



(51) International Patent Classification:

G06F 12/02 (2006.01) G06F 3/06 (2006.01)

(21) International Application Number:

PCT/US2019/067471

(22) International Filing Date:

19 December 2019 (19.12.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/425,455 29 May 2019 (29.05.2019) US

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

(72) Inventors: **MUTHIAH, Ramanathan**; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **RAMAMURTHY, Rankumar**; 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: HOST AND METHOD FOR STORAGE SYSTEM CALIBRATION




FIG. 1A

(57) Abstract: A storage system, host, and method for storage system calibration are provided. In one embodiment, a storage system is provided comprising a memory and a controller. The controller is configured to: determine a pattern of host writes to the memory; determine whether the pattern of host writes matches a granularity of a logical-to-physical address map used by the storage system; and in response to determining that the pattern of host writes does not match the granularity of the logical-to-physical address map used by the storage system, change the granularity of the logical-to-physical address map used by the storage system. In another embodiment, the storage system calibration is done by host directive. Other embodiments are provided.



Host and Method for Storage System Calibration

Background

[0001] Storage systems can come in a variety of forms and have a variety of performance characteristics. For example, Secure Digital (SD) cards with the same form factor can have various performance metrics for different product lines, such as Surveillance, Mobile Hosts, Gaming Hosts, and Video Speed Class. Each of these product lines has a different configuration in terms of memory and firmware binaries according to the specification for its produce line. Each of these product lines support storage of all types of data, but a given product line may be optimized for some data pattern, usually at the expense of other data patterns, due to the storage system's internal flash translation layer (FTL) design (i.e., its logical-to-physical address (L2P) map and allotted cache memory). For example, SD cards marketed as Video Speed Class (VSC) may be suboptimal for use with short, unaligned types of data. Similarly, SD cards marketed with increased random performance may be suboptimal for use with sequential data. Accordingly, a given storage system may not be tuned to perform optimally for all data patterns.

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 a flowchart 300 of a method an embodiment for runtime learning and FTL formatting.

[0008] Figure 4 is a flowchart of a runtime coarser flash translation later (FTL) process of an embodiment.

[0009] Figure 5 is a flowchart of a runtime finer FTL process of an embodiment.

[0010] Figure 6 is a flowchart of a method of an embodiment performed by a host for storage system calibration.

Detailed Description

[0011] Overview

[0012] By way of introduction, the below embodiments relate to a storage system and method for storage system calibration. In one embodiment, a storage system is provided comprising a memory and a controller. The controller is configured to: determine a pattern of host writes to the memory; determine whether the pattern of host writes matches a granularity of a logical-to-physical address map used by the storage system; and in response to determining that the pattern of host writes does not match the granularity of the logical-to-physical address map used by the storage system, change the granularity of the logical-to-physical address map used by the storage system.

[0013] In some embodiments, changing the granularity of the logical-to-physical address map comprises making the logical-to-physical address map coarser.

[0014] In some embodiments, making the logical-to-physical address map coarser comprises: performing host data consolidation; managing new write entries as a delta file until a coarser logical-to-physical address map is complete; and flushing the coarser logical-to-physical address map to the memory.

[0015] In some embodiments, changing the granularity of the logical-to-physical address map comprises making the logical-to-physical address map finer.

[0016] In some embodiments, making the logical-to-physical address map finer comprises: splitting the logical-to-physical address map into finer entries; managing new write entries as a delta file until the splitting is complete; and flushing the split logical-to-physical address map to the memory.

[0017] In some embodiments, the controller is further configured to determine whether the pattern of host writes matches the current granularity of the logical-to-physical address map used by the storage system in response to a number of host writes in the

pattern exceeding a threshold and/or in response to a positive health determination of the memory.

[0018] In some embodiments, the pattern of host writes to the memory is determined in runtime across a plurality of write cycles.

[0019] In some embodiments, the memory comprises a three-dimensional memory.

[0020] In some embodiments, the storage system is embedded in a host.

[0021] In some embodiments, the storage system is removably connected to a host.

[0022] In another embodiment, a method is provided that is performed in a storage system comprising a memory, wherein the storage system is in communication with a host. The method comprises: tracking host writes to the memory; determining whether a media management layer used by the storage system results in more sequential-type results or more random-type results than the host writes; in response to determining that the media management layer used by the storage system results in more sequential-type results, making translation mapping performed by the media management layer coarser; and in response to determining that the media management layer used by the storage system results in more random-type results, making translation mapping performed by the media management layer finer.

[0023] In some embodiments, making translation mapping performed by the media management layer coarser comprises: consolidating host data; managing new write entries as a delta file until a coarser media management layer is complete; and flushing the coarser media management layer to the memory.

