash chips 470a and 470b (collectively, “buffer NAND flash chips 470”). In some implementations, the NAND flash storage device 450 can include additional buffer NAND flash chips 470, as described further below. In some implementations, the NAND flash storage device 450 can use the RAM 465 for buffering data received over the I/O bus 430 while it is being written to one or more of the buffer NAND flash chips 470. The NAND flash storage device 450 also includes storage NAND flash chips 475a through 475e (collectively, “storage NAND flash chips 475”) that provide additional memory capacity over that of the buffer NAND flash chips 470. Data received over the I/O bus 430 can be buffered in the buffer NAND flash chips 470 before being copied to one of the storage NAND flash chips 475.

[0050] In some implementations, the buffer NAND flash chips 470 can be single level cell (SLC) NAND flash chips or low-latency NAND (LLNAND) flash chips while the storage NAND flash chips 475 may be multi-level cell (MLC) or triple-level cell (TLC) NAND flash chips. In some implementations, the buffer NAND flash chips 470 can be of smaller capacity than the storage NAND flash chips 475. In some implementations, a group of buffer NAND flash chips 470 can have a capacity equal to the capacity of a single storage NAND flash chip 475. The group of buffer NAND flash chips 470 can be used collectively as a single buffer chip.

[0051] The write bandwidth of NAND flash storage devices can be limited due to the relatively slow pace of performing write operations. However, in some implementations, write bandwidth can be improved by designating more buffer NAND flash chips 470 as write buffers. In some implementations, and similar to the NAND flash storage device 250, the NAND flash storage device 450 can include additional buffer NAND flash chips 470; for example, buffer NAND flash chips 470c and 470d. In some implementations, the NAND flash storage device 450 can include even more buffer NAND flash chips 470. The controller 460 can thus spread write operations among the buffer NAND flash chips 470 to multiply the write bandwidth.

[0052] The NAND flash storage device 450 can potentially face two drawbacks. First, the buffer NAND flash chips 470 will eventually reach capacity. Once all data in a buffer NAND flash chip 470 has been copied to one of the storage NAND flash chips 475, all data in the buffer NAND flash chip 470 can be deleted. While the buffer NAND flash chip 470 performs erase operations, it cannot perform a write operation. Therefore, the write bandwidth will be limited by the unavailability of that buffer NAND flash chip 470.

[0053] The second potential drawback relates to the finite number of write operations that a NAND flash chip can handle before exhibiting excessive bit errors. NAND flash chips experience wear from processing write operations. Eventually, after performing many write operations, the NAND flash chip will wear out and begin to exhibit failures.

[0054] In some implementations, the challenge of limited write bandwidth can be addressed by increasing the number of buffer NAND flash chips 470 as described above, and the NAND wear problem can be addressed by using higher endurance chips as buffer NAND flash chips 470. In some implementations, the NAND wear problem can be addressed by rotating the particular NAND flash chips designated as buffers for incoming data. Such a wear-leveling

scheme is described further below with reference to Figures 7 through 9. Figure 5 below illustrates operations of the NAND flash storage device 450.

[0055] Figure 5 shows operations 500 of a NAND flash storage device with a NAND buffer and additional NAND storage, according to an implementation consistent with the claims. These operations can be performed by a NAND flash storage device such as the NAND flash storage device 450 previously described. The operations 500 include stage 510, in which incoming data is written to the buffer NAND flash chip 470a and the storage NAND flash chip 475a, and data is copied with garbage collection from the storage NAND flash chip 475b to the buffer NAND flash chip 470b and subsequently copied back to the storage NAND flash chip 475b. The operations 500 include stage 520, in which incoming data is written to the buffer NAND flash chip 470b and the storage NAND flash chip 475b, and data is copied with garbage collection from the storage NAND flash chip 475c to the buffer NAND flash chip 470a and subsequently copied back to the storage NAND flash chip 475c. The operations 500 include stage 530, in which incoming data is written to the buffer NAND flash chip 470a and the storage NAND flash chip 475c, and the buffer NAND flash chip 470b receives data copied from the storage NAND flash chip 475d. The buffer NAND flash chips 470 are indicated by a buffer 580 box. The operation each NAND flash chip 470 or 475 is currently performing is indicated by an acronym. “PW” stands for “primary write” and indicates that the storage NAND flash chip 475 is currently designated for receiving data and performing write operations. “SW” stands for “secondary write” and indicates that the buffer NAND flash chip 470 is also currently designated for receiving data and performing write operations. “FPW” stands for “future primary write” and indicates that the NAND flash chip 475 is designated for being prepared for future service receiving data and performing write operations. “FSW” stands for “future secondary write” and indicates that the NAND flash chip 470 is designated for being prepared for future service receiving data and performing write operations. “RO” stands for “read-only” and indicates that the controller 460 will only call upon the RO storage NAND flash chips 475 to execute read operations. These labels are merely arbitrary, and are assigned only for the purpose of describing the operation of the various chips at particular points in time. Between each state of the operations 500, the role of each NAND flash chip 470 or 475 changes as different NAND flash chips 470 or 475 receive data and perform write operations.

