



(51) International Patent Classification:
G06F 3/06 (2006.01)

(21) International Application Number:
PCT/US2019/063241

(22) International Filing Date:
26 November 2019 (26.11.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/249,522 16 January 2019 (16.01.2019) US

(71) Applicant: **WESTERN DIGITAL TECHNOLOGIES, INC.** [US/US]; 5601 Great Oaks Parkway, San Jose, CA 95119 (US).

(72) Inventors: **ELIASH, Tomer**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **PRIVMAN, Ofer**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **GROSSMAN, Yuval**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **SOKOLOWSKI, Jonathan**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **BAZARSKY, Alexander**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **HAHN, Judah Gamliel**; c/o Western Digital Technologies Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US).

(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, JO, 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: SYSTEM AND METHOD FOR STORAGE SYSTEM PROPERTY DEVIATION




FIG. 1A

(57) Abstract: A system and method for storage system property deviation are provided. In one embodiment, a storage system is presented comprising a memory and a controller configured to: receive a request from a host to modify a write parameter used to write data to the memory; write data to the memory using the modified write parameter; restore the write parameter to its pre-modified state; and re-write the data to the memory in a background operation using the write parameter in its pre-modified state. Other embodiments are provided.

System and Method for Storage System Property Deviation

Background

[0001] Throughput, endurance, and data retention are factors that are often considered in memory design. Since a memory system can be used in many use cases, the system is often designed to provide acceptable throughput, endurance, and data retention for common use cases. While there may be variations between different manufacturers, all storage systems in a given manufacturer's product line usually have the same throughput, endurance, and data retention specifications.

Brief Description of the Drawings

[0002] Figure 1A is a block diagram of a non-volatile storage system of an embodiment.

[0003] Figure 1B is a block diagram illustrating a storage module of an embodiment.

[0004] Figure 1C is a block diagram illustrating a hierarchical storage system of an embodiment.

[0005] Figure 2A is a block diagram illustrating components of the controller of the non-volatile storage system illustrated in Figure 1A according to an embodiment.

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

[0007] Figure 3 is an illustration showing the relationship among endurance, throughput, and data retention of an embodiment.

[0008] Figure 4 is a flow chart of a method of an embodiment for host-controlled storage system property deviation.

[0009] Figure 5 is a flow chart of a method of an embodiment where a host chooses a configuration designed to optimize performance at the expense of endurance and data retention resistance.

[0010] Figure 6 is an illustration showing the relationship among endurance, throughput, and data retention of an embodiment.

[0011] Figure 7 is an illustration showing the relationship among endurance, throughput, and data retention of an embodiment.

Detailed Description

[0012] By way of introduction, the below embodiments relate to a system and method for storage system property deviation. In one embodiment, a storage system is presented comprising a memory and a controller configured to: receive a request from a host to modify a write parameter used to write data to the memory; write data to the memory using the modified write parameter; restore the write parameter to its pre-modified state; and re-write the data to the memory in a background operation using the write parameter in its pre-modified state.

[0013] In some embodiments, the data is marked so it can be identified as being written using the modified write parameter.

[0014] In some embodiments, the modified write parameter improves at least one of throughput, endurance, and data retention at an expense of another one or more of throughput, endurance, and data retention.

[0015] In some embodiments, the request comprises a selection of one of a pre-determined number of different write parameters.

[0016] In some embodiments, the request comprises a requirement, and wherein the controller is further configured to modify the write parameter to satisfy the requirement.

[0017] In some embodiments, the memory comprises a three-dimensional memory.

[0018] In some embodiments, the storage system is removably connectable to the host.

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

[0020] In another embodiment, a method for storage system property deviation is presented that is performed in a storage system comprising a memory. The method comprises altering a write parameter used to write data to the memory; writing data to the memory using the altered write parameter; and re-writing the data to the memory using the write parameter in its original state.

[0021] In some embodiments, the data is re-written in a background operation.

[0022] In some embodiments, the method further comprises receiving a request from a host to alter the write parameter.

[0023] In some embodiments, the request comprises a selection of one of a pre-determined number of different write parameters.