[0024] In some embodiments, making translation mapping performed by the media management layer finer comprises: splitting the media management layer into finer entries; managing new write entries as a delta file until the splitting is complete; and flushing the split media management layer to the memory.

[0025] In some embodiments, the method further comprises determining that a number of tracked host writes exceeds a threshold and/or making a positive health determination of the memory.

[0026] In some embodiments, the tracking is determined in runtime across a plurality of write cycles.

[0027] In another embodiment, a storage system is provided comprising: a memory; means for determining a pattern of host writes to the memory; means for determining whether the pattern of host writes matches a granularity of a logical-to-physical address map used by the storage system; and means for in response to determining that the pattern of host writes does not match the granularity of the logical-to-physical address map used by the storage system, changing the granularity of the logical-to-physical address map used by the storage system.

[0028] In some embodiments, means for changing the granularity of the logical-to-physical address map comprises means for making the logical-to-physical address map coarser.

[0029] In some embodiments, means for changing the granularity of the logical-to-physical address map comprises means for making the logical-to-physical address map finer.

[0030] In some embodiments, the storage system further comprises means for determining whether the pattern of host writes matches the current granularity of the logical-to-physical address map used by the storage system in response to a number of host writes in the pattern exceeding a threshold and/or in response to a positive health determination of the memory.

[0031] In some embodiments, the pattern of host writes to the memory is determined in runtime across a plurality of write cycles.

[0032] In another embodiment, a host is provided comprising a controller configured to: determine a pattern of host writes to a memory of a storage system in communication with the host; determine whether the pattern of host writes matches a granularity of a logical-to-physical address map used by the storage system; and in response to determining that the pattern of host writes does not match the granularity of the logical-to-physical address map used by the storage system, directing the storage system to change the granularity of the logical-to-physical address map used by the storage system.

[0033] In some embodiments, changing the granularity of the logical-to-physical address map comprises making the logical-to-physical address map coarser.

[0034] In some embodiments, making the logical-to-physical address map coarser comprises: performing host data consolidation; managing new write entries as a delta file

until a coarser logical-to-physical address map is complete; and flushing the coarser logical-to-physical address map to the memory.

[0035] In some embodiments, changing the granularity of the logical-to-physical address map comprises making the logical-to-physical address map finer.

[0036] In some embodiments, making the logical-to-physical address map finer comprises: splitting the logical-to-physical address map into finer entries; managing new write entries as a delta file until the splitting is complete; and flushing the split logical-to-physical address map to the memory.

[0037] In some embodiments, the controller is further configured to receive an indication from the storage system that the storage system rejected the direction from the host to change the granularity of the logical-to-physical address map used by the storage system.

[0038] In some embodiments, the controller is further configured to determine whether the pattern of host writes matches the current granularity of the logical-to-physical address map used by the storage system in response to a number of host writes in the pattern exceeding a threshold and/or in response to a positive health determination of the memory.

[0039] In some embodiments, the pattern of host writes to the memory is determined in runtime across a plurality of write cycles.

[0040] In some embodiments, the memory comprises a three-dimensional memory.

[0041] In some embodiments, the storage system is embedded in the host.

[0042] In some embodiments, the storage system is removably connected to the host.

[0043] In another embodiment, a method for storage system calibration is provided that is performed in a host in communication with a storage system comprising a memory.

The method comprises: tracking host writes to the memory; determining whether a media management layer used by the storage system results in more sequential-type results or more random-type results than the host writes; in response to determining that the media management layer used by the storage system results in more sequential-type results, instructing the storage system to make translation mapping performed by the media management layer coarser; and in response to determining that the media management layer used by the storage system results in more random-type results, instructing the

storage system to make translation mapping performed by the media management layer finer.

[0044] In some embodiments, making translation mapping performed by the media management layer coarser comprises: consolidating host data; managing new write entries as a delta file until a coarser media management layer is complete; and flushing the coarser media management layer to the memory.

[0045] In some embodiments, the method further comprises making translation mapping performed by the media management layer finer comprises: splitting the media management layer into finer entries; managing new write entries as a delta file until the splitting is complete; and flushing the split media management layer to the memory.

[0046] In some embodiments, the method further comprises determining that a number of tracked host writes exceeds a threshold and/or making a positive health determination of the memory.

[0047] In some embodiments, the method further comprises receiving an indication from the storage system that the storage system rejected the instruction from the host to make the media management layer finer.

[0048] In some embodiments, the tracking is determined in runtime across a plurality of write cycles.

[0049] In another embodiment, a host is provided comprising: means for determining a pattern of host writes to a memory of a storage system in communication with the host; means for determining whether the pattern of host writes matches a granularity of a logical-to-physical address map used by the storage system; and means for in response to determining that the pattern of host writes does not match the granularity of the logical-to-physical address map used by the storage system, directing the storage system to change the granularity of the logical-to-physical address map used by the storage system.