[0056] During stage 510, the buffer NAND flash chip 470a and the storage NAND flash chip 475a receive and write incoming data, as indicated by the arrows. The buffer NAND flash chip

470a and the storage NAND flash chip 475a receive and write the same data, as the SW and the PW chips, respectively; however, they can perform the write operations during non-overlapping time periods. In this manner, both PW and the SW chips contain substantially the same data, and a NAND flash chip that is not performing a write operation can respond to a read command and perform the read operation promptly without having to wait for completion of the relatively long write operation. In some implementations, the PW and SW chips can continue receiving data and performing write operations until they are full or substantially full. Meanwhile, the FPW storage NAND flash chip 475b copies its data to the FSW buffer NAND flash chip 470b in preparation for receiving external data in the next stage 520. During the copy operation, the NAND flash storage device 450 can perform garbage collection on the FPW and the FSW chips. When fragmented data is copied from the FPW chip to the FSW chip, it is consolidated so that it is written to the FSW chip in a continuous or substantially continuous memory space, while free space is similarly consolidated in a second continuous or substantially continuous memory space. This makes the data representation more compact and leaves space at the end for new data to be written at the next stage 520. After all valid data is copied from the FPW chip to the FSW chip, the free space on the FPW is erased, and the valid data is copied back from the FSW chip to the FPW chip. When the PW and SW chips have completed their write operations and the FPW and the FSW chips have completed their copying and garbage collection, the operations 500 can continue with the next stage 520.

[0057] During stage 510 (and stages 520 and 530 as well) the NAND flash storage device 450 may from time to time receive a read request. The controller 460 may handle the read request in one of three ways depending on whether the requested data is on a RO chip, one of the PW or SW chips, or one of the FPW or FSW chips. If the requested data is on a RO chip, the controller 460 can simply execute the read operation on the RO chip and returns the data—there is no write operation occurring that would interfere with or delay execution of the read operation. If the requested data is on the PW or SW chip, the controller 460 can execute the read operation on whichever of the PW or SW chips is not currently executing a write operation such that read operations are not queued behind write operations. If the requested data is on the FPW or FSW chip, the controller 460 can execute the read operation on whichever of the FPW or FSW is not currently executing a write operation such that read operations are not queued behind write operations. That is, if data is currently being copied with garbage collection from the FPW chip to the FSW chip, the controller 460 will be executing write operations on the FSW chip, and so

can execute the read operation on the FPW chip. When the data is being copied back from the FSW chip to the FPW chip, the controller 460 can execute the read operation on the FSW chip. In this manner, regardless of which NAND flash chip contains the requested data and which NAND flash chip[s] is/are executing write operations, the NAND flash device 450 can read and return the requested data without undue delay caused by queueing the read operation behind a write operation.

[0058] During stage 520, the storage NAND flash chip 475b is designated as the PW chip, the buffer NAND flash chip 470b is designated as the SW chip, the storage NAND flash chip 475c is designated as the FPW chip, the buffer NAND flash chip 470a is designated as the FSW chip, and the storage NAND flash chip 475a is designated RO. The functions of the PW, SW, FPW, and FSW chips are similar to functions performed by the chips with the same designations in the previous stage 510, but with different NAND flash chips performing each designated role. The SW and the PW chips receive and write the same data as indicated by the arrows; however, they can perform the write operations during non-overlapping time periods so that they can respond to a read command and perform the read operation promptly without having to wait for completion of the relatively long write operation. Meanwhile, the FPW chip copies its data to the FSW chip in preparation for receiving external data in the next stage 530. During the copy operation, the FPW and the FSW chips can perform garbage collection operations. When the PW and SW chips have completed their write operations and the FPW and the FSW chips have completed their copying and garbage collection, the operations 500 can continue with the next stage 530.

[0059] During stage 530, the storage NAND flash chip 475c is designated as the PW chip, the buffer NAND flash chip 470a is designated as the SW chip, the storage NAND flash chip 475d is designated as the FPW chip, the buffer NAND flash chip 470b is designated as the FSW chip, and the storage NAND flash chip 475b is designated RO. The functions of the PW, SW, FPW, and FSW chips are similar to functions performed by the chips with the same designations in the previous stages 510 and 520, but with different NAND flash chips performing each designated role. The SW and the PW chips receive and write the same data as indicated by the arrows; however, they can perform the write operations during non-overlapping time periods so that they can respond to a read command and perform the read operation promptly without having to wait for completion of the relatively long write operation. Meanwhile, the FPW chip copies its data to the FSW chip in preparation for receiving external data in the next stage.

