



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

(21) International Application Number:
PCT/US2022/028226

(22) International Filing Date:
07 May 2022 (07.05.2022)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
17/471,608 10 September 2021 (10.09.2021) US

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

(72) Inventors: **IONIN, Michael**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, California 95119 (US). **BAZARSKY, Alexander**; c/o Western

Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, California 95119 (US).

(74) Agent: **VER STEEG, Steven, H.** et al.; Patterson & Sheridan, LLP, 24 Greenway Plaza, Suite 1600, Houston, TX 77046 (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, IQ, IR, IS, IT, JM, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(54) Title: SELECTIVE HMB BACKUP IN NVM FOR LOW POWER MODE




FIG. 1

(57) Abstract: A data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitor a health of the memory device, determine that a threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device. The controller is further configured to receive a low power mode indication from the host device and enter the low power mode after copying the XOR parity data from the HMB to the memory device. The controller is further configured to correct read failures using the XOR parity data retrieved from the HMB.

(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))*

Selective HMB Backup In NVM For Low Power Mode**CROSS-REFERENCE TO RELATED APPLICATION(S)**

This application claims the benefit of and hereby incorporates by reference, for all purposes, the entirety of the contents of U.S. Nonprovisional Application No.17/471,608, filed September 10, 2021, and entitled "SELECTIVE HMB BACKUP IN NVM FOR LOW POWER MODE."

BACKGROUND OF THE DISCLOSURE**Field of the Disclosure**

Embodiments of the present disclosure generally relate to data storage devices, such as solid state drives (SSDs), and, more specifically, data recovery methods during low power modes.

Description of the Related Art

Non-volatile memory (NVM) express (NVMe) is based on a paired submission queue and completion queue mechanism. A host device places commands into the submission queue and a controller of a data storage device places completion messages associated with completed commands into the completion queue. In general, submission queues and completion queues are allocated in the memory of a host device, where each queue may be physically located contiguously or non-contiguously in the memory.

Because the amount of random access memory (RAM) in the data storage device is limited, the controller may be conservative about using RAM as a cache for user data (to limit the damage of a sudden power loss) and may store organizational metadata, such as a logical to physical (L2P) table, in the RAM. For fast and consistent performance, data structures are chosen to minimize the amount of computation and number of RAM lookups required at the expense of requiring more RAM. As a common cost reduction, the RAM may be allocated from a shared memory pool located in the host device, instead of adding additional RAM to the data storage device. A host memory buffer (HMB) may be allocated by the host device for use by the data storage device in PCIe/NVMe devices, which leverages the PCIe protocol to allow the controller direct access to the HMB.

The HMB is allocated as a fixed buffer during host device initialization and is fully given over to the data storage device as a dedicated memory location. The data storage device may then directly access the HMB and treat the HMB as an extension of the memory of the data storage device. When the data storage device enters a low power mode, the link between the host device and the controller is shut down. However, the controller may still continue to perform data

management operations in the background (BKOPS) and may require data stored in the HMB. Because the link between the host device and the controller is shut down, the HMB is unavailable to the controller.

Therefore, there is a need in the art for an improved data recovery method when a data storage device is in a low power mode.

SUMMARY OF THE DISCLOSURE

The present disclosure generally relates to data storage devices, such as solid state drives (SSDs), and, more specifically, data recovery methods during low power modes. A data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitor the health of the memory device, determine that a threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device. The controller is further configured to receive a low power mode indication from the host device and enter the low power mode after copying the XOR parity data from the HMB to the memory device. The controller is further configured to correct read failures using the XOR parity data retrieved from the HMB.

In one embodiment, a data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitor health of the memory device, determine that a threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device.

In another embodiment, a data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitoring at least one of a bit error rate (BER) of one or more open blocks of the memory device, a number of program erase cycles (PECs) of the one or more open blocks of the memory device, a temperature variation of the one or more open blocks of the memory device, and an abnormality of one or more open blocks relative to a majority of open blocks of the memory device, determine that a threshold corresponding to the monitoring has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device.

In another embodiment, a data storage device includes memory means and a controller coupled to the memory means. The controller is configured to determine that one or more blocks has

exceeded a temperature threshold or has an abnormality with one or more blocks as compared to a majority of other blocks of the one or more blocks, retrieve XOR parity data from the memory means, perform a read operation to read first data from the one or more blocks or an adjacent block to the one or more blocks, perform error correction on first data with the XOR parity data
5 retrieved from the memory means, determine that the one or more blocks that either exceeded the temperature threshold or had an abnormality no longer exceeds the temperature threshold or no longer has an abnormality, retrieve XOR parity data from a host memory buffer (HMB), perform a read operation to read second data from the one or more blocks or an adjacent block to the one or more blocks, and perform error correction on the second data with the XOR parity
10 data retrieved from HMB.

BRIEF DESCRIPTION OF THE DRAWINGS

So that the manner in which the above recited features of the present disclosure can be understood in detail, a more particular description of the disclosure, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended
15 drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments of this disclosure and are therefore not to be considered limiting of its scope, for the disclosure may admit to other equally effective embodiments.