[0050] In some embodiments, means for changing the granularity of the logical-to-physical address map comprises means for making the logical-to-physical address map coarser.

[0051] In some embodiments, means for changing the granularity of the logical-to-physical address map comprises means for making the logical-to-physical address map finer.

[0052] In some embodiments, the host further comprises means for receiving an indication from the storage system that the storage system rejected the direction from the host to change the granularity of the logical-to-physical address map used by the storage system.

[0053] In some embodiments, the host further comprises means for determining whether the pattern of host writes matches the current granularity of the logical-to-physical address map used by the storage system in response to a number of host writes in the pattern exceeding a threshold and/or in response to a positive health determination of the memory.

[0054] In some embodiments, the pattern of host writes to the memory is determined in runtime across a plurality of write cycles.

[0055] 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.

[0056] Embodiments

[0057] Storage systems suitable for use in implementing aspects of these embodiments are shown in Figures 1A-1C. Figure 1A is a block diagram illustrating a non-volatile storage system 100 according to an embodiment of the subject matter described herein. Referring to Figure 1A, non-volatile storage system 100 includes a controller 102 and non-volatile memory that may be made up of one or more non-volatile memory die 104. As used herein, the term die refers to the collection of non-volatile memory cells, and associated circuitry for managing the physical operation of those non-volatile memory cells, that are formed on a single semiconductor substrate. Controller 102 interfaces with a host system and transmits command sequences for read, program, and erase operations to non-volatile memory die 104.

[0058] The controller 102 (which may be a non-volatile memory controller (e.g., a flash, resistive random-access memory (ReRAM), phase-change memory (PCM), or magnetoresistive random-access memory (MRAM) 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.

[0059] As used herein, a non-volatile memory controller is a device that manages data stored on non-volatile memory and communicates with a host, such as a computer or electronic device. A non-volatile memory controller can have various functionality in addition to the specific functionality described herein. For example, the non-volatile memory controller can format the non-volatile memory to ensure the memory is operating properly, map out bad non-volatile 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 non-volatile memory controller and implement other features. In operation, when a host needs to read data from or write data to the non-volatile memory, it can communicate with the non-volatile memory controller. If the host provides a logical address to which data is to be read/written, the non-volatile memory controller can convert the logical address received from the host to a physical address in the non-volatile memory. (Alternatively, the host can provide the physical address.) The non-volatile 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).

[0060] Non-volatile memory die 104 may include any suitable non-volatile storage medium, including resistive random-access memory (ReRAM), magnetoresistive random-access memory (MRAM), phase-change memory (PCM), 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.

[0061] 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.

[0062] 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 storage system architectures (such as the ones shown in Figures 1B and 1C), 2, 4, 8 or more memory channels may exist between the controller and the 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.

[0063] 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), peripheral component interconnect express (PCIe) interface, or double-data-rate (DDR) interface. Storage module 200, in one embodiment, may be a solid state drive (SSD), or non-volatile dual in-line memory module (NVDIMM), such as found in server PC or portable computing devices, such as laptop computers, and tablet computers.

[0064] 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.

[0065] 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. 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 of related functions, or a self-contained hardware or software component that interfaces with a larger system, for example. Modules of the controller 102 may include a flash translation layer (FTL) formatter 111, which is discussed in more detail below, and can be implemented in hardware or software/firmware. The FTL formatter 111 can be configured to perform the algorithms and methods discussed below and shown in the attached drawings.

[0066] 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.

[0067] 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, serially attached small computer system interface (SAS), Fibre

Channel, universal serial bus (USB), PCIe, and NVMe. The host interface 120 typically facilitates transfer for data, control signals, and timing signals.

[0068] 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.

[0069] 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.

[0070] 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 ReRAM, MRAM, PCM, 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. Peripheral circuitry 141 includes a state machine 152 that provides status information to the controller 102.

[0071] Returning again to Figure 2A, the flash control layer 132 (which will be referred to herein as the flash translation layer (FTL) or, more generally, the “media management layer,” as the memory may not be flash) handles flash errors and interfaces with the host. In particular, the FTL, which may be an algorithm in firmware, is responsible for the internals of memory management and translates writes from the host into writes to the memory 104. The FTL may be needed because the memory 104 may have limited endurance, may only be written in multiples of pages, and/or may not be written unless it is erased as a block. The FTL understands these potential limitations of the memory 104, which may not be visible to the host. Accordingly, the FTL attempts to translate the writes from host into writes into the memory 104.

[0072] The FTL may include a logical-to-physical address (L2P) map and allotted cache memory. In this way, the FTL translates logical block addresses (“LBAs”) from the host to physical addresses in the memory 104. The FTL can include other features, such as, but not limited to, power-off recovery (so that the data structures of the FTL can be recovered in the event of a sudden power loss) and wear leveling (so that the wear across memory blocks is even to prevent certain blocks from excessive wear, which would result in a greater chance of failure).