During the copy operation, the FPW and the FSW chips can perform garbage collection operations. When the PW and SW chips have completed their write operations and the FPW and the FSW chips have completed their copying and garbage collection, the operations 500 can continue with the next stage. The operations 500 can continue in this manner until all of the storage NAND flash chips 475 have been written to, at which time the operations 500 can repeat starting at the stage 510.

[0060] Figure 6 is a flowchart showing a method 600 of using a NAND flash storage device with a NAND buffer and additional NAND storage, according to an implementation consistent with the claims. The method 600 can provide for NAND flash storage space beyond the storage space provided by the buffer NAND chips. For example, the method 600 can be performed by the NAND flash storage device 450 previously described with reference to Figure 4. The NAND flash storage device 450 can include buffer NAND flash chips 470 and storage NAND flash chips 475. The method 600 includes designating a first storage NAND flash chip as a primary write (PW) chip, a second storage NAND flash chip as a future primary write (FPW) chip, a first buffer NAND flash chip as a secondary write (SW) chip, a second buffer NAND flash chip as a future secondary write (FSW) chip, and any remaining NAND flash chips as read-only (RO). The method 600 includes receiving data over an input/output (I/O) bus (stage 620). The method 600 includes writing the received data to the PW chip and the SW chip successively such that the write operations to each NAND flash chip do not overlap in time (stage 630). The method 600 includes receiving a read command over the I/O bus (stage 640). The method 600 includes, responsive to receiving the read command, reading data from a NAND flash chip that is not currently executing a write operation such that read operations are not queued behind write operations (stage 650). The method 600 includes copying valid data from the FPW chip to the FSW chip (stage 660). The method 600 includes, following the writing of the received data and the copying of the valid data, designating the first storage NAND flash chip as RO, the second storage NAND flash chip as the PW chip, the first buffer NAND flash chip as the FSW chip, and the second buffer NAND flash chip as the SW chip (stage 670).

[0061] The method 600 includes designating a first storage NAND flash chip as a primary write chip, a second storage NAND flash chip as a future primary write chip, a first buffer NAND flash chip as a secondary write chip, a second buffer NAND flash chip as a future secondary write chip, and any remaining NAND flash chips as read-only (RO).

[0062] The method 600 includes receiving data over an I/O bus (stage 620). Stage 620 can be similar to stage 310 described previously.

[0063] The method 600 includes writing the received data to the PW chip and the SW chip successively such that the write operations to each NAND flash chip do not overlap in time (stage 630). Stage 630 can be similar to stage 320 described previously. In some implementations, the receiving stage 620 and the writing stage 630 can be repeated until the PW chip and the SW chip are full or substantially full.

[0064] The method 600 includes receiving a read command over the I/O bus (stage 640). Stage 630 can be similar to stage 430 described previously.

[0065] The method 600 includes, responsive to receiving the read command, reading data from a NAND flash chip that is not currently executing a write operation such that read operations are not queued behind write operations (stage 650). Stage 650 can be similar to stage 340 described previously. In practice, the requested data may reside on a RO chip, one of the PW or SW chips, or one of the FPW or FSW chips. The controller may handle the read request in one of three ways depending on where the requested data resides. If the requested data is on the PW or SW chip, the controller can execute the read operation on whichever of the PW or SW chips is not currently executing a write operation such that read operations are not queued behind write operations, and the method 600 can proceed to the next stage. If the requested data is on a RO chip, the controller can simply execute the read operation on the RO chip and returns the data—there is no write operation occurring that would interfere with or delay execution of the read operation. If the requested data is on the FPW or FSW chip, the controller can execute the read operation on whichever of the FPW or FSW is not currently executing a write operation such that read operations are not queued behind write operations. That is, if data is currently being copied with garbage collection from the FPW chip to the FSW chip, the controller 460 will be executing write operations on the FSW chip, and so can execute the read operation on the FPW chip. When the data is being copied back from the FSW chip to the FPW chip, the controller can execute the read operation on the FSW chip. In this manner, regardless of which NAND flash chip contains the requested data and which NAND flash chip[s] is/are executing write operations, the NAND flash device can read and return the requested data without undue delay caused by queueing the read operation behind a write operation. In this manner, stages 640 and 650 can proceed concurrently with stages 620 and 630, and read operations can be processed

promptly without being queued behind slower write operations. The stages 620 through 650 can repeat as additional data and read commands are received over the I/O bus.