Figure 1 is a schematic block diagram illustrating a storage system in which a data storage device may function as a storage device for a host device, according to certain embodiments.

20 Figure 2 is an illustration of a collection of graphs showing a failed bit count (FBC) distribution per cycling condition and open block retention, according to certain embodiments.

Figure 3 is an illustration of a link power management flow diagram, according to certain embodiments.

25 Figure 4 is a flow diagram illustrating a method of correcting a read failure using XOR parity data, according to certain embodiments.

Figure 5 is a flow diagram illustrating a method of correcting a read failure using XOR parity data, according to certain embodiments.

To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures. It is contemplated that elements

disclosed in one embodiment may be beneficially utilized on other embodiments without specific recitation.

DETAILED DESCRIPTION

In the following, reference is made to embodiments of the disclosure. However, it should be understood that the disclosure is not limited to specifically described embodiments. Instead, any combination of the following features and elements, whether related to different embodiments or not, is contemplated to implement and practice the disclosure. Furthermore, although embodiments of the disclosure may achieve advantages over other possible solutions and/or over the prior art, whether or not a particular advantage is achieved by a given embodiment is not limiting of the disclosure. Thus, the following aspects, features, embodiments, and advantages are merely illustrative and are not considered elements or limitations of the appended claims except where explicitly recited in a claim(s). Likewise, reference to “the disclosure” shall not be construed as a generalization of any inventive subject matter disclosed herein and shall not be considered to be an element or limitation of the appended claims except where explicitly recited in a claim(s).

The present disclosure generally relates to data storage devices, such as solid state drives (SSDs), and, more specifically, data recovery methods during low power modes. A data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitor a health of the memory device, determine that a threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device. The controller is further configured to receive a low power mode indication from the host device and enter the low power mode after copying the XOR parity data from the HMB to the memory device. The controller is further configured to correct read failures using the XOR parity data retrieved from the HMB.

Figure 1 is a schematic block diagram illustrating a storage system 100 in which a host device 104 is in communication with a data storage device 106, according to certain embodiments. For instance, the host device 104 may utilize a non-volatile memory (NVM) 110 included in data storage device 106 to store and retrieve data. The host device 104 comprises a host dynamic random access memory (DRAM) 138. The host DRAM 138 includes a host memory buffer (HMB) 150, which is a portion of the host DRAM 138 that is allocated to the data storage device 106 for use by a controller 108. For example, rather than storing data, such as metadata, XOR

parity data, and the like in a volatile memory 112 or in an internal memory of the controller 108, such as DRAM, the controller 108 may instead store the data in the HMB 150. The host device 104 provides full access HMB 150 to the controller 108, such that the HMB 150 is an extension of the memory of the data storage device 106.

- 5 In some examples, the storage system 100 may include a plurality of storage devices, such as the data storage device 106, which may operate as a storage array. For instance, the storage system 100 may include a plurality of data storage devices 106 configured as a redundant array of inexpensive/independent disks (RAID) that collectively function as a mass storage device for the host device 104.
- 10 The host device 104 may store and/or retrieve data to and/or from one or more storage devices, such as the data storage device 106. As illustrated in Figure 1, the host device 104 may communicate with the data storage device 106 via an interface 114. The host device 104 may comprise any of a wide range of devices, including computer servers, network-attached storage (NAS) units, desktop computers, notebook (i.e., laptop) computers, tablet computers, set-top
- 15 boxes, telephone handsets such as so-called “smart” phones, so-called “smart” pads, televisions, cameras, display devices, digital media players, video gaming consoles, video streaming device, or other devices capable of sending or receiving data from a data storage device.

- The data storage device 106 includes the controller 108, NVM 110, a power supply 111, volatile memory 112, the interface 114, and a write buffer 116. In some examples, the data storage
- 20 device 106 may include additional components not shown in Figure 1 for the sake of clarity. For example, the data storage device 106 may include a printed circuit board (PCB) to which components of the data storage device 106 are mechanically attached and which includes electrically conductive traces that electrically interconnect components of the data storage device 106 or the like. In some examples, the physical dimensions and connector configurations of the
- 25 data storage device 106 may conform to one or more standard form factors. Some example standard form factors include, but are not limited to, 3.5” data storage device (e.g., an HDD or SSD), 2.5” data storage device, 1.8” data storage device, peripheral component interconnect (PCI), PCI-extended (PCI-X), PCI Express (PCIe) (e.g., PCIe x1, x4, x8, x16, PCIe Mini Card, MiniPCI, etc.). In some examples, the data storage device 106 may be directly coupled (e.g.,
- 30 directly soldered or plugged into a connector) to a motherboard of the host device 104.