[0073] The granularity of the logical-to-physical address (L2P) map is typically either pre-set in the factory or chosen prior to run-time of the storage system. The granularity at which the logical-to-physical address (L2P) map is configured may depend on the intended use of the storage system. For example, if the storage system is intended to be used with sequential data (data that is contiguous in the logical address space/ not fragmented), a large logical group size (e.g., 4 megabytes (MB)) can be used. Using a large logical group size enables addressing a large amount of data with a single index or entry in both RAM and memory for reducing processing overhead. That is, indexing a large logical group size can maintain the efficiencies gained by fewer entries in the logical-to-physical address table. In contrast, if the storage system is intended to be used with fragmented data, a smaller logical group size (e.g., 4 kilobytes (KB)) can be used. That is, for data that is fragmented (not sequential), a large logical group block is not ideal.

[0074] As mentioned above, storage systems typically support storage of all types of data, but a storage system may be optimized for some data pattern, usually at the expense of other data patterns. For example, some storage systems may be suboptimal for use with short, unaligned types of data, while other storage systems may be suboptimal for use with sequential data. This is due to the granularity of the storage system's logical-to-physical address map in the FTL. Accordingly, a typical storage systems is optimized for some data patterns, and its behavior with other data patterns is sub-optimal. That is, the in-built FTL system in some storage systems is usually optimal for certain data patterns and not optimal for most others. The retail device serving a product purpose may not be able to satisfy another host (of a completely different type) of its requirement since it is not dynamically configurable.

[0075] To address this problem, the following embodiments provide a storage system 100 (sometimes referred to herein as a "device") that can, based on data analysis in runtime, support logical-to-physical (L2P) association at one of a plurality of different logical group granularities (e.g., one sector (512 bytes), 4 KB, 8 KB, 16 KB, or 32 KB). The lower granularities help in increased random performance, and the higher granularities offer higher sequential performance for the same FTL cache. Because the storage system 100 can adapt itself to host use cases based on the data pattern on-the-fly during runtime (and change them throughout the course of operation), the same storage system 100 can be sold in multiple product classes.

[0076] Further, the decision to recalibrate the FTL from its default to a newer FTL may or may not involve a full-fledged FTL format in the storage system 100, which directly affects the time taken for such an adjustment. This enables high endurance retail removable storage systems to be usable by multiple types of hosts at different points. The storage system 100 may not want to calibrate the FTL often owing to latency involved. But, having a dynamic way of changing the FTL granularity enables less FTL calibration overhead, optimum FTL cache usage, and optimal performance across several data storage types during the life of the storage system 100. This approach optimizes writes and reads, as well as garbage collection.

[0077] So, in these embodiments, the dynamic calibration of the storage system 100 involves the storage system 100 recalibrating the FTL based on learning in accordance

with its dynamic usage requirement through the FTL formal (in one embodiment, only the FTL format and not the data format, as existing data is safe and intact). The learning could be based on a history of a threshold number of commands, patterns, or full-fledged data analysis across a configurable threshold number of storage system 100 cycles.

[0078] The following paragraphs and Figures 3-6 illustrate an example of an embodiment. It should be noted that this is merely an example and other implementation can be used. Turning now to Figure 3, Figure 3 is a flowchart 300 of a method of an embodiment for runtime learning and FTL formatting. As shown in Figure 3, the storage system 100 receives a host write or input during runtime of the storage system (act 310) and (a) builds a command histogram and (b) performs continuous learning in “runtime” across multiple device cycles/threshold (act 320). In this way, the storage system 100 continuously monitors the data being accessed by the host. During this process, the storage system 100 determines if there have been enough commands received to decide on a FTL pattern and/or if the memory 104 (here, flash) health is good enough to consider FTL formatting (act 330). The number of commands can be a number chosen to represent a robust data set. The health of the memory 104 can be determined in any suitable way. For example, in one embodiment, the health of the memory 104 is determined by a number of read/write/erase cycles the memory 104 has endured. For instance, if the memory 104 has an expected lifetime of 1,000 cycles, and the memory 104 has already gone through 900 cycles, the memory 104 may be considered “end of life,” where it would not be worth the time and effort to reformat the FTL.

[0079] If the answer to those inquiries is no, the storage system 100 continues with its “routine” work using the existing FTL (act 340). However, if the answer to those inquiries is yes, the storage system 100 then determines if the threshold history pattern matches the FTL granularity (act 350). For example, the storage system 100 can look for some pattern, such as Video Speed Class patterns, in the storage system 100 over a threshold number of device cycles. If the storage system 100 determines the threshold history pattern matches the FTL granularity, the storage system 100 determines that the pattern is already in sync with the storage system’s FTL design (act 360) and proceeds with its “routine” work with the existing FTL (act 365).