[0066] The method 600 includes copying valid data from the FPW chip to the FSW chip (stage 660). In some implementations, during the copy operation, the FPW and FSW chips can perform garbage collection to consolidate the valid data and the free space, respectively. For example, the NAND flash storage device can first copy the valid data from the FPW chip to the FSW chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the FSW chip. The NAND flash storage device can copy the valid data from the FSW chip back to the FPW chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the FPW chip. The NAND flash storage device can erase the memory spaces of the FPW and FSW chips occupied by the free space. The copying operations of stage 660 can be performed during an overlapping time period with the buffering operations of stages 620 and 630. In addition, read commands can be received and executed by whichever of the FPW and the FSW is not currently executing write operations as part of the copy operations. In this manner, the copying operations of stage 660 can also be performed during an overlapping time period with the read operations of stages 640 and 650.

[0067] The method 600 includes, following the writing of the received data and the copying of the valid data, designating the first storage NAND flash chip as RO, the second storage NAND flash chip as the PW chip, the first buffer NAND flash chip as the FSW chip, and the second buffer NAND flash chip as the SW chip (stage 670). The stages of the method 600 can repeat as additional data and read commands are received over the I/O bus. In some implementations, the role played by each particular NAND flash chip can cycle, with different NAND flash chips performing different buffer and storage functions at different times.

[0068] Figure 7 is block diagram of an example NAND flash storage device 750 with a rotating NAND buffer for wear leveling, according to an implementation consistent with the claims. The NAND flash storage device 750 includes a controller 760 for processing read and write commands received over the I/O bus 730. The controller 760 can include RAM 765, or be in communication with the RAM 765 via a memory bus. The controller 760 can include programmable logic such as an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a microcontroller, or a microprocessor. The controller 760

can be integral with, or external to the NAND flash storage device 750. The NAND flash storage device 750 includes NAND flash chips 770a to 770f (collectively, “NAND flash chips 770”). In some implementations, the NAND flash storage device 750 can include more or fewer NAND flash chips 770. In some implementations, the NAND flash storage device 750 can use the RAM 765 for buffering data received over the I/O bus 730 while it is being written to one or more of the NAND flash chips 770. The NAND flash storage device 750 can employ wear leveling by occasionally changing the NAND flash chips 770 used for buffering. For example, during a first stage of operation, NAND flash chips 770a and 770b can serve as the NAND buffers, with the controller 760 writing data to each one in turn. Example operations of the NAND flash storage device 750 are described in more detail below with reference to Figure 8.

[0069] Figure 8 shows example operations 800 of a NAND flash storage device with a rotating NAND buffer employing wear leveling, according to illustrative implementations. These operations can be performed by a NAND flash storage device such as the NAND flash storage device 750 previously described. The operations 800 include stage 810, in which incoming data is written to the NAND flash chips 770a and 770c, while data is copied with garbage collection from the NAND flash chip 770d to the NAND flash chip 770b and subsequently copied back to the NAND flash chip 770d. The operations 800 include stage 820, in which incoming data is written to the NAND flash chips 770b and 770d, while data is copied with garbage collection from the NAND flash chip 770e to the NAND flash chip 770c and subsequently copied back to the NAND flash chip 770e. The operations 800 include stage 830, in which incoming data is written to the NAND flash chips 770c and 770e, while data is copied with garbage collection from the NAND flash chip 770f to the NAND flash chip 770d and subsequently copied back to the NAND flash chip 770f. During each operation 800, the secondary write and the future secondary write NAND flash chips 770 are considered the buffer NAND flash chips, and are indicated by a buffer 880 box. The operation each NAND flash chip 770 is currently performing is indicated by an acronym. “PW” stands for “primary write” and indicates that the NAND flash chip 770 is currently receiving data and performing write operations. “SW” stands for “secondary write” and indicates that the NAND flash chip 770 is currently receiving data and performing write operations as one of the buffer NAND flash chips. “FPW” stands for “future primary write” and indicates that the NAND flash chip 770 is being prepared for future service as the primary write chip by copying its data to the future secondary write chip. “FSW” stands for “future secondary write” and indicates that the NAND flash chip

770 is receiving data from the FPW chip in preparation for future service as a SW buffer NAND flash chip. “RO” stands for “read-only” and indicates that the controller 760 will only call upon the RO NAND flash chip 770 to execute read operations. These labels are merely arbitrary, and are assigned only for the purpose of describing the function of the various chips during a given stage of operation 800. Between each stage of the operations 800, the role of each NAND flash chip 770 can change, as indicated by the labels below each chip and the buffer 880 box rotating through the available NAND flash chips 770.