Interface 114 may include one or both of a data bus for exchanging data with the host device 104 and a control bus for exchanging commands with the host device 104. Interface 114 may

operate in accordance with any suitable protocol. For example, the interface 114 may operate in accordance with one or more of the following protocols: advanced technology attachment (ATA) (e.g., serial-ATA (SATA) and parallel-ATA (PATA)), Fibre Channel Protocol (FCP), small computer system interface (SCSI), serially attached SCSI (SAS), PCI, and PCIe, non-volatile memory express (NVMe), OpenCAPI, GenZ, Cache Coherent Interface Accelerator (CCIX), Open Channel SSD (OCSSD), or the like. Interface 114 (e.g., the data bus, the control bus, or both) is electrically connected to the controller 108, providing an electrical connection between the host device 104 and the controller 108, allowing data to be exchanged between the host device 104 and the controller 108. In some examples, the electrical connection of interface 114 may also permit the data storage device 106 to receive power from the host device 104. For example, as illustrated in Figure 1, the power supply 111 may receive power from the host device 104 via interface 114.

The NVM 110 may include a plurality of memory devices or memory units. NVM 110 may be configured to store and/or retrieve data. For instance, a memory unit of NVM 110 may receive data and a message from controller 108 that instructs the memory unit to store the data.

Similarly, the memory unit may receive a message from controller 108 that instructs the memory unit to retrieve data. In some examples, each of the memory units may be referred to as a die. In some examples, the NVM 110 may include a plurality of dies (i.e., a plurality of memory units). In some examples, each memory unit may be configured to store relatively large amounts of data (e.g., 128MB, 256MB, 512MB, 1GB, 2GB, 4GB, 8GB, 16GB, 32GB, 64GB, 128GB, 256GB, 512GB, 1TB, etc.).

In some examples, each memory unit may include any type of non-volatile memory devices, such as flash memory devices, phase-change memory (PCM) devices, resistive random-access memory (ReRAM) devices, magneto-resistive random-access memory (MRAM) devices, ferroelectric random-access memory (F-RAM), holographic memory devices, and any other type of non-volatile memory devices.

The NVM 110 may comprise a plurality of flash memory devices or memory units. NVM Flash memory devices may include NAND or NOR-based flash memory devices and may store data based on a charge contained in a floating gate of a transistor for each flash memory cell. In NVM flash memory devices, the flash memory device may be divided into a plurality of dies, where each die of the plurality of dies includes a plurality of physical or logical blocks, which may be further divided into a plurality of pages. Each block of the plurality of blocks within a particular memory device may include a plurality of NVM cells. Rows of NVM cells may be

electrically connected using a word line to define a page of a plurality of pages. Respective cells in each of the plurality of pages may be electrically connected to respective bit lines.

Furthermore, NVM flash memory devices may be 2D or 3D devices and may be single level cell (SLC), multi-level cell (MLC), triple level cell (TLC), or quad level cell (QLC). The controller

5 108 may write data to and read data from NVM flash memory devices at the page level and erase data from NVM flash memory devices at the block level.

The power supply 111 may provide power to one or more components of the data storage device 106. When operating in a standard mode, the power supply 111 may provide power to one or more components using power provided by an external device, such as the host device 104. For
10 instance, the power supply 111 may provide power to the one or more components using power received from the host device 104 via interface 114. In some examples, the power supply 111 may include one or more power storage components configured to provide power to the one or more components when operating in a shutdown mode, such as where power ceases to be received from the external device. In this way, the power supply 111 may function as an
15 onboard backup power source. Some examples of the one or more power storage components include, but are not limited to, capacitors, super-capacitors, batteries, and the like. In some examples, the amount of power that may be stored by the one or more power storage components may be a function of the cost and/or the size (e.g., area/volume) of the one or more power storage components. In other words, as the amount of power stored by the one or more
20 power storage components increases, the cost and/or the size of the one or more power storage components also increases.

The volatile memory 112 may be used by controller 108 to store information. Volatile memory 112 may include one or more volatile memory devices. In some examples, controller 108 may use volatile memory 112 as a cache. For instance, controller 108 may store cached information
25 in volatile memory 112 until the cached information is written to the NVM 110. As illustrated in Figure 1, volatile memory 112 may consume power received from the power supply 111. Examples of volatile memory 112 include, but are not limited to, random-access memory (RAM), dynamic random access memory (DRAM), static RAM (SRAM), and synchronous dynamic RAM (SDRAM (e.g., DDR1, DDR2, DDR3, DDR3L, LPDDR3, DDR4, LPDDR4, and
30 the like)).

Controller 108 may manage one or more operations of the data storage device 106. For instance, controller 108 may manage the reading of data from and/or the writing of data to the NVM 110. In some embodiments, when the data storage device 106 receives a write command from the

host device 104, the controller 108 may initiate a data storage command to store data to the NVM 110 and monitor the progress of the data storage command. Controller 108 may determine at least one operational characteristic of the storage system 100 and store at least one operational characteristic in the NVM 110.

