

(51) International Patent Classification:
G06F 13/16 (2006.01)(21) International Application Number:
PCT/US2017/018888(22) International Filing Date:
22 February 2017 (22.02.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/189,684 22 June 2016 (22.06.2016) US(71) Applicant: SANDISK TECHNOLOGIES LLC
[US/US]; 6900 North Dallas Parkway, Suite 325, Plano, TX
75024 (US).(72) Inventors: SHAHARABANY, Amir; Gilad 4 Street,
Kochav Yair (IL). MARKUS, Yoav; 23 Helsingy Street,
Tel Aviv (IL). HELLER, Tal; Alonay Abba, PO 192 (IL).
OSHINSKY, Hadas; Feirberg 1, Kfar Saba (IL).(74) Agent: HETZ, Joseph, F.; Brinks Gilson & Lione, P.O.
Box 10087, Chicago, IL 60610 (US).(81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR,
KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.(84) Designated States (*unless otherwise indicated, for every
kind of regional protection available*): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: STORAGE SYSTEM AND METHOD FOR BURST MODE MANAGEMENT USING TRANSFER RAM




FIG. 1A




FIG. 1B




FIG. 1C

(57) Abstract: A storage system uses consumption of transfer RAM as a trigger to enter and exit burst mode. In one embodiment, the storage system stores, in volatile memory, data to be written in non-volatile memory; monitors an allocation level of the volatile memory to determine a first amount of time that the allocation level is at a first level and a second amount of time that the allocation level is at second level; enters burst mode when a ratio of the first amount of time and the second amount of time is above a first threshold; and exits burst mode when the ratio of the first amount of time and the second amount of time is below a second threshold. Other embodiments are possible, and each of the embodiments can be used alone or together in combination.

Storage System and Method for Burst Mode Management Using Transfer RAM

Background

[0001] Storage systems may be used in different conditions, which place different performance requirements on the memory in the storage system. To account for these different conditions, memories in some storage systems can be operated in a “normal” mode or in a “burst” mode, in which a higher-than-average performance by the storage system is required to satisfy a higher-than-average number of write commands from a host controller. For example, in some storage systems, the burst mode may be capable of supporting a storage rate of 160 MB/s for 4 GB (e.g., when raw video shot by a camera is sent to the storage system for storage), while the normal mode may support rates of up to 80 MB/s for 12 GB. To achieve the higher-than-average performance and provide high throughput in burst mode, the storage system can store data in single-level cells (SLCs) instead of multi-level cells (MLCs), as writing to SLC cells (X1) is faster than writing to MLC cells (e.g., X3).

[0002] To trigger burst mode, a host can send a request or other indication to signal that the storage system should switch to burst mode. However, in many environments, the host does not provide such an indication, and the storage system needs to manage the transition to and from burst mode autonomously. A storage system may infer burst mode by monitoring the data transfer rate over an interface from the host, and if the throughput is above a first threshold, a transition to burst mode may be performed. Similarly, terminating burst mode and transitioning back to normal mode may be done when the transfer rate is below a second threshold. As another technique, the storage system may “sniff” the interface and infer from a header of incoming data if the data is of a type that is suitable for burst mode (e.g., video or data in an mpeg format).

Brief Description of the Drawings

[0003] Figure 1A is a block diagram of a non-volatile storage system of an embodiment.

[0004] Figure 1B is a block diagram illustrating an exemplary storage module of an embodiment.

[0005] Figure 1C is a block diagram illustrating a hierarchical storage system of an embodiment.

[0006] Figure 2A is a block diagram illustrating exemplary components of the controller of the non-volatile storage system illustrated in Figure 1A according to an embodiment.

[0007] Figure 2B is a block diagram illustrating exemplary components of the non-volatile memory storage system illustrated in Figure 1A according to an embodiment.

[0008] Figure 3 is a transfer RAM allocation table of an embodiment.

[0009] Figure 4 is a fill-level indicator of an embodiment.

[0010] Figure 5 is flow chart of a method of an embodiment for burst mode management using transfer RAM.

[0011] Figure 6 is a graph of an embodiment of first burst activity versus time.

[0012] Figure 7 is a graph of an embodiment of second burst activity versus time.

