



(51) International Patent Classification:

G06F 3/05 (2006.01) G06F 3/06 (2006.01)

(21) International Application Number:

PCT/US2019/067445

(22) International Filing Date:

19 December 2019 (19.12.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/371,613 01 April 2019 (01.04.2019) US

(71) Applicant: **SANDISK TECHNOLOGIES LLC**
[US/US]; 5080 Spectrum Drive, Suite 1050w, Addison, TX
75001 (US).

(72) Inventors: **SHAKED, Tal**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **GILAD, Omer**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **HOD, Liat**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **SOBOL, Eyal**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **ZILBERSTEIN,**

Einav; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **HAHN, Judah, Gamliel**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US).

(74) Agent: **HETZ, Joseph, F.**; Brinks Gilson & Lione, P.O. Box 10087, Chicago, IL 60610 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

(54) Title: METHOD AND SYSTEM FOR VISUALIZING A CORRELATION BETWEEN HOST COMMANDS AND STORAGE SYSTEM PERFORMANCE




FIG. 1A

(57) Abstract: A method and system for visualizing a correlation between host commands and storage system performance are provided. In one embodiment, a method comprises receiving information concerning host operations of a host performed over a time period; receiving information concerning storage system operations of a storage system performed over the time period; and simultaneously displaying both the host operations and the storage system operations over the time period. Other embodiments are possible, and each of the embodiments can be used alone or together in combination.



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))*

Method and System for Visualizing a Correlation Between Host Commands and Storage System Performance

Cross-Reference to Related Application

[0001] This application is a continuation-in-part of U.S. Patent Application No. 15/347,565, filed November 9, 2016, which is hereby incorporated by reference.

Background

[0002] One metric used when designing a storage system is the write amplification factor. The write amplification factor is defined as the amount of data written to the memory of the storage system divided by the amount of data written by a host. A write amplification factor of one would be ideal as it would provide the best response time and promote high endurance of the memory. However, writing host data often comes with write overhead, such as writing control data in the memory for flash management and possibly relocating data from one pool of blocks to another. Various methods can be used to measure the write amplification factor.

Brief Description of the Drawings

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

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

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

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

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

[0008] Figure 3 is a diagram showing factors that can affect a write amplification factor of an embodiment.

[0009] Figure 4 is a block diagram of a system using a write amplification tool of an embodiment.

[0010] Figure 5 is a flow chart of a method of an embodiment for measuring an amount of data written by a host.

[0011] Figure 6 is a flow chart of a method of an embodiment for measuring an amount of data written to a memory of a storage system.

[0012] Figure 7 is a flow chart of a method of an embodiment for calculating a write amplification factor.

[0013] Figures 8A-8D are graphs generated and displayed by a write amplification tool of an embodiment.

[0014] Figures 9A-9G are graphs generated and displayed by a write amplification tool of an embodiment.

[0015] Figure 10 is a flow chart of a method of an embodiment for visualizing a correlation between host commands and storage system performance.

[0016] Figure 11 is a graph of an embodiment that simultaneously displays information about host operations and information about storage system operations over a time period.

[0017] Figure 12 is a graph of an embodiment of tasks performed by a storage device over time.

[0018] Figure 13 is a graph of an embodiment of tasks performed by a storage device and a host over time.

[0019] Figure 14 is a graph of an embodiment of power consumption.

Detailed Description

[0020] By way of introduction, the below embodiments relate to a method and system for write amplification analysis. In one embodiment, a method is provided that is performed in a computing device. The method comprises determining an amount of data written from the computing device to a storage system over a time period, wherein the storage system comprises a memory; determining an amount of data written to the memory by the storage system over the time period; calculating a write amplification factor over the time period; and simultaneously displaying graphs of the amount of data written from the computing device over the time period, the amount of data written to the memory over the time period, and the write amplification factor over the time period.

[0021] In some embodiments, the method further comprises displaying a graph of consumed capacity over the time period.

[0022] In some embodiments, the method further comprises displaying a graph of over-provisioned blocks of the memory over the time period.

[0023] In some embodiments, the method further comprises displaying a graph of a size of control writes written over the time period.

[0024] In some embodiments, the method further comprises displaying a graph of relocation data over the time period.