[0024] In some embodiments, the method further comprises marking the data so it can be identified as being written using the altered write parameter.

[0025] In some embodiments, the altered write parameter improves at least one of throughput, endurance, and data retention at an expense of another one or more of throughput, endurance, and data retention.

[0026] In some embodiments, the memory comprises a three-dimensional memory.

[0027] In another embodiment, a storage system is provided comprising a memory; means for writing data to the memory using a modified write parameter; and means for re-writing the data to the memory using the write parameter in its pre-modified state.

[0028] In some embodiments, the storage system further comprises means for marking the data so it can be identified as being written using the modified write parameter.

[0029] In some embodiments, the data is re-written in a background operation.

[0030] In some embodiments, the memory comprises a three-dimensional memory.

[0031] In some embodiments, the storage system is embedded in a host.

[0032] 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.

[0033] Turning now to the drawing, 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. (The terms “memory” and “media” may be used interchangeably herein.) Controller 102 interfaces with a host system and transmits command sequences for read, program, and erase operations to non-volatile memory die 104.

[0034] 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.

[0035] 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).

[0036] 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.

[0037] 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.

[0038] 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.

[0039] 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 interconnect 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.

[0040] 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 a Non-Volatile Memory Express (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.

[0041] Figure 2A is a block diagram illustrating 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. For example, in this embodiment, the controller 102 comprises a properties modifier 111, which will be discussed in more detail below. A module 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 or related functions, or a self-contained hardware or software component that interfaces with a larger system, for example.

[0042] 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.

[0043] 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, Serial Attached SCSI (SAS), Fibre Channel, Universal Serial Bus (USB), PCIe, and NVMe. The host interface 120 typically facilitates transfer for data, control signals, and timing signals.

[0044] Back end module 110 includes an error correction code (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.

[0045] 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.

[0046] Figure 2B is a block diagram illustrating 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. Non-volatile memory die 104 further includes a data cache 156 that caches data.

[0047] Storage systems (sometimes referred to herein as storage devices or just devices) are often required to withstand certain conditions, according to host specifications. These conditions include throughput, quality of service (QoS), endurance, data retention (DR), and others. These requirements are often generic and do not apply to many of users. For example, a gamer who uses the storage device often and replaces his entire computer every few years would have little use of the device's data retention withstanding capabilities. In turn, he may prefer faster response time, better throughput, and QoS. However, when the device is designed, its properties (e.g., read parameters and/or write parameters) are designed to withstand all conditions, without caveats. This

requirement forces the device designers to impose certain limitations that may not apply to all users. In practice, many of the devices in the field do not realize their full potential in view of the individual consumer's needs.

[0048] Flash memory design often imposes a three-way trade-off between major factors, as shown in Figure 3. These factors are throughput, endurance, and data retention. Improving the device in one of these aspects (improving throughput, increasing the program/erase (P/E) cycles the device can withstand, and increasing the period (or temperature) without power the device can withstand) adversely affects the other two aspects. In Figure 3, these design aspects are shown along with the standard requirement fulfillment rates: very good, acceptable, and not acceptable. Previously, the generic requirements were satisfied by all devices that belong to the same product line, sometimes differing by original equipment manufacturing (OEM)-specific requirements. There was no option to modify properties during a single device's lifetime by compromising other properties. The below embodiments present a way to (temporarily or permanently) affect this relationship and improve one of the factors at the expense of the other two (or just one of them).

[0049] In one embodiment, the storage system 100 allows modification of its properties after its design, and a host can (temporarily or permanently) improve one of the device's properties at the expense of (temporarily or permanently) potentially impairing other properties. The idea is based on the fact that many users do not utilize some of the OEM requirements regarding data retention or endurance, while centering on throughput. These users may be gamers that require fast device response time or financial applications using temporary data. Other users may prefer to store data in a "vault mode," which allows longer data retention while writing the data slowly, at the expense of write performance.