[0080] However, if the storage system 100 determines the threshold history pattern does not match the FTL granularity, the storage system 100 determines if the host pattern is more sequential than the current FTL design (act 370). If it is, the storage system 100 triggers a “runtime coarser FTL process,” which is shown in the flowchart 400 in Figure 4. If it is not, the storage system 100 triggers a “runtime finer FTL process,” which is shown in the flowchart 500 in Figure 5. That is, if it turns out that the default FTL is more towards sequential data, but the write pattern from the host has been random data for a long time (e.g., many flash cycles), such as with gaming hosts, then a finer FTL is preferred.

[0081] Turning now to Figure 4, Figure 4 is a flowchart 400 of a “runtime coarser FTL process” of an embodiment. As shown in Figure 4, in this process, the storage system 100 performs host data consolidation (i.e., host data movement) in sync with new, coarse FTL tables (act 410). For example, garbage collection can be used to compact data. Next, the storage system 100 manages new write entries as an “FTL delta” until completion of the FTL format (act 420). An FTL delta is used to manage intermediate data flow during runtime. Finally, the storage system 100 flushes the new FTL tables in the memory (flash) and posts a “done” status (act 430). As shown by this example, a decision to go coarser compared to a default FTL association (say from a default sector level to a new 4K level, or even more further, to a 16K or 32K level) would result in less FTL entries, hence a smaller FTL block budget. However, going coarser involves host data movement, as host data desegregation is required to adjust for new FTL links. This mode is comparatively more suitable for sequential data (compared to sector level association). Since the storage system 100 is aware of the pattern, it can afford to release some of the FTL blocks to improve device performance.

[0082] Turning now to Figure 5, Figure 5 is a flowchart 500 of a “runtime finer FTL process” of an embodiment. As shown in Figure 5, the storage system 100 performs an FTL table split into finer entries as decided by runtime learning (act 510). Next, the storage system 100 manages new entries as an FTL delta until the completion of the FTL format (act 520). Finally, the storage system 100 flushes the new FTL tables and posts a “done” status (act 530). As shown in this example, a decision to calibrate to a finer FTL (e.g., new associations at the sector level from a default 4K level; or say, to finer 4K from

32K) would involve additions of new FTL updates alone, but no host data movement in the memory 104, since the data is desegregated owing to a coarser default FTL association. In the case of new, finer FTL, the existing FTL logical-to-physical association is split, which increases the number of entries. Increased FTL pages leads to an increased block budget for the FTL. The finer possible FTL calibration defines the worst case FTL block budget in the storage system 100, which can be used as overprovisioning blocks when the FTL is calibrated to be coarser. Further, because going to finer FTL does not require host data movement, the storage system 100 can perform such an operation with a spread ratio that is phased within protocol time-out limits. It should be noted that the “protocol time-out limits” mentioned above are applicable both for going to coarser or finer FTL. Some work may be performed in both cases, such as in a phased manner.

[0083] Consider the following examples. Assuming the FTL page size to be fixed at 4KB, consider the logical block granularity at 4 KB, one FTL page has 1,024 L2P entries, each entry being 4 bytes. At a 4KB data flash management unit (FMU) granularity, 1,024 entries will cover 4 MB of logical region. So, in a 4 KB RAM, one FTL pages covers a total of 4 MB of logical data. Accordingly, finer granularity occupies more FTL space and giving better random performance. In another example, the change in granularity is from a 4 KB logical block to a 32 KB logical block, where the FTL page size is still 4 KB. Again, one FTL page is 1,024 L2P entries, each entry being 4 bytes. So, the FTL page size is 4 KB. At a 32 KB data FMU granularity, 1,024 entries will cover 32 MB of logical region. So, in a 4 KB RAM, one FTL page covers a total of 32 MB of logical data. Accordingly, coarser granularity occupies lesser FTL space, giving better sequential performance.

[0084] It should be noted that different jumps in calibration levels from the current default may result in different latencies and depend on the amount of data present in the storage system 100. As such, the storage system 100 may not want to calibrate FTL often, owing to the latency involved. But, it opens up a way where retail storage systems, such as SD cards, can be used across multiple hosts having different types of requirements, such as Surveillance, Mobile Hosts, gaming, and low-end camera hosts operating under the Video Speed Class specification. These embodiments enable less

FTL calibration overhead and optimal performance across several data storage types. Optimum usage of the FTL RAM cache minimizes hardware cost.