Detailed Description

[0013] Overview

[0014] By way of introduction, the below embodiments relate to a storage system and method for burst mode management using transfer RAM. In one embodiment, a method is provided that is performed in a storage system comprising a non-volatile memory and a volatile memory. The method comprises storing, in the volatile memory, data to be written in the non-volatile memory; monitoring an allocation level of the volatile memory to determine a first amount of time that the allocation level of the volatile memory is at a first level and a second amount of time that the allocation level of the volatile memory is at second level; entering burst mode when a ratio of the first amount of time and the second amount of time is above a first threshold; and exiting burst mode when the ratio of the first amount of time and the second amount of time is below a second threshold.

[0015] In some embodiments, the first level is when the non-volatile memory is full, and the second level is when the non-volatile memory is less than full.

[0016] In some embodiments, the method further comprises storing, in the volatile memory, data read from the non-volatile memory, wherein the allocation level of the

volatile memory accounts for both data to be stored in the non-volatile memory and data read from the non-volatile memory.

[0017] In some embodiments, after entering burst mode, the storage system stores the data from the volatile memory in single-level cells in the non-volatile memory, and wherein after exiting burst mode, the storage system uses multi-level cells in the non-volatile memory.

[0018] In some embodiments, the non-volatile memory comprises a three-dimensional memory.

[0019] In some embodiments, the storage system is embedded in a host.

[0020] In some embodiments, the storage system is removably connected to a host.

[0021] In another embodiment, a storage system is provided comprising a first memory; a second memory configured to temporarily store data before it is stored in the first memory; and a controller in communication with the first and second memories. The controller is configured to determine a fill level of the second memory and enter or exit burst mode based on the fill-level of the second memory.

[0022] In some embodiments, the first memory comprises a non-volatile memory and the second memory comprises a volatile memory.

[0023] In some embodiments, the controller is configured to determine the fill level of the second memory by determining how much of the second memory has been allocated.

[0024] In some embodiments, the controller is further configured to determine a first amount of time that the fill level of the second memory is at a first level and a second amount of time that the fill level of the second memory is at second level, wherein the controller is configured to enter burst mode when a ratio of the first amount of time and the second amount of time is above a first threshold, and wherein the controller is configured to exit burst mode when the ratio of the first amount of time and the second amount of time is below a second threshold.

[0025] In some embodiments, the first level is when the second memory is full, and wherein the second level is when the second memory is less than full.

[0026] In some embodiments, the second memory is further configured to temporarily store data read from the first memory, and wherein the fill level of the second memory is due to both data to be stored in the first memory and data read from the first memory.

[0027] In some embodiments, the controller is further configured to store the data from the second memory in single-level cells in the first memory after entering burst mode and use multi-level cells in the first memory after exiting burst mode.

[0028] In some embodiments, the first memory comprises a three-dimensional memory.

[0029] In some embodiments, the storage system is embedded in a host.

[0030] In some embodiments, the storage system is removably connected to a host.

[0031] In another embodiment, a storage system is provided comprising a memory; a transfer memory configured to hold data before the data is sent to the memory for storage; and means for using consumption of the transfer memory as a trigger to enter and exit burst mode.

[0032] In some embodiments, the transfer memory is further configured to hold data read from the memory.

[0033] In some embodiments, the means for using comprises a controller.

[0034] In some embodiments, the memory comprises a three-dimensional memory.

[0035] In some embodiments, the storage system is embedded in a host.

[0036] In some embodiments, the storage system is removably connected to a host.

[0037] Other embodiments are possible, and each of the embodiments can be used alone or together in combination. Accordingly, various embodiments will now be described with reference to the attached drawings.

[0038] Exemplary Embodiments

[0039] Storage systems suitable for use in implementing aspects of these embodiments are shown in Figures 1A -1C. Figure 1A is a block diagram illustrating a non-volatile storage system 100 according to an embodiment of the subject matter described herein. Referring to Figure 1A, non-volatile storage system 100 includes a controller 102 and non-volatile memory that may be made up of one or more non-volatile memory die 104. As used herein, the term die refers to the collection of non-volatile memory cells, and associated circuitry for managing the physical operation of those non-volatile memory cells, that are formed on a single semiconductor substrate. Controller 102 interfaces with a host system and transmits command sequences for read, program, and erase operations to non-volatile memory die 104.

[0040] The controller 102 (which may be a flash memory controller) can take the form of processing circuitry, a microprocessor or processor, and a computer-readable medium that stores computer-readable program code (e.g., firmware) executable by the (micro)processor, logic gates, switches, an application specific integrated circuit (ASIC), a programmable logic controller, and an embedded microcontroller, for example. The controller 102 can be configured with hardware and/or firmware to perform the various functions described below and shown in the flow diagrams. Also, some of the components shown as being internal to the controller can also be stored external to the controller, and other components can be used. Additionally, the phrase “operatively in communication with” could mean directly in communication with or indirectly (wired or wireless) in communication with through one or more components, which may or may not be shown or described herein.