[0025] In some embodiments, the amount of data written from the computing device over the time period is determined by monitoring a bus between the computing device and storage system for write commands.

[0026] In some embodiments, the amount of data written to the memory over the time period is determined by monitoring a bus between the memory and a controller of the storage system.

[0027] In some embodiments, the storage system is a simulation model of a storage system.

[0028] In some embodiments, the storage system is a real storage system.

[0029] In some embodiments, the method further comprises calculating an optimization function for a flash management algorithm in the storage system to reduce the write amplification factor.

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

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

[0032] In some embodiments, the storage system is removably connected to a host.

[0033] In another embodiment, a computing device is provided comprising means for gathering information about an amount of data written to a memory of a storage system over a time period and information about an amount of data written from a host to the storage system over the time period; and means for displaying graphical representations synchronized over the time period of activity in the storage system that contributes to the amount of data written to the memory of the storage system over the time period being more than the amount of data written from the host to the storage system over the time period.

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

[0035] In some embodiments, the storage system is an embedded storage system.

[0036] In some embodiments, the storage system is a removable storage system.

[0037] In another embodiment, a computer-readable storage medium storing computer-readable program code that, when executed by a processor, causes the processor to: collect information associated with a write amplification factor, wherein the information is collected over a time period for different write activities; generate graphs based on the information; and display the graphs together on a display device.

[0038] In some embodiments, the storage system comprises a three-dimensional memory.

[0039] In some embodiments, the storage system is embedded in a host.

[0040] In some embodiments, the storage system is removably connected to a host.

[0041] The below embodiments also relate to a method and system for visualizing a correlation between host commands and storage device performance. In one embodiment, a method is presented that is performed in a computing device. The method comprises receiving information concerning host operations of a host performed over a time period; receiving information concerning storage system operations of a storage system performed over the time period; and simultaneously displaying both the host operations and the storage system operations over the time period.

[0042] In some embodiments, the host operations and the storage system operations are simultaneously displayed in a graph.

[0043] In some embodiments, the graph shows when host operations are being performed without storage system operations.

[0044] In some embodiments, the graph indicates a start and a stop of a sync operation.

[0045] In some embodiments, the method further comprises using information displayed on the graph to determine a size of a write buffer.

[0046] In some embodiments, the method further comprises using information displayed on the graph to determine a capacitor size.

[0047] In some embodiments, the method further comprises displaying a graph of power consumption over the time period.

[0048] In some embodiments, the storage system comprises a three-dimensional memory.

[0049] In some embodiments, the storage system is embedded in the host.

[0050] In some embodiments, the storage system is removably connected to the host.

[0051] In another embodiment, a method is provided comprising receiving information regarding activity of a host after the host initiates a process to flush commands to a storage system; receiving information regarding activity of the storage system after the host initiates the process to flush commands to the storage system; and simultaneously displaying the information regarding the activity of the host and the information regarding the activity of the storage system.

[0052] In some embodiments, the simultaneous displaying shows periods of time in which there is host activity without storage system activity after the host flushes commands to the storage system.

[0053] In some embodiments, the information regarding the activity of the host and the information regarding the activity of the storage system are displayed on a graph.

[0054] In some embodiments, the method further comprises displaying a plot of power consumption on the graph.

[0055] In some embodiments, the method further comprises displaying indicators on the graph of when the host initiates and ends the process to flush commands to the storage system.

[0056] In some embodiments, the method further comprises determining a size of a write buffer based on the information that is simultaneously displayed.

[0057] In some embodiments, the method further comprises determining a capacitor size based on the information that is simultaneously displayed.

[0058] In another embodiment, a computing device is provided comprising means for receiving information concerning host operations of a host performed over a time period; means for receiving information concerning storage system operations of a storage system performed over the time period; and means for simultaneously displaying both the host operations and the storage system operations over the time period.

[0059] In some embodiments, the computing device further comprises means for displaying a graph showing power consumption over the time period.

[0060] In some embodiments, the computing device further comprises means for displaying indicators of a start and stop of a flush operation.

[0061] 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.