[0085] In the above examples, the storage system 100 made the determination as to whether to change the FTL. In another embodiment, the storage system 100 performs FTL formatting based on a host directive, where the decision to recalibrate is taken as a result of learning/requirements at the host side. For example, the host (e.g., a controller in the host) might decide to perform FTL format/recalibration when the data pattern is not optimally in sync with the storage environment and/or if the value proposition is in favor. The host can be provided with a table of possible calibrations which the underlying firmware or hardware can support in the storage system 100, and the host can decide on the right logical-to-physical address association to use based on current as well as future needs after assessing the value proposition for a given FTL format latency. This embodiment may involve a handshake and protocol standardization between the host and the storage system 100, and the storage system 100 may have the option to reject the host directive.

[0086] Turning again to the drawings, Figure 6 is a flowchart 600 of this embodiment. As shown in Figure 6, the host (a) performs learning and (b) directs the storage system 100 to perform runtime FTL formatting (act 610). The storage system 100 evaluates the memory (here, NAND flash) health and (b) decides to “accept” or “reject” the host directive based on the health evaluation (act 620). If the storage system 100 decides not to accept the host directive (act 630), the host knows that the memory is in bad health (or due to other reasons), and no new FTL format will be used (act 640). If the storage system 100 decides to accept the host directive (act 630), the storage system 100 then determines if a coarser or finer FTL is needed (act 650). If a coarser FTL is needed, the storage system 100 triggers the “runtime coarser FTL process” shown in the flowchart 400 in Figure 4 (act 660). If a finer FTL is needed, the storage system 100 triggers the “runtime finer FTL process” shown in the flowchart 500 in Figure 5 (act 670).

[0087] There are several advantages associated with these embodiments. For example, one advantage is that the FTL granularity can be dynamically changed (e.g., a 4 byte physical address may point to 4 KB logical data in one FTL configuration, and the same 4 byte may point to 32 KB logical data in another FTL configuration). So, RAM can be

efficiently used by the storage system 100, among other advantages. As another advantage, the storage system 100 can dynamically define the retail storage system market based on the underlying data pattern to extract the best performance as well as optimize FTL cache and block usages. Further, the storage system 100 can adapt itself to the host use case based on the data pattern, so that the same storage system 100 can be sold to multiple product classes. Yet another advantage is that the storage system 100 start-up and garbage collection time can be lowered considerably with these embodiments. Additionally, as a byproduct of learning, since the storage system 100 has prior information on the nature of data that is being stored, it can adjust its resources according to optimize various performances. For example, by changing its FTL caching policy according to a chosen FTL granularity, the storage system 100 can enable cache and block resources to be used elsewhere in the system leading to enhanced performance. As another example, the extra blocks that are gained from coarser FTL calibration can be used for capacity overprovisioning or for burst performance if it turns out that the analysis is mostly sequential biased in the storage system.

[0088] It should be noted that these embodiments can be used in conjunction with storage systems that uses a layered architecture approach, such as the one described in U.S. Patent Application Publication No. US 2015/0143029A1, which is hereby incorporated by reference. In the case of layered architecture, both the FTL directory and FTL sub-page granularity can be formatted in run-time to a new value leading to a one-time FTL table re-arrangement for both directory and sub-pages through garbage collection. In the case of a non-layered FTL architecture, the single level of granularity can be formatted to a newer value at runtime. The FTL can be modified in real time (e.g., in the 10th cycle of storage system usage in a 500 cycle endurance product). FTL formatting can lead to table changes and may or may not lead to host data movement. During real-time FTL formatting, the host data can be saved in the storage system's memory as delta mapping with corresponding run-lengths to avoid confusion as well as to avoid data loss. The decision can be made during the life of the storage system after sufficient learning to avoid FTL choking (i.e., unnecessary frequent overheads with FTL formatting). Since the decision is run-time, any format may lead to extra garbage

collection work. As a policy, when the FTL format work is incomplete, a new decision to reformat can be put on hold to avoid device choking.

[0089] To describe an example, consider U.S. Patent Application Publication No. US 2015/0143029A1, which describes a layered architecture, a global address table (GAT) directory covering 4 MB, and GAT pages (split for random data within that directory) covering 4 KB. For the same example case, based on learning from a pattern or a host directive, the storage system can change the GAT directory in runtime to cover 32 MB of logical region and change the corresponding GAT page to cover 32 KB of data somewhere in the middle of device's life (e.g., the 10th cycle of product, learning from first 9 cycles).

[0090] In this case, the new chosen granularity has become coarser. Along with FTL changes, the host data would have to be relocated to complete the FTL format procedure. Until that point of FTL format completion, the FTL can be stored in memory as an FTL delta without any dependency on the tables to avoid confusion. Once the format is complete, the FTL delta entries can be merged back to the newly-created FTL tables.

[0091] In the case of choosing finer granularity (e.g., through learning), the host data need not be moved, but FTL new tables are created in an extended FTL block budget (a finer FTL table would need a greater block budget). Until the creation of the new FTL tables in extended blocks, any new logical-to-physical updates can be done as an FTL delta and stored in memory without being merged to the main FTL tables.