[0041] As used herein, a flash memory controller is a device that manages data stored on flash memory and communicates with a host, such as a computer or electronic device. A flash memory controller can have various functionality in addition to the specific functionality described herein. For example, the flash memory controller can format the flash memory to ensure the memory is operating properly, map out bad flash memory cells, and allocate spare cells to be substituted for future failed cells. Some part of the spare cells can be used to hold firmware to operate the flash memory controller and implement other features. In operation, when a host needs to read data from or write data to the flash memory, it will communicate with the flash memory controller. If the host provides a logical address to which data is to be read/written, the flash memory controller can convert the logical address received from the host to a physical address in the flash memory. (Alternatively, the host can provide the physical address.) The flash memory controller can also perform various memory management functions, such as, but not limited to, wear leveling (distributing writes to avoid wearing out specific blocks of memory that would otherwise be repeatedly written to) and garbage collection (after a block is full, moving only the valid pages of data to a new block, so the full block can be erased and reused).

[0042] Non-volatile memory die 104 may include any suitable non-volatile storage medium, including NAND flash memory cells and/or NOR flash memory cells. The

memory cells can take the form of solid-state (e.g., flash) memory cells and can be one-time programmable, few-time programmable, or many-time programmable. The memory cells can also be single-level cells (SLC), multiple-level cells (MLC), triple-level cells (TLC), or use other memory cell level technologies, now known or later developed. Also, the memory cells can be fabricated in a two-dimensional or three-dimensional fashion.

[0043] The interface between controller 102 and non-volatile memory die 104 may be any suitable flash interface, such as Toggle Mode 200, 400, or 800. In one embodiment, storage system 100 may be a card based system, such as a secure digital (SD) or a micro secure digital (micro-SD) card. In an alternate embodiment, storage system 100 may be part of an embedded storage system.

[0044] Although, in the example illustrated in Figure 1A, non-volatile storage system 100 (sometimes referred to herein as a storage module) includes a single channel between controller 102 and non-volatile memory die 104, the subject matter described herein is not limited to having a single memory channel. For example, in some NAND storage system architectures (such as the ones shown in Figures 1B and 1C), 2, 4, 8 or more NAND channels may exist between the controller and the NAND memory device, depending on controller capabilities. In any of the embodiments described herein, more than a single channel may exist between the controller and the memory die, even if a single channel is shown in the drawings.

[0045] Figure 1B illustrates a storage module 200 that includes plural non-volatile storage systems 100. As such, storage module 200 may include a storage controller 202 that interfaces with a host and with storage system 204, which includes a plurality of non-volatile storage systems 100. The interface between storage controller 202 and non-volatile storage systems 100 may be a bus interface, such as a serial advanced technology attachment (SATA) or peripheral component interface express (PCIe) interface. Storage module 200, in one embodiment, may be a solid state drive (SSD), such as found in portable computing devices, such as laptop computers, and tablet computers.

[0046] Figure 1C is a block diagram illustrating a hierarchical storage system. A hierarchical storage system 250 includes a plurality of storage controllers 202, each of which controls a respective storage system 204. Host systems 252 may access memories

within the storage system via a bus interface. In one embodiment, the bus interface may be an NVMe or fiber channel over Ethernet (FCoE) interface. In one embodiment, the system illustrated in Figure 1C may be a rack mountable mass storage system that is accessible by multiple host computers, such as would be found in a data center or other location where mass storage is needed.

[0047] Figure 2A is a block diagram illustrating exemplary components of controller 102 in more detail. Controller 102 includes a front end module 108 that interfaces with a host, a back end module 110 that interfaces with the one or more non-volatile memory die 104, and various other modules that perform functions which will now be described in detail. Modules of the controller 102 may include a RAM allocator 111 that is configured to allocate space in the RAM 116 and to keep track of the fill-level of the RAM 116. The RAM allocator 111 will be discussed in more detail below. It should be noted, though, that a module, circuit, or monitor may take the form of a packaged functional hardware unit designed for use with other components, a portion of a program code (e.g., software or firmware) executable by a (micro)processor or processing circuitry that usually performs a particular function of related functions, or a self-contained hardware or software component that interfaces with a larger system, for example.