[0062] Turning now to the drawing, storage systems suitable for use in implementing aspects of these embodiments are shown in Figures 1A -1C. Figure 1A is a block diagram illustrating a non-volatile storage system 100 according to an embodiment of the subject matter described herein. Referring to Figure 1A, non-volatile storage system 100 includes a controller 102 and non-volatile memory that may be made up of one or more non-volatile memory die 104. As used herein, the term die refers to the collection of non-volatile memory cells, and associated circuitry for managing the physical operation of those non-volatile memory cells, that are formed on a single semiconductor substrate. (The terms “memory” and “media” may be used interchangeably herein.) Controller 102 interfaces with a host system and transmits command sequences for read, program, and erase operations to non-volatile memory die 104.

[0063] The controller 102 (which may be a flash memory controller) can take the form of processing circuitry, a microprocessor or processor, and a computer-readable medium that stores computer-readable program code (e.g., firmware) executable by the (micro)processor, logic gates, switches, an application specific integrated circuit (ASIC), a programmable logic controller, and an embedded microcontroller, for example. The controller 102 can be configured with hardware and/or firmware to perform the various functions described below and shown in the flow diagrams. Also, some of the components shown as being internal to the controller can also be stored external to the controller, and other components can be used. Additionally, the phrase “operatively in communication with” could mean directly in communication with or indirectly (wired or wireless) in communication with through one or more components, which may or may not be shown or described herein.

[0064] As used herein, a flash memory controller is a device that manages data stored on flash memory and communicates with a host, such as a computer or electronic device. A flash memory controller can have various functionality in addition to the specific

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

[0065] Non-volatile memory die 104 may include any suitable non-volatile storage medium, including NAND flash memory cells and/or NOR flash memory cells. The memory cells can take the form of solid-state (e.g., flash) memory cells and can be one-time programmable, few-time programmable, or many-time programmable. The memory cells can also be single-level cells (SLC), multiple-level cells (MLC), triple-level cells (TLC), or use other memory cell level technologies, now known or later developed. Also, the memory cells can be fabricated in a two-dimensional or three-dimensional fashion.

[0066] 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.

[0067] Although, in the example illustrated in Figure 1A, non-volatile storage system 100 (sometimes referred to herein as a storage module) includes a single channel between controller 102 and non-volatile memory die 104, the subject matter described herein is not limited to having a single memory channel. For example, in some NAND storage

system architectures (such as the ones shown in Figures 1B and 1C), 2, 4, 8 or more NAND channels may exist between the controller and the NAND memory device, depending on controller capabilities. In any of the embodiments described herein, more than a single channel may exist between the controller and the memory die, even if a single channel is shown in the drawings.

[0068] Figure 1B illustrates a storage module 200 that includes plural non-volatile storage systems 100. As such, storage module 200 may include a storage controller 202 that interfaces with a host and with storage system 204, which includes a plurality of non-volatile storage systems 100. The interface between storage controller 202 and non-volatile storage systems 100 may be a bus interface, such as a serial advanced technology attachment (SATA) or peripheral component interconnect express (PCIe) interface. Storage module 200, in one embodiment, may be a solid state drive (SSD), such as found in portable computing devices, such as laptop computers, and tablet computers.

[0069] 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.

[0070] Figure 2A is a block diagram illustrating components of controller 102 in more detail. Controller 102 includes a front end module 108 that interfaces with a host, a back end module 110 that interfaces with the one or more non-volatile memory die 104, and various other modules that perform functions which will now be described in detail. For example, in this embodiment, the controller 102 comprises a NAND bus recorder 111, which can be implemented in hardware or software/firmware and is configured to record traffic on the bus to the memory 104. The use of the NAND bus recorder 111 will be discussed in more detail below. Instead of or in addition to the NAND bus recorder 111 in the storage system 100, an external NAND bus monitor and/or an internal host bus monitor can be used. Also, in one embodiment, part of the computing element of write

amplification is in the controller 102 of the storage system 100 (by using the NAND bus monitor and host bus monitor). 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.

[0071] 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.

[0072] 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, a serial advanced technology attachment (SATA), SATA Express, serially attached small computer system interface (SAS), Fibre Channel, universal serial bus (USB), peripheral component interface express (PCIe), and Non-Volatile Memory Express (NVMe). The host interface 120 typically facilitates transfer for data, control signals, and timing signals.