[0092] As and when the memory's health deteriorates, some storage systems may not want to perform FTL formatting owing to larger background time involved and lower return on investment. So, effectively the storage system may reject a host directive keeping in mind the garbage collection workload, ongoing FTL format, or memory health to avoid choking.

[0093] In summary, the embodiments presented here can change the storage system's FTL according to a real-time pattern to enhance performance with random or sequential writes. This provides advantages for a given memory's use in multiple products with the same form factor. FTL formatting may mean only FTL changes in some cases and may include host data movement and sometimes idle-time control movement. Learning the real-time pattern can be done over a certain number of cycles. These embodiments can

decide when to calibrate and when not to calibrate to avoid choking caused by unnecessary calibration. These embodiments can use protocol standardization for host-side learning. Different FTL caching policies can be used to optimize caches, and different latencies can be used for different amounts of data, as latency and the value proposition can depend on how much data is written in a cycle. Resource adjustments in the memory can be enabled as a byproduct of FTL learning, which can lead to power savings.

[0094] 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.

[0095] 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.

[0096] 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.

[0097] 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.

[0098] 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.

[0099] 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.

[00100] 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).

[00101] 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.

[00102] 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.

[00103] 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.

[00104] 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.

[00105] 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.

[00106] 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.

[00107] 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 host comprising:
a controller configured to:
determine a pattern of host writes to a memory of a storage system
configured to communicate with the host;
determine whether the pattern of host writes matches a granularity of a
logical-to-physical address map used by the storage system; and
in response to determining that the pattern of host writes does not match
the granularity of the logical-to-physical address map used by the storage system,
direct the storage system to change the granularity of the logical-to-physical
address map used by the storage system.
2. The host of Claim 1, wherein changing the granularity of the logical-to-physical
address map comprises making the logical-to-physical address map coarser.
3. The host of Claim 2, wherein making the logical-to-physical address map coarser
comprises:
consolidating host data;
managing new write entries as a delta file until a coarser logical-to-physical
address map is complete; and
flushing the coarser logical-to-physical address map to the memory.

4. The host of Claim 1, wherein changing the granularity of the logical-to-physical address map comprises making the logical-to-physical address map finer.
5. The host of Claim 4, wherein making the logical-to-physical address map finer comprises:
 - splitting the logical-to-physical address map into finer entries;
 - managing new write entries as a delta file until the splitting is complete; and
 - flushing the split logical-to-physical address map to the memory.
6. The host of Claim 1, wherein the controller is further configured to determine whether the pattern of host writes matches current granularity of the logical-to-physical address map used by the storage system in response to a number of host writes in the pattern exceeding a threshold and/or in response to a positive health determination of the memory.
7. The host of Claim 1, wherein the pattern of host writes to the memory is determined in runtime across a plurality of write cycles.
8. The host of Claim 1, wherein the controller is further configured to receive an indication from the storage system that the storage system rejected the direction from the host to change the granularity of the logical-to-physical address map used by the storage system.

9. The host of Claim 1, wherein the memory comprises a three-dimensional memory.
10. The host of Claim 1, wherein the storage system is configured to be removably connected to the host.
11. A method for storage system calibration, the method comprising:
performing the following in a host in communication with a storage system comprising a memory:
tracking host writes to the memory;
determining whether a media management layer used by the storage system results in more sequential-type results or more random-type results than the host writes;
in response to determining that the media management layer used by the storage system results in more sequential-type results, instructing the storage system to make translation mapping performed by the media management layer coarser; and
in response to determining that the media management layer used by the storage system results in more random-type results, instructing the storage system to make translation mapping performed by the media management layer finer.
12. The method of Claim 11, wherein making translation mapping performed by the media management layer coarser comprises:

consolidating host data;
managing new write entries as a delta file until a coarser media management layer is complete; and
flushing the coarser media management layer to the memory.

13. The method of Claim 11, wherein making translation mapping performed by the media management layer finer comprises:

splitting the media management layer into finer entries;
managing new write entries as a delta file until the splitting is complete; and
flushing the split media management layer to the memory.

14. The method of Claim 11, further comprising determining that a number of tracked host writes exceeds a threshold and/or making a positive health determination of the memory.

15. The method of Claim 11, further comprising receiving an indication from the storage system that the storage system rejected the instruction from the host to make the media management layer finer.

16. The method of Claim 11, wherein the tracking is determined in runtime across a plurality of write cycles.

17. A host comprising:

means for determining a pattern of host writes to a memory of a storage system configured to be in communication with the host;

means for determining whether the pattern of host writes matches a granularity of a logical-to-physical address map used by the storage system; and

means for, in response to determining that the pattern of host writes does not match the granularity of the logical-to-physical address map used by the storage system, directing the storage system to change the granularity of the logical-to-physical address map used by the storage system.