[0070] During stage 810, the NAND flash chip 770a serves as the SW chip, the NAND flash chip 770b serves as the FSW chip, the NAND flash chip 770c serves as the PW chip, the NAND flash chip 770d serves as the FPW chip, and the remaining NAND flash chips serve as RO chips. Incoming data is written to the NAND flash chips 770a and 770c—the SW and PW chips, respectively. Meanwhile, data is copied with garbage collection from the NAND flash chip 770d to the NAND flash chip 770b—the FPW and FSW chips, respectively—and subsequently copied back to the NAND flash chip 770b. During stage 810, the buffer 880 includes the NAND flash chips 770a and 770b. The NAND flash chip 770b receives a copy of the data from the NAND flash chip 770d, so it may be at least partially unavailable for executing read operations, and may be affirmatively prevented by the controller 760 from executing read or write operations not associated with the copy operation. Garbage collection may be performed during the copy operation. This means that only valid data is copied and any gaps left by invalid data are removed in that process. This makes the data representation more compact and leaves space at the end for new data to be written at the next stage. In some implementations, this compacted form of data in NAND flash chip 770b may be copied back to NAND flash chip 770d. The NAND flash chips 770e and 770f, as well as any other NAND flash chips 770 not pictured, can serve as RO chips. At an appropriate time, the controller 760 can update the NAND flash chip 770 designations, and proceed to stage 820 of the operations 800. This can be when the buffer NAND flash chips 770a and 770c become full or substantially full, after a predetermined number of write operations have been performed, or following a predetermined interval of time.

[0071] During stage 810 (and stages 820 and 830 as well) the NAND flash storage device 750 may from time to time receive a read request. The controller 760 may handle the read request in one of three ways depending on whether the requested data is on a RO chip, one of the PW or SW chips, or one of the FPW or FSW chips. If the requested data is on a RO chip, the controller 760 can simply execute the read operation on the RO chip and returns the data—there

is no write operation occurring that would interfere with or delay execution of the read operation. If the requested data is on the PW or SW chip, the controller 760 can execute the read operation on whichever of the PW or SW chips is not currently executing a write operation such that read operations are not queued behind write operations. If the requested data is on the FPW or FSW chip, the controller 760 can execute the read operation on whichever of the FPW or FSW is not currently executing a write operation such that read operations are not queued behind write operations. That is, if data is currently being copied with garbage collection from the FPW chip to the FSW chip, the controller 760 will be executing write operations on the FSW chip, and so can execute the read operation on the FPW chip. When the data is being copied back from the FSW chip to the FPW chip, the controller 760 can execute the read operation on the FSW chip. In this manner, regardless of which NAND flash chip contains the requested data and which NAND flash chip[s] is/are executing write operations, the NAND flash device 750 can read and return the requested data without undue delay caused by queueing the read operation behind a write operation.

[0072] During stage 820, the NAND flash chip 770b serves as the SW chip, the NAND flash chip 770c serves as the FSW chip, the NAND flash chip 770d serves as the PW chip, the NAND flash chip 770e serves as the FPW chip, and the remaining NAND flash chips serve as RO chips. Incoming data is written to the NAND flash chips 770b and 770d, and data is copied with garbage collection from the NAND flash chip 770e to the NAND flash chip 770c, and subsequently copied back to the NAND flash chip 770e. The NAND flash chip 770a is designated as an RO chip, and is available to execute read operations without interruption from write operations. The buffer 880 now includes the NAND flash chips 770b and 770c. Because the NAND flash chip 770c previously served as the PW chip, the data it contains will not be unique because the data had previously been written to the previous SW as well. Accordingly, the NAND flash chip 770c can receive duplicate data from the NAND flash chip 770e in preparation for being designated SW in the next stage of operation. During the copying operation, the NAND flash chips 770c and 770e can perform garbage collection operations.

[0073] During stage 830, the NAND flash chip 770c serves as the SW chip, the NAND flash chip 770d serves as the FSW chip, the NAND flash chip 770e serves as the PW chip, the NAND flash chip 770f serves as the FPW chip, and the remaining NAND flash chips serve as RO flash chips. Incoming data is written to the NAND flash chips 770c and 770e, while data is copied with garbage collection from the NAND flash chip 770f to the NAND flash chip 770d, and

subsequently copied back to the NAND flash chip 770f. The NAND flash chip 770b is designated as an RO chip, and is available to execute read operations without interruptions from write operations. The buffer 880 now includes the NAND flash chips 770c and 770d. Because the NAND flash chip 770d previously served as the PW chip, the data it contains will not be unique because the data had previously been written to the previous SW as well. Accordingly, the NAND flash chip 770d can receive duplicate data from the NAND flash chip 770f in preparation for being designated SW in a subsequent stage of operation.