[0073] 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 Redundant Array of Independent Drives (RAID) 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.

[0074] 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.

[0075] Figure 2B is a block diagram illustrating components of non-volatile memory die 104 in more detail. Non-volatile memory die 104 includes peripheral circuitry 141 and non-volatile memory array 142. Non-volatile memory array 142 includes the non-volatile memory cells used to store data. The non-volatile memory cells may be any suitable non-volatile memory cells, including NAND flash memory cells and/or NOR flash memory cells in a two dimensional and/or three dimensional configuration. Non-volatile memory die 104 further includes a data cache 156 that caches data.

[0076] As discussed above, one metric used when designing a storage system is the write amplification (WA) factor. The write amplification factor is defined as the amount of data written to the memory of the storage system divided by the amount of data written by a host. Expressed mathematically, the write amplification factor (WAF) is defined as:

$$WAF = \frac{\text{Amount of Data written to NAND}}{\text{Amount of Data written by Host}}$$

[0077] A write amplification factor of one would be ideal as it would provide the best response time and promote high endurance of the memory. However, writing host data often comes with write overhead, and Figure 3 shows some of the factors that can affect the write amplification factor. These factors include writing control data in the memory for flash management and possibly relocating data from one pool of blocks to another. As used herein, control data can refer to the additional writes required for data structures

for flash management. The amount of control data written depends on the write scenario. For example, random writes on wide address ranges may require more frequent updates to the logical-to-physical address table than sequential writes.

[0078] Relocating data can take the form of folding or garbage collection. In one embodiment, folding refers to moving data from a block of single-level cells (SLC) to a block of multi-level or triple-level cells (MLC or TLC). In contrast, garbage collection refers to moving data between blocks of the same memory cell type (e.g., MLC to MLC, or TLC to TLC). As shown in Figure 3, garbage collection can depend on overprovisioning (e.g., availability of free TLC/SLC blocks beyond the exported capacity of the memory). Overprovisioning can be determined by the amount of spare blocks at production time and also by the amount of blocks unmapped by the host and refers to how many spare blocks the system has for writing new incoming data. For example, when a host wants to write data to a block that already contains data, the storage system can write the incoming data to a spare block and then invalidate the old block in the logical-to-physical address table. If there were no spare blocks available in this scenario, the old data would need to be evacuated from the target block before the incoming host data could be written. This could lead to performance issues. Accordingly, the more spare blocks available in memory, the easier it is to accommodate random write scenarios.

[0079] Fragmentation of media is another vector under garbage collection and is measured by a metric referred to valid count (VC). Fragmentation occurs when there are blocks that are partially written because some of them were invalidated. A block that has a relatively-small amount of valid data is a good candidate for garbage collections because not much data needs to be moved. In contrast, if a 1 MB block has only 16 KB of invalid data, much more data would need to be moved, resulting in more write overhead.

[0080] Returning again to the drawings, Figure 4 is a block diagram of a system of an embodiment for calculating and analyzing a write amplification factor. As shown in Figure 4, in this embodiment, a computer 400 (also referred to herein as a computing device) and display 410 are provided. These components can take any suitable form. For example, the computer 400 can be a personal computer or server, and the display 410 can

be a stand-alone monitor. Alternatively, the computer 400 can be a mobile device that has the display 410 integrated in it (e.g., as a touch screen). Of course, these are merely examples, and other implementations can be used. For example, in an embodiment where part of the computing element of write amplification is in the controller 102 of the storage system (by using the NAND bus monitor and host bus), the computer 400 can just read the data and monitor it periodically.

[0081] The computer 400 in this embodiment comprises a processor 420 and a memory 430. The processor 420 is configured to implement a write amplification analysis tool 440. In one embodiment, computer-readable program code of the write amplification analysis tool 440 is stored in the memory 430 and is executed by the processor 420 (i.e., the write amplification analysis tool 440 can be software/firmware executed by hardware). In another embodiment, the write amplification analysis tool 440 is implemented exclusively in hardware. In any event, in one embodiment, the write amplification analysis tool 440 can be used to implement the algorithms shown in the attached flowcharts and described herein. In one embodiment, the write amplification analysis tool 440 can be used to perform write amplification analysis on a simulation of a storage system or can be used to perform write amplification analysis on an actual storage system 100 connected to the computer 400 by a bus 450. The computer 400 can also be used to generate and display terabytes written (TBW), which specifies how many terabytes can be written to the memory 104 until it cannot absorb any more data.