- 5 In some embodiments, when the data storage device 106 receives a write command from the host device 104 via a host bus 152, the controller 108 temporarily stores the data associated with the write command in the internal memory or write buffer 116 before sending the data to the NVM 110 via an NVM bus 122 for transferring data. The host bus 152 may be referred to as a “link” herein for exemplary purposes. In other embodiments, when the data storage device 106
10 receives a write command from the host device 104, the controller 108 may send the data to the NVM 110 via the NVM bus 122 for transferring data without temporarily storing the data in the buffer 116.

- The controller 108 includes an exclusive OR (XOR) parity engine 118 and a low power mode control unit 120. When data associated with a write command is received by the controller 108,
15 the XOR parity engine 118 generates XOR parity data for the received data. The controller 108 then writes the XOR parity data to the HMB 150 as a backup to other data correction mechanisms, such as error correction code (ECC), low-density parity-check (LDPC), and the like. When data read from the NVM 110 has a read failure due to a bit error rate (BER) or a failed bit count (FBC) exceeding a threshold of the data correction mechanisms, the controller
20 108 may retrieve the XOR parity data from the HMB 150 to attempt to correct the data read from the NVM 110. In some examples, the XOR parity data of the last six written blocks are stored in the HMB 150. The previously listed number of written blocks is not intended to be limiting, and it is contemplated that the controller 108 may store XOR parity data associated with more or less than the listed number of blocks to the HMB 150.

- 25 The low power mode control unit 120 is configured to monitor a health of one or more blocks of the NVM 110 and copy XOR parity data from the HMB 150 to the NVM 110 in critical cases prior to the data storage device 106 transitioning to a low power state mode. In some examples, the low power mode control unit 120 may be configured to transfer XOR parity data from the HMB 150 to the NVM 110 prior to receiving a low power state mode transition request (i.e.,
30 during normal operation). The monitoring of the health of the one or more open blocks includes monitoring a BER of open blocks, which may include a time that the open blocks has been open, a program erase cycle (PEC) of the open blocks, extreme external conditions (e.g., temperature variations) of the open blocks or adjacent open blocks, and an abnormality of the one or more

open blocks relative to a majority of open blocks or relative to adjacent open blocks of the NVM 110. It is to be understood that other factors related to the health of the one or more open blocks may be monitored and may be applicable to the embodiments described herein.

A critical case may be described as any monitored health factor of the one or more open blocks reaching or exceeding a threshold value associated with each of the health factors. The threshold value may be adjusted dynamically by the controller 108 during the operation of the data storage device 106 based on the health and characteristics of the NVM 110. Thus, once a critical case is determined, the controller 108 may retrieve the related XOR parity data from the HMB 150 and store the XOR parity data in the NVM 110. For example, if the low power mode control unit 120 determines that a block (either the block being read or an adjacent block) has an abnormal condition, such as an extreme temperature variation, the low power mode control unit 120 may alert to the controller 108 to retrieve the corresponding XOR parity data. Once the abnormal condition has subsided, the controller may update the XOR parity data stored in the HMB.

Figure 2 is an illustration of a collection of graphs 200 showing a failed bit count (FBC) distribution per cycle (i.e., PEC) and open block retention, according to certain embodiments. It is to be understood that the collection of graphs 200 shown is intended to be an exemplary illustration and is non-limiting. Each title of a graph represents an open block retention time, such that “Time0” refers to an unused memory device or NVM, such as the NVM 110 of Figure 1, and “After_400Hrs” refers to, for example, 400 total hours of use that may be continuous or non-continuous. Furthermore, each curve represents a number of PECs. For example, a PEC of 1 refers to 1 program erase cycle and a PEC of 500 refers to 500 program erase cycles. The X axis, which is FBC/4KB, represents a BER associated with the open blocks. The Y axis is the folder sector failure rate (Folded_SER) which is a probability of having a corresponding FBC.

As shown in the collection of graphs 200, as the number of PECs increases, the FBC/4KB increases. Likewise, as the open block retention time increases, the FBC/4KB value increases. The LDPC line shown at about 700 FBC/4KB represents a LDPC decoding capability or a threshold FBC/4KB value. The LDPC line may be adjusted dynamically based on operational conditions of the data storage device 106. Although, the FBC/4KB for each of the cycles does not exceed or pass the LDPC line, the FBC/4KB for each of the cycles grows close to exceeding or passing the LDPC line. Referring back to Figure 1, the low power mode control unit 120 may determine the probability for a block to pass a BER threshold associated with a decoding capability. Thus, if the low power mode control unit 120 determines that the probability is not likely or that there is a high likelihood of a decoding/read failure occurring, the low power mode

control unit 120 may cause the controller 108 to retrieve the XOR parity data associated with one or more blocks from the HMB 150 prior to entering a low power state mode.

Figure 3 is an illustration of a link power management flow diagram 300, according to certain embodiments. Aspects of the storage system 100 may be referenced in the description herein for exemplary purposes. The data storage device 106 includes several power states. For example, the data storage device 106 may have the following 5 power states: L0, L0s, L1, L2, and L3, where L1 includes a L1.1 sub-state and a L1.2 sub-state. Each of the power states are associated with a distinct operation of the data storage device 106. Power states L0, L0s, and L1 are considered operational power states and utilize about 1 W to about 8 W of power, whereas power states L2 and L3 are considered non-operational power states, utilizing about 2 mW to about 50 mW of power.