[0050] Turning again to the drawings, Figure 4 is a flow chart 400 of a method of an embodiment for host-controlled storage system property deviation. As shown in Figure 4, the host (or the user) assesses the requirements of the current application with respect to its performance, endurance, and data retention (act 410). Next, the host decides whether the requirements are fulfilled using the standard device configuration (act 420). If the requirements are fulfilled, the standard device configuration employing the

standard device properties are used (act 430). Otherwise, the host may supply additional device requirement through some protocol to adapt standard device configuration (act 440). The custom device configuration may be either determined by the device or picked (by either device or host) from a pre-set number of different configurations. In one embodiment, the controller 102 (e.g., the properties modifier 111) takes the input and modifies the properties accordingly.

[0051] The following paragraphs provide two uses cases as example implementations of these embodiments.

[0052] The first use case relates to a “gaming mode,” where it is desired to have high performance at the expense of endurance and/or data retention. In many cases, storage devices are used to store data for a short time, rather than use long term storage. Personal computers (PCs) are usually replaced after several years, and new mobile devices released by vendors every year replace the old models regularly. On the other hand, many consumers need better performance (throughput) to improve frame rate while gaming, shoot longer movies, shoot more pictures, and so on. In this environment, this embodiment can provide users with the hardware they need, customizing it to their needs.

[0053] Figure 5 is a flow chart 500 of a method of an embodiment where a host chooses a configuration designed to optimize performance at the expense of endurance and data retention resistance, and Figure 6 visualizes the modifications made to the properties using this method. Turning first to Figure 5, the host chooses a configuration that was pre-defined to increase performance, according to its current needs (act 510). Next, the storage device 100 applies the configuration to the subsequent host requests, such as write and read (act 520). For example, the configuration (i.e., the write parameter) may apply to the verification strength and the voltage window of the write operation (verify pulse number and voltage window size are in inverse relation to bit error rate (BER)). It may also apply the configuration to different management operations and set its schedule accordingly.

[0054] If the configuration the host selected impaired endurance and data retention significantly, it may be of use to replace the data once the host finished its high performance phase. More specifically, the device 100 can mark data written using a suboptimal configuration (act 530) and then re-write the marked data with an optimal

endurance configuration as part of the management (background) operation (act 540). Figure 6 shows that, by using this embodiment, the throughput has moved to the very good zone, while endurance and data retention are impaired – shown in the acceptable/not-acceptable zones.

[0055] Turning now to the second use case, it may be desired to use the storage system 100 in a “vault mode,” which allows high data retention at the expense of performance and endurance. In this use case, the storage system 100 is used to store data for a long time rather than used for short-term storage. For example, the storage system 100 may be disconnected from a power source and even kept in a physical vault. Advanced external storage devices already have solid state drives (SSDs) inside, and this trend is expected to expand.

[0056] In this mode of storage, there is no use for standard storage demands, such as throughput or endurance (number of P/E cycles), but it is desired to store the data for a long time, requiring good data retention resistance. In such a use case, the host may select a write configuration that prefers a large number of programming and verify pulses and also a larger voltage window. The storage system 100 may also write the data several times, which impairs device endurance and write throughput but improves data retention resistance. Figure 7 visualizes the modification made to the properties. One can see that data retention resistance is improved at the expense of throughput and endurance.

[0057] There are several advantages associated with these embodiments. For example, these embodiments can allow better utilization of a storage device’s properties. Users that do not require some of the device’s properties can be given an option to suit their needs exactly and, by this, gain other improved properties.

[0058] 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.

[0059] 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.

[0060] 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 examples, and memory elements may be otherwise configured.

[0061] 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.

[0062] 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.

[0063] 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 wordlines.

[0064] 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).

[0065] 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.

[0066] 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.

[0067] 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.

[0068] 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.

[0069] 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.

[0070] One of skill in the art will recognize that this invention is not limited to the two dimensional and three dimensional 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.

[0071] 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 embodiments described herein can be used alone or in combination with one another.