[0074] Additional stages can progress in this manner indefinitely to spread wear among some or all of the available NAND flash chips 770. In performing these operations, the NAND flash storage device 750 will write each received page of data to the NAND flash chips 770 twice, doubling the amount of wear inflicted; however, when the wear is spread evenly across many NAND flash chips 770, this amount of additional wear may be acceptable.

[0075] Figure 9 is a flowchart showing a method 900 of using a NAND flash storage device with a rotating NAND buffer for wear leveling, according to an implementation consistent with the claims. For example, the method 900 can be performed by the NAND flash storage device 750 previously described with reference to Figure 7. The method 900 can rotate through available NAND flash chips, designating each in turn as a buffer NAND chip, in order to distribute the wear caused by executing write operations and avoid failures potentially caused by focusing write operations on a small number of NAND flash chips. The method 900 includes designating a first NAND flash chip as a secondary write (SW) chip, a second NAND flash chip as a future secondary write (FSW) chip, a third NAND flash chip as a primary write (PW) chip, and a fourth NAND flash chip as a future primary write (FPW) chip (stage 910). The method 900 includes receiving data over an input/output (I/O) bus (stage 920). The method 900 includes writing the received data to the PW chip and the SW chip successively such that write operations to each chip do not overlap in time (stage 930). The method 900 includes receiving a read command over the I/O bus (stage 940). The method 900 includes, responsive to receiving the read command, reading data from a NAND flash chip that is not currently executing a write operation such that read operations are not queued behind write operations (stage 950). The method 900 includes copying valid data from the FPW chip to the FSW chip (stage 960). The method 900 includes, following the writing of the received data and the copying of the valid data: designating the first NAND flash chip as read-only, the second NAND flash chip as the SW

chip, the third NAND flash chip as the FSW chip, and the fourth NAND flash chip as the PW chip (stage 970).

[0076] The method 900 includes designating the first NAND flash chip as the SW chip, the second NAND flash chip as the FSW chip, the third NAND flash chip as the PW chip, and the fourth NAND flash chip as the FPW chip (stage 910). The designations were described previously with reference to Figure 8.

[0077] The method 900 includes receiving data over the I/O bus (stage 920). Stage 920 is similar to stages 310 and 610 described previously.

[0078] The method 900 includes writing the received data to the PW chip and the SW chip successively such that write operations to each chip do not overlap in time (stage 930). Stage 930 is similar to stages 320 and 620 described previously. In some implementations, the method 900 may continue performing receive operations of stage 920 and the write operations of stage 930 until the PW chip and the SW chip are full or substantially full before advancing to stage 950. The operations of stages 940 and 950, described below, can be performed concurrently with the operations of stages 920 and 930.

[0079] The method 900 includes receiving a read command over the I/O bus (stage 940). Stage 940 is similar to stages 330 and 640 described previously.

[0080] The method 900 includes, responsive to receiving the read command, reading data from a NAND flash chip that is not currently executing a write operation such that read operations are not queued behind write operations (stage 950). Stage 950 is similar to stages 340 and 650 previously described. In practice, the requested data may reside on a RO chip, one of the PW or SW chips, or one of the FPW or FSW chips. The controller may handle the read request in one of three ways depending on where the requested data resides. If the requested data is on the PW or SW chip, the controller can execute the read operation on whichever of the PW or SW chips is not currently executing a write operation such that read operations are not queued behind write operations, and the method 600 can proceed to the next stage. If the requested data is on a RO chip, the controller can simply execute the read operation on the RO chip and returns the data—there is no write operation occurring that would interfere with or delay execution of the read operation. If the requested data is on the FPW or FSW chip, the controller can execute the read operation on whichever of the FPW or FSW is not currently executing a write operation

such that read operations are not queued behind write operations. That is, if data is currently being copied with garbage collection from the FPW chip to the FSW chip, the controller 460 will be executing write operations on the FSW chip, and so can execute the read operation on the FPW chip. When the data is being copied back from the FSW chip to the FPW chip, the controller can execute the read operation on the FSW chip. In this manner, regardless of which NAND flash chip contains the requested data and which NAND flash chip[s] is/are executing write operations, the NAND flash device can read and return the requested data without undue delay caused by queueing the read operation behind a write operation. In this manner, stages 940 and 950 can proceed concurrently with stages 920 and 930, and read operations can be processed promptly without being queued behind slower write operations. The stages 920 through 950 can repeat as additional data and read commands are received over the I/O bus.