[0082] Turning again to the drawings, Figures 5-7 are flow charts 500, 600, 700 of methods of an embodiment that can be used to calculate the write application factor. As discussed above, the underlying data to generate the write application factor is the amount of data written by a host and the amount of data written to the memory 104. Examples of methods that can be used to collect this data are shown in Figures 5 and 6, respectively.

[0083] Starting first with Figure 5, Figure 5 is a flow chart 500 of a method of an embodiment for measuring an amount of data written by a host (e.g., the computer 400 in Figure 4). As shown in Figure 5, the write amplification analysis tool 440 first divides the time domain into fixed window sizes (e.g., 100 milliseconds (ms)) (act 510). Then, the write amplification analysis tool 440 monitors the bus 450 between the computer 400

and the storage system 100 to get an input trace with the following information: command type, command size, command start time, and command duration (act 520). Starting with time window #1, the write amplification analysis tool 440 finds all the write commands whose start times fall into this window (act 530). The write amplification analysis tool 440 then calculates the amount of data written by using the following formula: $\text{Sum}(\text{number_of_sectors} \times \text{sector_size})$ (act 540). Then, the write amplification analysis tool 440 calculates the write performance by dividing the amount of data by the window size (act 550). The write amplification analysis tool 440 repeats these acts (530, 540, 550) for each window. After all the windows have been processed, the write amplification analysis tool 440 has measured the amount of data written by the computer 400.

[0084] As shown in the flow chart 600 in Figure 6, the write amplification analysis tool 440 then measures the amount of data written to the memory 104 of storage system 100. While the measurement is taken on an actual storage system 100 in this example, as mentioned above, the analysis can be performed on a simulation. As shown in Figure 6, in this embodiment, the write amplification analysis tool 440 divides the time domain into fixed-size windows (e.g., 100 ms) (act 610). The write amplification analysis tool 440 then receives input from the NAND bus recorder 111 and uses a protocol analyzer on the record activity on the NAND bus to extract command type and size (act 620). For example, the amount of data for control or relocation can be taken from NAND bus recorder 111, which occurs in parallel and is synchronized to the host activity, and address behavior can be taken from a recording utility in the host, such as FTrace.

[0085] Starting with time window #1, the write amplification analysis tool 440 finds all the write commands those start times fall into this window (act 630). The write amplification analysis tool 440 then calculates the amount of data written by using the following formula: $\text{Sum}(\text{number_of_sectors} \times \text{sector_size})$ (act 640). Then, the write amplification analysis tool 440 calculates the write performance by dividing the amount of data by the window size (act 650). The write amplification analysis tool 440 repeats these acts (630, 640, 650) for each window. After all the windows have been processed, the write amplification analysis tool 440 has measured the amount of data written to the memory 104 of storage system 100.

[0086] Turning now to the flow chart in Figure 7, to calculate the write amplification factor, the write amplification factor analysis tool 440 starts with time window #1 (act 710) can divide the NAND performance for that window (as determined in the flow chart 600 of Figure 6) with the host performance of that window (as determined in the flow chart 500 of Figure 5) (act 720). The write amplification factor analysis tool 440 repeats this process of the other windows.

[0087] With this information now gathered, the write amplification factor analysis tool 440 can simultaneously display graphs on the display 410 of the amount of data written from the host, the amount of data written to the memory, and the write amplification value synchronized over the time period for various write scenarios, so a storage system designer can see the various effects on the write analysis factor over time. For example, Figure 8A shows five different write and erase activities over time (sequential write over the entire media, random write on 4GB, discard/unmap ½ the media, random write on 4GB, sequential write over ½ the media, and random write over the entire media), and Figures 8B-8D shown the effect of those activities, over time, on measured data written by the host (Figure 8B), measured data written to the NAND (Figure 8C), and the calculated write amplification factor (Figure 8D).

[0088] Of course the above graphs are merely examples, and different or other types of graphs can be used. Some of these additional graphs are shown in Figures 9A-9G.