An operational power state refers to the ability of the host device 104 to communicate with the NVM 110 of the data storage device 106. A non-operational power state refers to the inability of the host device 104 to communicate with the NVM 110 of the data storage device 106 due to a shut down or disconnection of a link between the host device 104 and the controller 108. The listed non-operational power states are not intended to be limiting and may include other power states, such as the L1.1 and L1.2 power states. Furthermore, it is contemplated that more or less power states than the number of power states shown in the link power management flow diagram 300 may be available and more or less low power states may be applicable to the embodiments described herein.

Power states are numbered sequentially, where higher numbers represent lower power requirements due to a greater number of offline circuits and corresponding higher exit latencies. Furthermore, each power state has an associated power requirement and an exit latency. L0 and L0s may require 4.5 W with the lowest exit latency. L1 may require less power than L0, such as 3 W, and may have an exit latency equal to or higher than the exit latency of L0. L2 may require less power than L1 and may have an exit latency equal to or higher than the exit latency of L1. L3 may require less power than L2 and may have an exit latency equal to or higher than the exit latency of L2. The values for the power states and exit latencies are not intended to be limiting, but to provide an example of possible embodiments.

L0 is referred to as a fully operational state, where I/O commands are enabled, and the device may generate interrupts. Interrupts are an automatic transfer of firmware execution due to a system timer or a user command. Power states L0s and L1 are also operational states; however,

L0s and L1 may have a lower functionality than that of L0. For example, L0s has a similar power requirement as that of the L0, but only allows for a serial link in one direction. However, L1 allows for a bidirectional serial link and allows for a greater reduction in the power requirement, but has a higher exit latency than that of L0 and L0s. Power states L2 and L3 are non-operational power states have a power requirement less than that of the operational power states. Furthermore, memory devices 206 not used are placed in a non-operational power state, L2 and L3, to limit the idle power consumption to a minimal value.

In order for I/O commands to occur, the NVM 110 is woken up and placed into either L0, L0s, or L1. The controller 108 changes the power state of the NVM 110 from the operational power states, such as L0, L0s, or L1, to a different operational power state, such as L0, L0s, or L1, or to a non-operational power state, such as L2 or L3, depending on the situation. However, in order for the NVM 110 to be placed into L2 or L3, the NVM 110 will need to be in power state L2/L3 ready, which is a pseudo-state to prepare the component for a loss in power and reference clock(s). The controller 108 allocates the appropriate amount of power to return all power states L0s, L1, L2, L3 into power state L0 when a full operational state is required. For example, to return to L0 from L2 or L3, the NVM 110 transitions to a transient pseudo-state, LDn, before transitioning to L0. The LDn state may be a fundamental reset state, a hot reset state, or a link disable transmission state by the upstream component (e.g., the host device 104).

The power state L1, in some embodiments, includes additional sub-states, L1.1 and L1.2, where the power state L1 is referred to as L1.0. The L1 sub-states (L1SS), L1.1 and L1.2, may require more power for operation than L2 and L3; however, the L1SS utilizes less power than the L1.0 state. At an L1SS, the NVM 110 remains operational and requires less power to return to a more operational state, such as L1.0 or L0. Furthermore, the L1SS requires less time than the L2 and/or the L3 power states to return to a full active power state L0.

Figure 4 is a flow diagram illustrating a method 400 of correcting a read failure using XOR parity data, according to certain embodiments. Aspects of the storage system 100 of Figure 1 may be referenced in the description herein for exemplary purposes. For example, method 400 may be employed by the controller 108.

At block 402, the controller 108 stores XOR parity data generated by the XOR parity engine 118 in the HMB 150. The XOR parity data stored may be associated with a number of recently written blocks, such as the six most recently written to blocks, or up to as many data blocks in the NVM 110 that include valid data. At block 404, the low power mode control unit 120

monitors a health of one or more blocks and determines if the health is above a threshold. For example, the health may be associated with a BER, a PEC, a temperature variation, an abnormality of an open block when compared to other open blocks, and the like. At block 406, the controller 108 receives an indication to enter a low power state mode. The low power state mode may be any of the low power states, such as a L1SS, L2, or a L3 power state.

At block 408, the controller determines if the health of the one or more recently written to blocks is above a threshold value (e.g., a critical case). If the health is not above the threshold value at block 408, then method 400 continues to method 420, where method 420 may also describe a current method of correcting read failures using XOR parity data. At block 422, the data storage device 106 enters the low power state mode, where the link to the host device 104 is shut down and background operations (e.g., maintenance operations) are triggered. The background operations may include garbage collection, data relocation, wear leveling, and the like.

At block 424, the controller 108 determines that a read failure has occurred due to a BER or FBC exceeding a decoding capability of a decoder. Because the XOR parity data is stored in the HMB 150 and the link to the HMB is down, the controller 108 cannot access the HMB. Thus, the controller 108, at block 426, waits for the link to the host device 104 to be re-initialized.