[0048] Referring again to modules of the controller 102, a buffer manager/bus controller 114 manages buffers in random access memory (RAM) 116 and controls the internal bus arbitration of controller 102. A read only memory (ROM) 118 stores system boot code. Although illustrated in Figure 2A as located separately from the controller 102, in other embodiments one or both of the RAM 116 and ROM 118 may be located within the controller. In yet other embodiments, portions of RAM and ROM may be located both within the controller 102 and outside the controller.

[0049] Front end module 108 includes a host interface 120 and a physical layer interface (PHY) 122 that provide the electrical interface with the host or next level storage controller. The choice of the type of host interface 120 can depend on the type of memory being used. Examples of host interfaces 120 include, but are not limited to, SATA, SATA Express, SAS, Fibre Channel, USB, PCIe, and NVMe. The host interface 120 typically facilitates transfer for data, control signals, and timing signals.

[0050] Back end module 110 includes an error correction controller (ECC) engine 124 that encodes the data bytes received from the host, and decodes and error corrects the data bytes read from the non-volatile memory. A command sequencer 126 generates command sequences, such as program and erase command sequences, to be transmitted to non-volatile memory die 104. A RAID (Redundant Array of Independent Drives) module 128 manages generation of RAID parity and recovery of failed data. The RAID parity may be used as an additional level of integrity protection for the data being written into the memory device 104. In some cases, the RAID module 128 may be a part of the ECC engine 124. A memory interface 130 provides the command sequences to non-volatile memory die 104 and receives status information from non-volatile memory die 104. In one embodiment, memory interface 130 may be a double data rate (DDR) interface, such as a Toggle Mode 200, 400, or 800 interface. A flash control layer 132 controls the overall operation of back end module 110.

[0051] The storage system 100 also includes other discrete components 140, such as external electrical interfaces, external RAM, resistors, capacitors, or other components that may interface with controller 102. In alternative embodiments, one or more of the physical layer interface 122, RAID module 128, media management layer 138 and buffer management/bus controller 114 are optional components that are not necessary in the controller 102.

[0052] Figure 2B is a block diagram illustrating exemplary components of non-volatile memory die 104 in more detail. Non-volatile memory die 104 includes peripheral circuitry 141 and non-volatile memory array 142. Non-volatile memory array 142 includes the non-volatile memory cells used to store data. The non-volatile memory cells may be any suitable non-volatile memory cells, including NAND flash memory cells and/or NOR flash memory cells in a two dimensional and/or three dimensional configuration. Peripheral circuitry 141 includes a state machine 152 that provides status information to controller 102.

[0053] In one embodiment, the storage system 100 can be used in a “normal” mode or in “burst” mode. In burst mode, the storage system 100 provides higher-than-average performance and higher-than-average throughput by storing data in single-level cells (SLCs) in the memory 104 instead of multi-level cells (MLCs), as writing to SLC cells

(X1) is faster than writing to MLC cells (e.g., X3). By way of example only and without intending any limitation on the claims, in one embodiment, the storage system 100 may be capable of supporting a storage rate of 160 MB/s for 4 GB (e.g., when raw video shot by a camera is sent to the storage system for storage) while in burst mode, while the normal mode may support rates of up to 80 MB/s for 12 GB. Optionally, the memory management software/firmware of the storage module 100 can be configured to avoid garbage collection operations during burst mode, as the number of memory cells available to store data is temporarily reduced during the garbage collection process. Also optionally, less-robust ECC parameters (e.g., the number of parity bits used, the minimum Hamming distance used, etc.) can be used to increase write performance in burst mode, or ECC may not even be used as all in burst mode. It should be noted that while the below embodiments will use the term “burst mode,” a single burst mode (with a certain storage rate) can be used or a plurality of different burst modes (with different storage rates) can be used. In this way, the storage system 100 can choose a burst mode that fits a desired performance.

[0054] There are several ways in which burst mode can be triggered. For example, a host can send a request or other indication to signal that the storage system should switch to burst mode. However, in many environments, the host does not provide such an indication, and the storage system 100 would need to manage the transition to and from burst mode autonomously. The storage system 100 may detect burst mode by monitoring the data transfer rate over an interface from the host, and if the throughput is above a first threshold, a transition to burst mode may be performed. Similarly, terminating burst mode and transitioning back to normal mode may be done when the transfer rate is below a second threshold. As another technique, the storage system 100 may “sniff” the interface and infer from a header of incoming data if the data is of a type that is suitable for burst mode (e.g., video or data in an mpeg format).