What is claimed is:

1. A storage system comprising:
a memory; and
a controller configured to:
receive a request from a host to modify a write parameter used to write data to the memory;
write data to the memory using the modified write parameter;
restore the write parameter to its pre-modified state; and
re-write the data to the memory in a background operation using the write parameter in its pre-modified state.
2. The storage system of Claim 1, wherein the data is marked so that the data can be identified as having been written using the modified write parameter.
3. The storage system of Claim 1, wherein the modified write parameter improves at least one of throughput, endurance, and data retention at an expense of another one or more of throughput, endurance, and data retention.
4. The storage system of Claim 1, wherein the request comprises a selection of one of a pre-determined number of different write parameters.
5. The storage system of Claim 1, wherein:
the request comprises a requirement; and

the controller is further configured to modify the write parameter to satisfy the requirement.

6. The storage system of Claim 1, wherein the memory comprises a three-dimensional memory.
7. The storage system of Claim 1, wherein the storage system is removably connectable to the host.
8. The storage system of Claim 1, wherein the storage system is embedded in the host.
9. A method for storage system property deviation, the method comprising:
performing the following in a storage system comprising a memory:
 altering a write parameter used to write data to the memory;
 writing data to the memory using the altered write parameter; and
 re-writing the data to the memory using the write parameter in its original state.
10. The method of Claim 9, wherein the data is re-written in a background operation.
11. The method of Claim 9, further comprising receiving a request from a host to alter the write parameter.

12. The method of Claim 11, wherein the request comprises a selection of one of a pre-determined number of different write parameters.

13. The method of Claim 9, further comprising marking the data so that the data can be identified as having been written using the altered write parameter.

14. The method of Claim 9, wherein the altered write parameter improves at least one of throughput, endurance, and data retention at an expense of another one or more of throughput, endurance, and data retention.

15. The method of Claim 9, wherein the memory comprises a three-dimensional memory.

16. A storage system comprising:
a memory;
means for writing data to the memory using a modified write parameter; and
means for re-writing the data to the memory using the write parameter in its pre-modified state.

17. The storage system of Claim 16, further comprising means for marking the data so that the data can be identified as having been written using the modified write parameter.

18. The storage system of Claim 16, wherein the data is re-written in a background operation.
19. The storage system of Claim 16, wherein the memory comprises a three-dimensional memory.
20. The storage system of Claim 16, wherein the storage system is embedded in a host.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2A




FIG. 2B




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/063241

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos.: 2,13,17
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
The recited feature of "marking" the data is overly obscure, to the extent that no meaningful search could be conducted in respect to said claims. The supporting disclosure fails to describe how such marking should be implemented, contrary to article 5. In what manner is such marking performed? Is it, for example, by adding actual indication to the data itself? Is it by using some of tracking at database or table?
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/063241

A. CLASSIFICATION OF SUBJECT MATTER IPC (20200101) G06F 3/06 CPC (20130101) G06F 3/061, G06F 3/0629, G06F 3/0634, G06F 3/0607 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) IPC (20200101) G06F 3/06 CPC (20130101) G06F 3/061, G06F 3/0629, G06F 3/0634, G06F 3/0607 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) Databases consulted: Orbit Search terms used: ((SSD OR FLASH)/AB AND (RE?WRITTEN)/AB/CLMS/OBJ/ADB) AND (G06F-003/0634)/CPC		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2016291883 A1 SANDISK TECHNOLOGIES 06 Oct 2016 (2016/10/06) fig. 1; ¶¶2,12,41	1,3-12,14-16,18-20
A	WO 2015185002 A1 LIANG et al. 10 Dec 2015 (2015/12/10) abstract	1,3-12,14-16,18-20
☐ 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 "D" document cited by the applicant in the international application "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 26 Feb 2020		Date of mailing of the international search report 26 Feb 2020
Name and mailing address of the ISA: Israel Patent Office Technology Park, Bldg.5, Malcha, Jerusalem, 9695101, Israel Email address: pctoffice@justice.gov.il		Authorized officer MAUDA Nissim Telephone No. 972-73-3927247

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2019/063241

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
WO 2015185002 A1	10 Dec 2015	WO 2015185002 A1	10 Dec 2015
		CN 105242871 A	13 Jan 2016
		CN 105242871 B	05 Jun 2018
US 2016291883 A1	06 Oct 2016	US 2016291883 A1	06 Oct 2016