[0081] The method 900 includes copying valid data from the FPW chip to the FSW chip (stage 960). To prepare two new NAND flash chips to serve as buffer NAND chips in a later stage, the method 900 can copy data from one NAND flash chip to another NAND flash chip. The copy source can be designated the FPW chip and the copy destination can be designated the FSW chip. In some implementations, the copy operation can be bidirectional; that is, data can be harmonized between the FPW chip and the FSW chip such that they each contain a duplicate of the other's data. During the copying operation, the FSW chip and the FPW chip can perform garbage collection operations to consolidate data and free space of data that is invalid or no longer needed. The copying operations of stage 960 can be performed during an overlapping time period with the buffering operations of stages 920 and 930. In addition, read commands can be received and executed by whichever of the FPW and the FSW is not currently executing write operations as part of the copy operations. In this manner, the copying operations of stage 960 can also be performed during an overlapping time period with the read operations of stages 940 and 950.

[0082] The method 900 includes, following the writing of the received data and the copying of the valid data: designating the first NAND flash chip as a read-only chip, the second NAND flash chip as the SW chip, the third NAND flash chip as the FSW chip, and the fourth NAND flash chip as the PW chip (stage 970). Once the FPW chip and the FSW chip contain duplicate data, they will be ready to take over as the PW chip and the SW chip, respectively. As long as the copying operations can be performed while the buffer NAND flash chips are still handling

write operations, the NAND flash storage device 770 can switch designations without delay and continue executing write operations without a significant loss in write bandwidth.

[0083] In some implementations, the rotation can continue. For example, the controller 760 can designate a fifth NAND flash chip 770e as the FPW chip. The controller 760 can write second data received over the I/O bus 730 to the PW chip and the SW chip such that write operations to each chip do not overlap in time. The controller 760 can copy second valid data from the FPW chip to the FSW chip. And, following the writing of the second received data and the copying of the second valid data, the controller 760 can designate the second NAND flash chip as read-only, the third NAND flash chip as the SW chip, the fourth NAND flash chip as the FSW chip, the fifth NAND flash chip as the PW chip, and a sixth NAND flash chip as the FPW chip.

[0084] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular implementations of particular inventions. Certain features that are described in this specification in the context of separate implementations can also be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation can also be implemented in multiple implementations separately or in any suitable sub-combination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a sub-combination or variation of a sub-combination.

[0085] References to "or" may be construed as inclusive so that any terms described using "or" may indicate any of a single, more than one, and all of the described terms. The labels "first," "second," "third," and so forth are not necessarily meant to indicate an ordering and are generally used merely to distinguish between like or similar items or elements.

[0086] Various modifications to the implementations described in this disclosure may be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other implementations without departing from the spirit or scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to

07 10 19

be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

[0087] What is claimed is:

CLAIMS

1. A NAND flash storage device comprising:

- a plurality of NAND flash chips; and
- a controller configured to:
 - designate a first NAND flash chip as a secondary write chip, a second NAND flash chip as a future secondary write chip, a third NAND flash chip as a primary write chip, a fourth NAND flash chip as a future primary write chip, and any remaining NAND flash chips as read-only;
 - receive data over an input/output (I/O) bus;
 - write the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
 - receive a read command over the I/O bus; and
- responsive to receiving the read command, read data from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations
- copy valid data from the future primary write chip to the future secondary write chip; and
- following the writing of the received data and the copying of the valid data, designate the first NAND flash chip as read-only, the second NAND flash chip as the secondary write chip, the third NAND flash chip as the future secondary write chip, and the fourth NAND flash chip as the primary write chip.

2. The NAND flash storage device of claim 1, wherein the controller is configured to:

- following the writing of the received data and the copying of the valid data:
 - designate a fifth NAND flash chip as the future primary write chip;
 - write second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
 - copy second valid data from the future primary write chip to the future secondary write chip; and
- following the writing of the second received data and the copying of the second valid data: designate the second NAND flash chip as read-only, the third NAND flash chip as the secondary write chip, the fourth NAND flash chip as the future secondary write chip, and the fifth NAND flash chip as the primary write chip.

9
10
06
21

3. The NAND flash storage device of claim 1, wherein the copy operation includes:

- copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip;
- copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip;
- and
- erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

4. The NAND flash storage device of claim 1, 2 or 3, wherein:

- the write operation and the copy operation are performed during overlapping time periods.

5. A NAND flash storage device comprising:

- a plurality of buffer NAND flash chips;
- a plurality of storage NAND flash chips; and
- a controller configured to:
 - designate a first storage NAND flash chip as a primary write chip, a second storage NAND flash chip as a future primary write chip, a first buffer NAND flash chip as a secondary write chip, a second buffer NAND flash chip as a future secondary write chip, and any remaining NAND flash chips as read-only;
 - receive incoming data over an input/output (I/O) bus;
 - write the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
 - receive a read command over the I/O bus; and
 - responsive to receiving the read command, read data from whichever of the first buffer NAND flash chip or the first storage NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations; and
- copy valid data from the future primary write chip to the future secondary write chip; and

21 06 19

following the writing of the received data and the copying of the valid data, designate the first storage NAND flash chip as read-only, the second storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the future secondary write chip, and the second buffer NAND flash chip as the secondary write chip.