[0055] One issue with these approaches is that they only consider what is coming from the host and not the storage limitations of the RAM 116. In one embodiment, the RAM 116 acts as a temporary storage unit for data (e.g., received from the host) to be written to the non-volatile memory 104 and also as a temporary storage unit for data read from the non-volatile memory 104 (e.g., before sending the data to the host). Because of this, the

RAM 116 is sometimes referred to herein as “transfer RAM,” “cache RAM,” or “buffer RAM”. However, it should be understood that another type of storage unit (e.g., a non-volatile memory) can be used to temporarily store data. But, for simplicity, the RAM 116 will be used to illustrate this embodiment.

[0056] The problem with burst activity is that the RAM 116 can fill up faster than the rate at which the data is transferred from RAM 116 to the non-volatile memory 104 for storage. In this situation, incoming data from the host will be prevented until there is available space in the RAM 116, which can degrade performance. By entering burst mode, data is transferred out of RAM 116 faster, so there is room to accept incoming data. In prior trigger approaches that only consider what is coming from the host, there may be data in the RAM 116 that was read from the memory 104 and consuming space in the RAM 116. That is, those prior approaches were not designed for mix mode (i.e., write and read commands interleaved), as they only consider write operations’ impact on the RAM 116. The read operations confuse the detection, even though they are part of the burst.

[0057] To address this situation, the following embodiments use consumption of the transfer memory (e.g., RAM 116) as a trigger to enter and exit burst mode. Since the transfer memory is an input/output data reservoir – both for data to be stored in the memory 104 and data read from the memory 104, its level (amount of stored data) is an appropriate indication for entering or exiting burst mode. Consider, for example, the situation in which the transfer RAM 116 has a size of 128 KB, and the transfer RAM allocator 111 (Figure 2A) manages the allocations from the transfer RAM 116. As shown in Figure 3, in one embodiment, the transfer RAM allocator 111 can store a table (or other data structure) that lists the addresses in the transfer RAM 116 and whether or not that address is allocated. In another embodiment, which is shown in Figure 4, the transfer RAM allocator 111 can have an indicator for the allocated amount of RAM 116 that shows the transfer RAM allocation level. Of course, these are just examples, and other mechanisms used to judge the allocation, fill-level, or consumption of the transfer RAM 116 can be used.

[0058] In general, if the transfer RAM allocation level is equal to the size of the transfer RAM 116 (i.e., the transfer RAM 116 is full) for a long period of time, the

storage controller 102 is experiencing a burst condition, since whenever the transfer RAM 116 is deallocated, it is immediately reallocated for the next command. On the other hand, if the transfer RAM allocation level is less than the size of the transfer RAM 116 (i.e., the transfer RAM 116 is not full) for a long period of time, the storage controller 102 is in a non-burst condition. Monitoring the allocation level indicator allow the storage controller 102 to decide whether or not a burst condition exists and, hence, whether or not the storage system 100 should go into burst mode.

[0059] In one embodiment, the storage system 100 generates two event triggers: a transfer RAM full event and a transfer RAM not-full event. These events can be defined in any desirable way. In one embodiment, a transfer RAM full event occur when the transfer RAM 116 is completely allocated, and the transfer RAM not-full event occurs after the first chunk of data (e.g., 4K) is evicted from the transfer RAM 116. The controller 102 can use these events to enter or exit burst mode.

[0060] In another embodiment, the storage system 100 can be configured to regularly accumulate the time between the transfer RAM full event and the transfer RAM not-full event (the numerator) within a time window (e.g., 50 milliseconds). Similarly, the storage system 100 can be configured to accumulate the complementary time between transfer RAM not-full event to the next transfer RAM full event (the denominator) in the same time window. In this embodiment, the ratio between these values (the transfer RAM full time divided by the transfer RAM not-full time) can be used as the trigger for entering and exiting burst mode. This embodiment will be described in more detail below in conjunction with the flow chart 500 of Figure 5.

[0061] As shown in Figure 5, the storage system 100 (e.g., the controller 102 generally or the RAM allocator 111 in particular) monitors an allocation level of the RAM 116 (act 510). As discussed above, this monitoring can be performed in any suitable way, such as, but not limited to, using an allocation table or an allocation fill-level indicator. Next, the storage system 100 determines an amount of time that the allocation level is at a first level (act 520). In this embodiment, the “first level” is the transfer RAM full level (i.e., when the transfer RAM 116 is completely allocated). However, in other embodiments, this “first level” can be set at a different level (e.g., 95% of full capacity). The storage system 100 also determines an amount of time that the allocation level is at a second level

(act 530). In this embodiment, the “second level” is any level other than the transfer RAM full level (i.e., even when the smallest “chunk” is available in the transfer RAM 115). In other embodiments, this “second level” can be set at a different level (e.g., 95% of full capacity).