[0089] Figure 9A shows the behavior over time of the commands in the address space. Figure 9B (like Figure 8D) shows the write amplification factor over time derived from all the factors discussed above. Figure 9C shows overprovisioning both in terms of free blocks and flash management units (FMUs) (e.g., 4K bytes). In case of a random write, the amount of free FMUs stays the same, but the amount of available blocks may decrease over time until a garbage collection threshold is achieved, thus activating the garbage collection process. Figure 9D shows data written to the NAND (like Figure 8C). Figure 9E shows NAND data – host data. Figure 9F shows the excessive data generated only from control activity. As will be discussed below, sequential writes usually have very little control updates, while pure random writes generate much more table updates. Additional garbage collection generates even more control data. Figure 9G shows relocation data when garbage collection is initiated.

[0090] As shown in these drawings, for a sequential write, data is written straight to the memory with sequential address (Figure 9A) with a minimum amount of control data needed due to minimum updates to the logical-to-physical address table (Figure 9F).

Thus, as shown in Figure 9B, the write amplification factor is almost 1, meaning that the amount of data written to the memory is almost all from the data written by the host. As shown in Figure 9C, in this example, writing to the exported capacity of the entire media causes the overprovisioned flash management units (FMUs) and the overprovisioned blocks to reach a level just above the garbage collection threshold (i.e., the memory is filled to its exported capacity, and all that remains is a minimum level of spare blocks).

[0091] Next, for the random write in the 4GB range, there is a jump up in the write amplification factor (Figure 9B) because random writes would invalidate data in different blocks, and the logical-to-physical address table would need to be updated due to this host activity (Figure 9F). However, as shown in Figure 9C, part way into the random write, the number of overprovisioned blocks drops below the garbage collection threshold. This will cause the controller 102 to perform garbage collection and reallocate data from old blocks to new blocks (Figure 9G), which will also cause the controller 102 to update the logical-to-physical address table from this garbage collection activity (Figure 9F). Writing this extra control data due to garbage collection causes another jump up in the write amplification factor (Figure 9B).

[0092] Next, the host sends a discards/unmap command (i.e., the address is not relevant) to half of the media. This means that, almost immediately, half of the exported capacity is available again for use, although the data in those blocks may or may not be erased (Figure 9A), and there is a jump in the overprovisioning (Figure 9C).

Additionally, as shown in Figure 9F, there is a small bump in the amount of control data due to an update in the logical-to-physical address table to make blocks as available.

[0093] Following the discard command, there is another random write in the 4GB range. This time, because half the media is available for storage, the data can be written without triggering garbage collection. As such, the logical-to-physical address table only needs to be updated with host activity (and not with garbage collection activity, as in the prior random write), so the write amplification factor is less than that of the prior random write (Figure 9F). Next, there is a sequential write to half the media. This is similar to

the sequential write discussed earlier. Finally, there is a random write to the entire media. This is also similar to the random write discussed earlier, but, since there is more data being written now, the write amplification factor, control data, and relocation data have increased.

[0094] There are several advantages associated with these embodiments. For example, the write amplification factor analysis tool disclosed herein provides a powerful analysis tool that shows the dynamic behavior of write amplification and can cross correlate write amplification with other available information. This tool allows a storage system designer to analyze real use cases having a plethora of read, write, discard, flush, etc. operations in different sizes and addresses. Unlike prior methods that just generate a value for the write amplification factor, these embodiments allow a storage system designer to see the effect that real-life write use cases have over time on the write amplification factor by displaying graphs of various metrics simultaneously and synchronized to the same time scale.

[0095] Displaying all this information in a synchronized way in the same time scale shows the dynamic behavior of write amplification over time and in correlation with various input scenarios. For example, showing how much data is written by a host versus how much data is written to the memory of the storage system for various types of write activity can show the storage system designer when a problem occurs and why it is occurring. Using these graphs, the storage system designer can determine what adjustments to make to the flash management algorithm in the storage system 100 in order to improve or optimize the write amplification factor to improve response time and avoid reducing the endurance of the memory. As used herein, the “flash management algorithm” refers to the hardware, software, and/or firmware in the storage system 100 that controls the operation of the storage system 100, and, in particular, affects the relocation and control contributors to write amplification. That is, after looking at the various graphs, a storage system designer can modify the flash management algorithm to reduce the write amplification factor (e.g., the average or peaks) for various scenarios.