6. The NAND flash storage device of claim 5, wherein the controller is configured to:
 - following the writing of the received data and the copying of the valid data:
 - designate a third storage NAND flash chip as the future primary write chip;
 - write second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
 - copy second valid data from the future primary write chip to the future secondary write chip; and
 - following the writing of the second received data and the copying of the second valid data, designate the second storage NAND flash chip as read-only, the third storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the secondary write chip, and the second buffer NAND flash chip as the future secondary write chip.
7. The NAND flash storage device of claim 5, wherein the copy operation includes:
 - copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip;
 - copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip;
 - and erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.
8. The NAND flash storage device of any one of claims 5 to 7, wherein:
 - the buffer NAND flash chips are single-level cell (SLC) NAND flash chips; and
 - the storage NAND flash chips are triple-level cell (TLC) NAND flash chips.

21 06 19

9. A method of buffering a NAND flash storage device comprising:

- designating a first NAND flash chip as a secondary write chip, a second NAND flash chip as a future secondary write chip, a third NAND flash chip as a primary write chip, a fourth NAND flash chip as a future primary write chip, and any remaining NAND flash chips as read-only
- receiving data over an input/output (I/O) bus;
- writing the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
- receiving a read command over the I/O bus; and
- responsive to receiving the read command, reading data from whichever of the first NAND flash chip or the second NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations
- copying valid data from the future primary write chip to the future secondary write chip;
- and

following the writing of the received data and the copying of the valid data: designating the first NAND flash chip as read-only, the second NAND flash chip as the secondary write chip, the third NAND flash chip as the future secondary write chip, and the fourth NAND flash chip as the primary write chip.

10. The method of claim 9, comprising:

- following the writing of the received data and the copying of the valid data:
 - designating a fifth NAND flash chip as the future primary write chip;
 - writing second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
 - copying second valid data from the future primary write chip to the future secondary write chip; and
- following the writing of the second received data and the copying of the second valid data, designating the second NAND flash chip as read-only, the third NAND flash chip as the secondary write chip, the third NAND flash chip as the further secondary write chip, the fourth NAND flash chip as the primary write chip, and the fifth NAND flash chip as the future primary write chip.

21 06 19

11. The method of claim 9, wherein the copy operation includes:

- copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip;
- copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip;
- and
- erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

12. The method of claim 9, 10 or 11, wherein:

- the writing operation and the copying operation are performed during overlapping time periods.

13. A method of buffering a NAND flash storage device comprising:

- designating a first storage NAND flash chip as a primary write chip, a second storage NAND flash chip as a future primary write chip, a first buffer NAND flash chip as a secondary write chip, a second buffer NAND flash chip as a future secondary write chip, and any remaining NAND flash chips as read-only;
- receiving data over an input/output (I/O) bus;
- writing the received data to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;
- receiving a read command over the I/O bus; and
- responsive to receiving the read command, reading data from whichever of the first buffer NAND flash chip or the first storage NAND flash chip is not currently executing a write operation such that read operations are not queued behind write operations; and
- copying valid data from the future primary write chip to the future secondary write chip;
- and
- following the writing of the received data and the copying of the valid data: designating the first storage NAND flash chip as read-only, the second storage NAND flash chip

21 06 19

as the primary write chip, the first buffer NAND flash chip as the future secondary write chip, and the second buffer NAND flash chip as the secondary write chip.

14. The method of claim 13, comprising:

following the writing of the received data and the copying of the valid data:

designating a third storage NAND flash chip as the future primary write chip;
writing second data received over the I/O bus to the primary write chip and the secondary write chip successively such that write operations to each chip do not overlap in time;

copying second valid data from the future primary write chip to the future secondary write chip; and

following the writing of the second received data and the copying of the second valid data, designating the second storage NAND flash chip as read-only, the third storage NAND flash chip as the primary write chip, the first buffer NAND flash chip as the secondary write chip, and the second buffer NAND flash chip as the future secondary write chip.

15. The method of claim 13, wherein the copy operation includes:

copying the valid data from the future primary write chip to the future secondary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future secondary write chip;

copying the valid data from the future secondary write chip back to the future primary write chip such that the valid data and the free space, respectively, occupy substantially contiguous respective memory spaces of the future primary write chip;
and

erasing the memory spaces of the future primary write chip and the future secondary write chip occupied by the free space.

16. The method of claim 13, 14 or 15, wherein:

the buffer NAND flash chips are single-level cell (SLC) NAND flash chips; and
the storage NAND flash chips are triple-level cell (TLC) NAND flash chips.

21 06 19