[0062] Figure 6 is a graph showing these measurements for exemplary write activity. As shown in this timing diagram, a TRAM full event occurs when the transfer RAM 116 becomes full, and a TRAM not-full event when the transfer RAM 116 is partially evicted. In this example, t1, t3, t5, t7, and t9 are the times that the allocation level is at a first level, and t2, t4, t6, t8, and t10 are the times that the allocation level is at a second level. So, in acts 520 and 530, the sum of the TRAM full times and the sum of TRAM not-full times are accumulated in a time window. Figure 7 is another example, where the transfer RAM 116 was filled completely and then the rate of data transferred for writing to the storage system 100 was reduced below the rate of evicting data, such that the transfer RAM 116 will not be full. (In this example, the ratio of full time versus not-full time will get smaller, and burst mode will be terminated.)

[0063] Returning to Figure 5, the storage system 100 next computes a ratio of the two times (act 540), and compares the ratio to two thresholds, which can be pre-defined based on performance parameters of the storage system 100 (acts 550 and 570). In this way, the ratios can act as an indication to enter and exist burst mode programming. The first threshold can be defined for entering burst mode, and the second threshold can be defined for exiting burst mode. So, as shown in Figure 5, if the ratio is greater than the first threshold (act 550), the storage system 100 enters burst mode (act 560). If the ratio is less than the second threshold (act 570), the storage system 100 exits burst mode (act 580). Put in mathematical terms, the storage system 100 enters burst mode when $\sum t \text{ TRAM full time} / \sum t \text{ TRAM not full time} > \text{threshold 1}$, and the storage system 100 exit burst mode when $\sum t \text{ TRAM full time} / \sum t \text{ TRAM not full time} < \text{threshold 2}$.

[0064] There are many advantages associated with these embodiments. For example, these embodiments can provide better performance in burst mode and better endurance of SLC blocks since burst mode will be entered and exited more properly according to host requirements and the size of the transfer RAM 116 in the storage system 100. Further, these embodiments are relatively simple to implement and require low firmware

overhead (e.g., CPU time). This simplicity can also lend these embodiments to be implemented with hardware. Further, these embodiments are suitable for a “mix mode” in which there are read and write commands interleaved in the burst. This is in contrast to prior burst mode detection mechanisms that use bandwidth as a trigger. In summary, these embodiments can provide better performance and endurance tailored to the transfer RAM 116 size.

[0065] Finally, as mentioned above, any suitable type of memory can be used. Semiconductor memory devices include volatile memory devices, such as dynamic random access memory (“DRAM”) or static random access memory (“SRAM”) devices, non-volatile memory devices, such as resistive random access memory (“ReRAM”), electrically erasable programmable read only memory (“EEPROM”), flash memory (which can also be considered a subset of EEPROM), ferroelectric random access memory (“FRAM”), and magnetoresistive random access memory (“MRAM”), and other semiconductor elements capable of storing information. Each type of memory device may have different configurations. For example, flash memory devices may be configured in a NAND or a NOR configuration.

[0066] The memory devices can be formed from passive and/or active elements, in any combinations. By way of non-limiting example, passive semiconductor memory elements include ReRAM device elements, which in some embodiments include a resistivity switching storage element, such as an anti-fuse, phase change material, etc., and optionally a steering element, such as a diode, etc. Further by way of non-limiting example, active semiconductor memory elements include EEPROM and flash memory device elements, which in some embodiments include elements containing a charge storage region, such as a floating gate, conductive nanoparticles, or a charge storage dielectric material.

[0067] Multiple memory elements may be configured so that they are connected in series or so that each element is individually accessible. By way of non-limiting example, flash memory devices in a NAND configuration (NAND memory) typically contain memory elements connected in series. A NAND memory array may be configured so that the array is composed of multiple strings of memory in which a string is composed of multiple memory elements sharing a single bit line and accessed as a

group. Alternatively, memory elements may be configured so that each element is individually accessible, e.g., a NOR memory array. NAND and NOR memory configurations are exemplary, and memory elements may be otherwise configured.