[0096] For example, as shown in the figures and discussed above, performing a random write of the entire media provides the worst write application factor. To reduce this, the storage system designer can change the structure of the logical-to-physical

address table or change the policy of when that table is updated. In general, updating the table less often increase the risk of a loss-of-data if there is a sudden power loss.

However, if a sudden power loss is unlikely, decreasing the amounts of table updates may be acceptable to improve performance. The write amplification factor analysis tool can also be configured to automatically or semi-automatically (e.g., with designer input) calculate an optimization function for the flash management algorithm to reduce the write amplification factor. For example, the write amplification analysis tool can be used to optimize the storage system's firmware code for folding, garbage collection, etc.

[0097] In another embodiment, a method and system are provided for visualizing a correlation between host commands and storage system performance. Storage systems, which are sometimes referred to herein as "storage devices" or "devices," (such as solid-state drives (SSDs) embedded in a mobile device, such as a phone, tablet, or wearable device) can execute many input/output operations (e.g., read, write, trim, and flush) during use of the mobile device. The operations may include characteristics (such as timestamps on initiation and completion) and peripheral data (such as power state and aggregate queue depth). Analysis of these input/output commands and their characteristics can be used to design and implement algorithms for data storage. In operation, an application running in the storage system can log the various operations that take place, and this log can be analyzed by a computing device (e.g., the host or another device). In data analysis environments for storage workloads, there are often millions of individual data points that represent specific characteristics of the input/output operations sent from the mobile device (host) to the storage system, and vice versa.

[0098] When there is a performance or power issue with the host interacting with the storage system, it may be assumed that the problem is with the storage system, and the data analysis mentioned above can be used to help identify and solve the issue. For example, there is inefficiency in the data storage algorithm, the analysis of the storage system's input-output operations can be used to identify the cause of the inefficiency and suggest a solution. Examples of such analysis can be found in U.S. Patent Application No. 15/347,565, filed November 9, 2016, and U.S. Patent Application No. 15/226,661, filed August 2, 2016, both of which are hereby incorporated by reference.

[0099] However, there can be situations where the performance or power issue is caused at least in part by the host. In such situations, analyzing storage system operations alone would not identify that the host is contributing to the problem or what the source of the problem is. The following embodiments can be used to address this situation.

[00100] As shown in the flow chart 1000 in Figure 10, in one embodiment, a computing device 400 (e.g., a personal computer (PC) or server) receives information from the host concerning host operations (events) performed over a time period (act 1010) and also receives information from the storage system 100 concerning storage system operations (events) performed over the time period (act 1020). For example, the host and storage system 100 can provide their logs to the computing device 400 over a wired or wireless connection. The logs can include a list of the various commands or operations that took place and various related data, such as, but not limited to, timestamps on initiation and completion, power state, aggregate queue depth, etc. In addition to or instead of using a log, host and/or storage system activity can be measured by monitoring traffic on a bus used with those components. In one embodiment, the computing device 400 is in communication with a server that provides the computing device 400 with computer-readable program code to execute software that performs the algorithm shown in the flow chart 1000 in Figure 10.

[00101] The computing device 400 plots these events in time order on the same graph (act 1030), the result of which is a simultaneous display of both the host operations and the storage system operations over the time period. Figure 11 is an example of such a graph, which shows host and storage system tasks over time. As shown in Figure 11, the host performs a number of file write operations internal to the host. These operations can be, for example, the host's operating system queuing-up write commands in the host's memory after a user calls a write function. In this example, there are five write operations, each represented by a different rectangle, with the first write operation being larger than the others. At this point in time, there are no storage system operations, as the host has not provided the write commands to the storage system 100.

[00102] At some point (e.g., because of a host-triggered time-out or because a host application calls for it), the host decides to flush the write commands from its internal memory to the storage system 100 by calling a kernel application program interface

(API) to start the synchronization (sync) operation. This is indicated by the Sync Start arrow in Figure 11. At that point, the host performs additional operations for a period of time. For example, the host might be coalescing one or more write commands together, putting the write commands in a queue, updating internal tables in the host's file system, determining which page of memory are "dirty," or even perform an operation that is unrelated to the write commands. Accordingly, in this example, there is a delay between Sync Start and the time the storage system 100 actually starts performing the write operations.