After the link is re-initialized at block 426, the controller 108 retrieves the XOR parity data from the HMB 150 corresponding to the data associated with the read failure at block 428. At block 430, the controller 108 corrects the data associated with the read failure using the relevant XOR parity data retrieved from the HMB 150.

If the health is above a threshold at block 408, then the low power mode control unit 120 instructs the controller 108 to copy the XOR parity data from the HMB 150 to the NVM 110 to be used as backup for critical cases (e.g., the one or more blocks above the threshold value) at block 410. It is to be understood that the controller 108 may copy only the relevant XOR parity data (i.e., the XOR parity data associated with the one or more open blocks, which may include adjacent open blocks, above the threshold) or the XOR parity data associated with a number of recently written to open blocks. At block 412, the data storage device 106 enters the low power state mode, where the link to the host device 104 is shut down and background operations (e.g., maintenance operations) are triggered. The background operations may include garbage collection, data relocation, wear leveling, and the like.

At block 414, the controller 108 determines that a read failure has occurred, where the read failure may be associated with data read during the back ground operations at block 412. At

block 416, the XOR parity data associated with the data corresponding to the read failure is retrieved from the NVM 110. At block 418, the controller 108 corrects the data associated with the read failure using the relevant XOR parity data retrieved from the NVM 110.

Figure 5 is a flow diagram illustrating a method 500 of correcting a read failure using XOR

5 parity data, according to certain embodiments. Aspects of the storage system 100 of Figure 1 may be referenced in the description herein for exemplary purposes. For example, method 500 may be employed by the controller 108.

At block 502, the controller 108 stores XOR parity data generated by the XOR parity engine 118 in the HMB 150. The XOR parity data stored may be associated with a number of recently
10 written blocks, such as the six most recently written to blocks, or up to as many data blocks in the NVM 110 that include valid data. At block 504, the low power mode control unit 120 monitors a health of one or more blocks and determines if the health is above a threshold. For example, the health may be associated with a BER, a PEC, a temperature variation, an abnormality of an open block when compared to other open blocks, and the like. At block 506,
15 the controller 108 determines if a low power state mode entrance indication (e.g., a command) has been received. If there has not been a low power state mode entrance indication at block 506, then the low power mode control unit 120 continues to monitor the health of the one or more open blocks at block 504. However, if there has been a low power state mode entrance indication at block 506, then the controller determines if a threshold corresponding to the health
20 of the one or more blocks has been reached or exceeded at block 508. For example, the threshold may be the LDPC line exemplified in the collection of graphs 200 of Figure 2.

If the threshold corresponding to the health of the one or more blocks has been reached or exceeded at block 508, then the corresponding XOR parity data is retrieved from the HMB 150 and programmed to the NVM 110 by the controller 108 at block 510. However, if the threshold
25 corresponding to the health of the one or more blocks has not been reached or exceeded at block 508 or after the relevant XOR parity data is retrieved from the HMB 150 and programmed to the NVM 110 at block 510, the data storage device 106 enters the low power state mode at block 512. When the data storage device 106 enters the low power state mode, the link between the HMB 150 and the controller 108 may be deactivated or shut down.

30 At block 514, the controller 108 determines if a read failure that has occurred during a background operation requires XOR parity data (i.e., cannot be corrected using other error correction methods). If the read failure can be corrected by other error correction methods other

than using XOR parity data at block 514, the data associated with the read failure is corrected using other error correction methods at block 516. However, if the read failure requires XOR parity data at block 514, then the controller 108 determines if there is XOR parity data available at block 518. For example, if the relevant XOR parity data was copied from the HMB 150 to the NVM 110 at block 510, then the controller 108 may use the relevant XOR parity data to correct the data associated with the read failure. However, if there has not been any XOR parity data transferred to the NVM 110 prior to entering the low power state mode at block 512 or if the XOR parity data does not correspond with the data associated read failure, then the controller 108 is not able to correct the read failure.

If there is not XOR parity data available at block 518, then the controller 108 waits for the link to be re-established between the host device 104 and the controller 108 at block 522. At block 524, the controller 108 determines that the link has been re-established, such as when the data storage device 106 enters an operational power state. At block 520, the data associated with the read failure is corrected using the relevant XOR parity data. For example, if the relevant XOR parity data was transferred to the NVM 110 at block 510 prior to the data storage device 106 entering the low power state mode at block 512, then the controller 108 may retrieve the relevant XOR parity data from the NVM 110 to correct the data corresponding to the read failure. Otherwise, the relevant XOR parity data is retrieved from the HMB 150 by the controller 108, where the controller 108 uses the retrieved XOR parity data from the HMB 150 to correct the data corresponding to the read failure.

By copying XOR parity data from the HMB to the NVM when the data storage device receives a low power state mode entrance indication, the data storage device may have increased error correction capability during times that the link between the HMB and the controller is unavailable.