18. The host of Claim 17, wherein means for changing the granularity of the logical-to-physical address map comprises means for making the logical-to-physical address map coarser or finer.

19. The host of Claim 17, further comprising means for receiving an indication from the storage system that the storage system rejected the direction from the host to change the granularity of the logical-to-physical address map used by the storage system.

20. The host of Claim 17, wherein the pattern of host writes to the memory is determined in runtime across a plurality of write cycles.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2A




FIG. 2B




FIG. 3




FIG. 4




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/067471

A. CLASSIFICATION OF SUBJECT MATTER

See extra sheet.

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 12/02, G06F 3/06

CPC (20130101) G06F 12/0246, G06F 2212/1016, G06F 2212/7201, G06F 2212/1044, G06F 3/0613, G06F 3/061, G06F 3/064, G06F 3/0638

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: Esp@cenet, Google Patents, Google Scholar, Orbit

Search terms used: storage system granularity pattern writes to the memory granularity finer coarser logical to physical address map sequential non-volatile ftl adjust dynamically map reconfiguration

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"μ-FTL: A Memory-Efficient Flash Translation Layer Supporting Multiple Mapping Granularities"; Proceedings of the 8th ACM international conference on Embedded software; p. 21-30; October 2008; Retrieved from the Internet:<URL: https://dl.acm.org/doi/pdf/10.1145/1450058.1450063 > LEE, Yong-Goo, et al. 31 Oct 2008 (2008/10/31) The whole document	1-20
X	"A dynamic switching flash translation layer based on page-level mapping"; IEICE TRANSACTIONS on Information and Systems, 2016, 99.6: 1502-1511; Retrieved from the Internet:<URL: https://www.jstage.jst.go.jp/article/transinf/E99.D/6/E99.D_2015EDP7406/_pdf > PARK, Dongchul; DEBNATH, Biplob; DU, David HC 06 Jun 2016 (2016/06/06) The whole document	1-10,17-20
A	US 2018293174 A1 JONG et al. 11 Oct 2018 (2018/10/11) The whole document	1-20
A	US 2018067961 A1 YANG et al. 08 Mar 2018 (2018/03/08) The whole document	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

03 Apr 2020

Date of mailing of the international search report

06 Apr 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

DRUCKER Ekaterina

Telephone No. 972-73-9327221

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/067471

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2019056881 A1 FELDMAN et al. 21 Feb 2019 (2019/02/21) The whole document	1-20
A	US 2015278004 A1 LI et al. 01 Oct 2015 (2015/10/01) The whole document	1-20

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/067471

A. CLASSIFICATION OF SUBJECT MATTER:

IPC (20200101) G06F 12/02, G06F 3/06

CPC (20130101) G06F 12/0246, G06F 2212/1016, G06F 2212/7201, G06F 2212/1044, G06F 3/0613, G06F 3/061, G06F 3/064, G06F 3/0638

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2019/067471

Patent document cited search report			Publication date	Patent family member(s)			Publication Date
US	2018293174	A1	11 Oct 2018	US	2018293174	A1	11 Oct 2018
				US	10565123	B2	18 Feb 2020
				WO	2018190962	A1	18 Oct 2018
US	2018067961	A1	08 Mar 2018	US	2018067961	A1	08 Mar 2018
				US	10467195	B2	05 Nov 2019
				KR	20180027326	A	14 Mar 2018
				KR	20180027327	A	14 Mar 2018
				KR	20180027328	A	14 Mar 2018
				KR	20180037571	A	12 Apr 2018
				US	2018067861	A1	08 Mar 2018
				US	10311025	B2	04 Jun 2019
				US	2018067869	A1	08 Mar 2018
				US	10372677	B2	06 Aug 2019
				US	2018067857	A1	08 Mar 2018
				US	10452612	B2	22 Oct 2019
				US	2018069944	A1	08 Mar 2018
				US	10455045	B2	22 Oct 2019
				US	2020028932	A1	23 Jan 2020
US	2019056881	A1	21 Feb 2019	US	2019056881	A1	21 Feb 2019
				US	10365849	B2	30 Jul 2019
				US	2019056882	A1	21 Feb 2019
				US	10503425	B2	10 Dec 2019
US	2015278004	A1	01 Oct 2015	US	2015278004	A1	01 Oct 2015
				US	10318365	B2	11 Jun 2019
				CN	104903864	A	09 Sep 2015
				CN	104903864	B	04 Sep 2018
				EP	2915045	A1	09 Sep 2015

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2019/067471

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
		EP 2915045 A4	14 Jun 2017
		EP 2915045 B1	02 Jan 2019
		TW 201432704 A	16 Aug 2014
		TW I514406 B	21 Dec 2015
		WO 2014070200 A1	08 May 2014