[0068] The semiconductor memory elements located within and/or over a substrate may be arranged in two or three dimensions, such as a two dimensional memory structure or a three dimensional memory structure.

[0069] In a two dimensional memory structure, the semiconductor memory elements are arranged in a single plane or a single memory device level. Typically, in a two dimensional memory structure, memory elements are arranged in a plane (e.g., in an x-z direction plane) which extends substantially parallel to a major surface of a substrate that supports the memory elements. The substrate may be a wafer over or in which the layer of the memory elements are formed or it may be a carrier substrate which is attached to the memory elements after they are formed. As a non-limiting example, the substrate may include a semiconductor such as silicon.

[0070] The memory elements may be arranged in the single memory device level in an ordered array, such as in a plurality of rows and/or columns. However, the memory elements may be arrayed in non-regular or non-orthogonal configurations. The memory elements may each have two or more electrodes or contact lines, such as bit lines and word lines.

[0071] A three dimensional memory array is arranged so that memory elements occupy multiple planes or multiple memory device levels, thereby forming a structure in three dimensions (i.e., in the x, y and z directions, where the y direction is substantially perpendicular and the x and z directions are substantially parallel to the major surface of the substrate).

[0072] As a non-limiting example, a three dimensional memory structure may be vertically arranged as a stack of multiple two dimensional memory device levels. As another non-limiting example, a three dimensional memory array may be arranged as multiple vertical columns (e.g., columns extending substantially perpendicular to the major surface of the substrate, i.e., in the y direction) with each column having multiple memory elements in each column. The columns may be arranged in a two dimensional configuration, e.g., in an x-z plane, resulting in a three dimensional arrangement of

memory elements with elements on multiple vertically stacked memory planes. Other configurations of memory elements in three dimensions can also constitute a three dimensional memory array.

[0073] By way of non-limiting example, in a three dimensional NAND memory array, the memory elements may be coupled together to form a NAND string within a single horizontal (e.g., x-z) memory device levels. Alternatively, the memory elements may be coupled together to form a vertical NAND string that traverses across multiple horizontal memory device levels. Other three dimensional configurations can be envisioned wherein some NAND strings contain memory elements in a single memory level while other strings contain memory elements which span through multiple memory levels. Three dimensional memory arrays may also be designed in a NOR configuration and in a ReRAM configuration.

[0074] Typically, in a monolithic three dimensional memory array, one or more memory device levels are formed above a single substrate. Optionally, the monolithic three dimensional memory array may also have one or more memory layers at least partially within the single substrate. As a non-limiting example, the substrate may include a semiconductor such as silicon. In a monolithic three dimensional array, the layers constituting each memory device level of the array are typically formed on the layers of the underlying memory device levels of the array. However, layers of adjacent memory device levels of a monolithic three dimensional memory array may be shared or have intervening layers between memory device levels.

[0075] Then again, two dimensional arrays may be formed separately and then packaged together to form a non-monolithic memory device having multiple layers of memory. For example, non-monolithic stacked memories can be constructed by forming memory levels on separate substrates and then stacking the memory levels atop each other. The substrates may be thinned or removed from the memory device levels before stacking, but as the memory device levels are initially formed over separate substrates, the resulting memory arrays are not monolithic three dimensional memory arrays. Further, multiple two dimensional memory arrays or three dimensional memory arrays (monolithic or non-monolithic) may be formed on separate chips and then packaged together to form a stacked-chip memory device.

[0076] Associated circuitry is typically required for operation of the memory elements and for communication with the memory elements. As non-limiting examples, memory devices may have circuitry used for controlling and driving memory elements to accomplish functions such as programming and reading. This associated circuitry may be on the same substrate as the memory elements and/or on a separate substrate. For example, a controller for memory read-write operations may be located on a separate controller chip and/or on the same substrate as the memory elements.

[0077] One of skill in the art will recognize that this invention is not limited to the two dimensional and three dimensional exemplary structures described but cover all relevant memory structures within the spirit and scope of the invention as described herein and as understood by one of skill in the art.

[0078] 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 the claimed invention. Finally, it should be noted that any aspect of any of the preferred embodiments described herein can be used alone or in combination with one another.