In one embodiment, a data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitor health of the memory device, determine that a threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device.

The monitoring the health of the memory device comprises monitoring a bit error rate (BER) of one or more open blocks of the memory device. The monitoring the BER comprises monitoring an amount of time that a block is open. The monitoring the health of the memory device

comprises monitoring a number of program erase cycles (PECs) of one or more open blocks of the memory device. The monitoring the health of the memory device comprises monitoring a temperature variation of one or more open blocks of the memory device. The monitoring the health of the memory device comprises monitoring an abnormality of one or more open blocks relative to a majority of open blocks of the memory device. The controller is further configured to receive an indication from the host device to enter a low power mode. Background operations are executed during the low power mode. The copying the XOR parity data from the HMB to the memory device occurs when the threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded and after receiving the indication from the host device to enter the low power mode. The controller is configured to utilize the XOR parity data copied to the memory device upon determining that a read failure has occurred.

In another embodiment, a data storage device includes a memory device and a controller coupled to the memory device. The controller is configured to store XOR parity data in a host memory buffer (HMB) of a host device, monitoring at least one (and in at one embodiment, at least two) of a bit error rate (BER) of one or more open blocks of the memory device, a number of program erase cycles (PECs) of the one or more open blocks of the memory device, a temperature variation of the one or more open blocks of the memory device, and an abnormality of one or more open blocks relative to a majority of open blocks of the memory device, determine that a threshold corresponding to the monitoring has been reached or exceeded, and copy the XOR parity data from the HMB to the memory device.

The controller is further configured to receive a command to enter a low power mode. The controller is further configured to, upon receiving the command to enter the low power mode, execute the copying. The controller is further configured to determine that a link between a host device and the memory device has been deactivated. The controller is configured to correct data of the one or more open blocks corresponding to the monitoring with the copied XOR parity data. The threshold is dynamic.

In another embodiment, a data storage device includes memory means and a controller coupled to the memory means. The controller is configured to determine that one or more blocks has exceeded a temperature threshold or has an abnormality with one or more blocks as compared to a majority of other blocks of the one or more blocks, retrieve XOR parity data from the memory means, perform a read operation to read first data from the one or more blocks or an adjacent block to the one or more blocks, perform error correction on first data with the XOR parity data retrieved from the memory means, determine that the one or more blocks that either exceeded

the temperature threshold or had an abnormality no longer exceeds the temperature threshold or no longer has an abnormality, retrieve XOR parity data from a host memory buffer (HMB), perform a read operation to read second data from the one or more blocks or an adjacent block to the one or more blocks, and perform error correction on the second data with the XOR parity data retrieved from HMB.

The controller is further configured to update the XOR parity data stored in the HMB upon determining that the one or more blocks that either exceeded the temperature threshold or had an abnormality no longer exceeds the temperature threshold or no longer has an abnormality. The controller is further configured to determine that a link between a host device and the data storage device has been shut down. The controller is configured to, upon determining that the link has been shut down, trigger background operations.

While the foregoing is directed to embodiments of the present disclosure, other and further embodiments of the disclosure may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

WHAT IS CLAIMED IS:

1. A data storage device, comprising:
a memory device; and
5 a controller coupled to the memory device, the controller configured to:
store XOR parity data in a host memory buffer (HMB) of a host device;
monitor health of the memory device;
determine that a threshold corresponding to the health of one or more blocks of
the memory device has been reached or exceeded; and
10 copy the XOR parity data from the HMB to the memory device.
2. The data storage device of claim 1, wherein the monitoring the health of the memory
device comprises monitoring a bit error rate (BER) of one or more open blocks of the memory
device.
- 15 3. The data storage device of claim 2, wherein the monitoring the BER comprises
monitoring an amount of time that a block is open.
4. The data storage device of claim 1, wherein the monitoring the health of the memory
20 device comprises monitoring a number of program erase cycles (PECs) of one or more open
blocks of the memory device.
5. The data storage device of claim 1, wherein the monitoring the health of the memory
device comprises monitoring a temperature variation of one or more open blocks of the memory
25 device.
6. The data storage device of claim 1, wherein the monitoring the health of the memory
device comprises monitoring an abnormality of one or more open blocks relative to a majority of
open blocks of the memory device.
- 30 7. The data storage device of claim 1, wherein the controller is further configured to receive
an indication from the host device to enter a low power mode.

8. The data storage device of claim 7, wherein background operations are executed during the low power mode.

9. The data storage device of claim 7, wherein the copying the XOR parity data from the HMB to the memory device occurs when the threshold corresponding to the health of one or more blocks of the memory device has been reached or exceeded and after receiving the indication from the host device to enter the low power mode.

10. The data storage device of claim 1, wherein the controller is configured to utilize the XOR parity data copied to the memory device upon determining that a read failure has occurred.

11. A data storage device, comprising:

a memory device; and

a controller coupled to the memory device, the controller configured to:

store XOR parity data in a host memory buffer (HMB) of a host device;

monitoring at least one of a bit error rate (BER) of one or more open blocks of the memory device, a number of program erase cycles (PECs) of the one or more open blocks of the memory device, a temperature variation of the one or more open blocks of the memory device, and an abnormality of one or more open blocks relative to a majority of open blocks of the memory device;

determine that a threshold corresponding to the monitoring has been reached or exceeded; and

copy the XOR parity data from the HMB to the memory device.

12. The data storage device of claim 11, wherein the controller is further configured to receive a command to enter a low power mode.

13. The data storage device of claim 12, wherein the controller is further configured to, upon receiving the command to enter the low power mode, execute the copying.

14. The data storage device of claim 11, wherein the controller is further configured to determine that a link between a host device and the memory device has been deactivated.

15. The data storage device of claim 14, wherein the controller is configured to correct data of the one or more open blocks corresponding to the monitoring with the copied XOR parity data.

5 16. The data storage device of claim 11, wherein the threshold is dynamic.

17. A data storage device, comprising:
memory means; and

a controller coupled to the memory means, the controller configured to:

10 determine that one or more blocks has exceeded a temperature threshold or has an abnormality with one or more blocks as compared to a majority of other blocks of the one or more blocks;

retrieve XOR parity data from the memory means;

15 perform a read operation to read first data from the one or more blocks or an adjacent block to the one or more blocks;

perform error correction on first data with the XOR parity data retrieved from the memory means;

20 determine that the one or more blocks that either exceeded the temperature threshold or had an abnormality no longer exceeds the temperature threshold or no longer has an abnormality;

retrieve XOR parity data from a host memory buffer (HMB);

perform a read operation to read second data from the one or more blocks or an adjacent block to the one or more blocks; and

25 perform error correction on the second data with the XOR parity data retrieved from HMB.

18. The data storage device of claim 17, wherein the controller is further configured to update the XOR parity data stored in the HMB upon determining that the one or more blocks that either exceeded the temperature threshold or had an abnormality no longer exceeds the
30 temperature threshold or no longer has an abnormality.

19. The data storage device of claim 17, wherein the controller is further configured to determine that a link between a host device and the data storage device has been shut down.

20. The data storage device of claim 19, wherein the controller is configured to, upon determining that the link has been shut down, trigger background operations.




FIG. 1




FIG. 2

300





FIG. 3






FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2022/028226

A. CLASSIFICATION OF SUBJECT MATTER G06F 3/06(2006.01)i 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 3/06(2006.01); G06F 11/10(2006.01); G11C 16/34(2006.01); G11C 29/42(2006.01); G11C 29/52(2006.01) Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Korean utility models and applications for utility models Japanese utility models and applications for utility models Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) eKOMPASS(KIPO internal) & Keywords: memory device, host memory buffer (HMB), XOR parity data, monitor, health, threshold, low power mode		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2020-0151055 A1 (SAMSUNG ELECTRONICS CO., LTD.) 14 May 2020 (2020-05-14) paragraphs [0035], [0059], [0070], [0098], [0107], [0110], [0135]; and figures 1, 4-5, 9, 11, 15	1-5,10-11,14-16
A		6-9,12-13,17-20
Y	US 2019-0265888 A1 (WESTERN DIGITAL TECHNOLOGIES, INC.) 29 August 2019 (2019-08-29) paragraphs [0049], [0059], [0068]; claims 1, 4, 11; and figures 1A, 2A	1-5,10-11,14-16
A	US 2018-0018101 A1 (SANDISK TECHNOLOGIES LLC) 18 January 2018 (2018-01-18) paragraphs [0061]-[0067]; and figures 5A-5B	1-20
A	US 2020-0218605 A1 (WESTERN DIGITAL TECHNOLOGIES, INC.) 09 July 2020 (2020-07-09) paragraphs [0023]-[0028]; and figure 1	1-20
A	KR 10-2020-0085519 A (SK HYNIX INC.) 15 July 2020 (2020-07-15) claims 1-6	1-20
☐ 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 08 September 2022		Date of mailing of the international search report 15 September 2022
Name and mailing address of the ISA/KR Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon 35208, Republic of Korea Facsimile No. +82-42-481-8578		Authorized officer YANG, JEONG ROK Telephone No. +82-42-481-5709

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/US2022/028226

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
US	2020-0151055	A1	14 May 2020	CN	111177040	A	19 May 2020
				DE	102019111132	A1	14 May 2020
				KR	10-2020-0054402	A	20 May 2020
				US	11016846	B2	25 May 2021
				US	2021-0248035	A1	12 August 2021
US	2019-0265888	A1	29 August 2019	US	10698610	B2	30 June 2020
				WO	2019-164573	A1	29 August 2019
US	2018-0018101	A1	18 January 2018	US	10521118	B2	31 December 2019
US	2020-0218605	A1	09 July 2020	US	10872012	B2	22 December 2020
KR	10-2020-0085519	A	15 July 2020	CN	111415699	A	14 July 2020
				US	10963339	B2	30 March 2021
				US	2020-0218606	A1	09 July 2020