What is claimed is:

1. A method of burst detection, the method comprising:
performing the following in a storage system comprising a non-volatile memory and a volatile memory:
storing, in the volatile memory, data to be written in the non-volatile memory;
monitoring an allocation level of the volatile memory to determine a first amount of time that the allocation level of the volatile memory is at a first level and a second amount of time that the allocation level of the volatile memory is at second level;
entering burst mode when a ratio of the first amount of time and the second amount of time is above a first threshold; and
exiting burst mode when the ratio of the first amount of time and the second amount of time is below a second threshold.
2. The method of Claim 1, wherein the first level is when the non-volatile memory is full, and wherein the second level is when the non-volatile memory is less than full.
3. The method of Claim 1 further comprising:
storing, in the volatile memory, data read from the non-volatile memory;
wherein the allocation level of the volatile memory accounts for both data to be stored in the non-volatile memory and data read from the non-volatile memory.

4. The method of Claim 1, wherein after entering burst mode, the storage system stores the data from the volatile memory in single-level cells in the non-volatile memory, and wherein after exiting burst mode, the storage system uses multi-level cells in the non-volatile memory.
5. The method of Claim 1, wherein the non-volatile memory comprises a three-dimensional memory.
6. The method of Claim 1, wherein the storage system is embedded in a host.
7. The method of Claim 1, wherein the storage system is removably connected to a host.
8. A storage system comprising:
 - a first memory;
 - a second memory configured to temporarily store data before it is stored in the first memory; and
 - a controller in communication with the first and second memories, the controller configured to:
 - determine a fill level of the second memory; and
 - enter or exit burst mode based on the fill-level of the second memory.

9. The storage system of Claim 8, wherein the first memory comprises a non-volatile memory and the second memory comprises a volatile memory.
10. The storage system of Claim 8, wherein the controller is configured to determine the fill level of the second memory by determining how much of the second memory has been allocated.
11. The storage system of Claim 8, wherein the controller is further configured to:
determine a first amount of time that the fill level of the second memory is at a first level and a second amount of time that the fill level of the second memory is at second level;
wherein the controller is configured to enter burst mode when a ratio of the first amount of time and the second amount of time is above a first threshold; and
wherein the controller is configured to exit burst mode when the ratio of the first amount of time and the second amount of time is below a second threshold.
12. The storage system of Claim 11, wherein the first level is when the second memory is full, and wherein the second level is when the second memory is less than full.
13. The storage system of Claim 8, wherein the second memory is further configured to temporarily store data read from the first memory, and wherein the fill level of the second memory is due to both data to be stored in the first memory and data read from the first memory.

14. The storage system of Claim 8, wherein the controller is further configured to store the data from the second memory in single-level cells in the first memory after entering burst mode and use multi-level cells in the first memory after exiting burst mode.
15. The storage system of Claim 8, wherein the first memory comprises a three-dimensional memory.
16. The storage system of Claim 8, wherein the storage system is embedded in a host.
17. The storage system of Claim 8, wherein the storage system is removably connected to a host.
18. A storage system comprising:
 - a memory;
 - a transfer memory configured to hold data before the data is sent to the memory for storage; and
 - means for using consumption of the transfer memory as a trigger to enter and exit burst mode.
19. The storage system of Claim 18, wherein the transfer memory is further configured to hold data read from the memory.

20. The storage system of Claim 18, wherein the means for using comprises a controller.
21. The storage system of Claim 18, wherein the memory comprises a three-dimensional memory.
22. The storage system of Claim 18, wherein the storage system is embedded in a host.
23. The storage system of Claim 18, wherein the storage system is removably connected to a host.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2A




FIG. 2B

ADDR	STATUS
0	ALLOCATED
1	ALLOCATED
2	ALLOCATED
3	FREE
-	
-	
-	

FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2017/018888

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06F13/16
 ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/301763 A1 (SHAHARABANY AMIR [IL] ET AL) 22 October 2015 (2015-10-22) figures 1,5 paragraphs [0023], [0024], [0028], [0033] -----	1-23
A	US 2014/281127 A1 (MARCU ALON [IL] ET AL) 18 September 2014 (2014-09-18) the whole document -----	1-23
A	US 2015/324119 A1 (ROMANOVSKY LEON [IL] ET AL) 12 November 2015 (2015-11-12) the whole document -----	1-23



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

3 May 2017

Date of mailing of the international search report

11/05/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer

van der Meulen, E

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/018888

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015301763 A1	22-10-2015	KR 20160147744 A	23-12-2016
		US 2015301763 A1	22-10-2015
		WO 2015160496 A1	22-10-2015

US 2014281127 A1	18-09-2014	NONE	

US 2015324119 A1	12-11-2015	NONE	