[00103] At some point, the storage system 100 is ready to perform the write operations. Figure 11 shows that the storage system 100 performs four write operations sequentially (in this example, two of the five write operations were coalesced). Figure 11 also shows that, in this example, the host performs additional operations after the storage system 100 has completed the write operations. At some time later, the host calls the kernel API to end the sync operation. This is indicated by the Sync End arrow in Figure 11. Accordingly, in this example, there is a delay between the time the storage system 100 finishes performing the write operations and Sync End. This delay and the delay after Sync Start is a measure of file system overhead.

[00104] As shown by Figure 11, by simultaneously displaying information about host and storage system operations over the same time period, this embodiment helps a technician visualize a correlation between host commands and storage system performance. Here, the graph in Figure 11 shows that relatively little time between Sync Start and Sync End is consumed by storage system operations (act 1040). If the technician did not have the benefit of this graph, he may assume that the storage system 100 used the entire time between Sync Start and Sync End to perform the write operation and not realized that the delay is mostly from host overhead. According, if there is a performance issue experienced when the host interacts with the storage system 100, this graph indicates that the problem is on the host side – not the storage system side. With this knowledge, the technician can focus his efforts in improving the efficiency of the host instead of the storage system 100. For example, the technician may use the information on the graph to determine the optimal size of the write buffer and/or the capacitor size used for flushing data in case of ungraceful shutdown (UGSD).

[00105] This embodiment can also be used to derive storage system power and throughput as a cost of host activity (act 1050). In this way, this embodiment can be used to visualize the correlation between host commands and device performance and power consumption per command sequence and storage device and host platform. This embodiment will now be discussed in conjunction with Figures 12-14.

[00106] Figure 12 shows a graph of tasks performed by the storage device 100 over time. This graph does not show host activity (i.e., the host activity has been filtered out of this graph). Figure 13 adds the host activity. Figure 13 is similar to Figure 11 in that it shows writes to the host memory before the sync operation, as well as host activity during the sync operation. However, the host activity during the sync operation is less than in Figure 11. Finally, Figure 14 adds a plot of power consumption over time onto the graph. As can be seen in Figure 14, power consumption is a function of both host operations and storage system operations. This graph shows how much of the power consumed is due to host overhead versus the storage system during a write operation. Such a graph may be helpful in trying to identify what is contributing to draining the battery life of the host.

[00107] There are several advantages associated with these embodiments. For example, the task-over-time plot generated by these embodiments can be used to evaluate file system performance, power, and throughput per command sequence and per storage system. It also allows for comparing different platform file system overheads. Further, there are several alternatives that can be used with these embodiments. For example, while the host and storage system activity was shown as being simultaneously displayed on the same graph, in other embodiment, the activity is simultaneously displayed on different graphs that are displayed along with one another, or even in a chart or some other non-graphical form.

[00108] 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.

[00109] 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.

[00110] 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.

[00111] 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.

[00112] 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.

[00113] The memory elements may be arranged in the single memory device level in an ordered array, such as in a plurality of rows and/or columns. However, the memory elements may be arrayed in non-regular or non-orthogonal configurations. The memory elements may each have two or more electrodes or contact lines, such as bit lines and word lines.

[00114] 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).

[00115] 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.

[00116] 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.

[00117] 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.

[00118] 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.

[00119] 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.

[00120] 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.

[00121] 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 method comprising:
performing the following in a computing device:
receiving information concerning host operations of a host performed over a time period;
receiving information concerning storage system operations of a storage system performed over the time period; and
simultaneously displaying both the host operations and the storage system operations over the time period.
2. The method of Claim 1, wherein the host operations and the storage system operations are simultaneously displayed in a graph.
3. The method of Claim 2, wherein the graph shows when host operations are being performed without storage system operations.
4. The method of Claim 2, wherein the graph indicates a start and a stop of a synchronization operation.
5. The method of Claim 2 further comprising using information displayed on the graph to determine a size of a write buffer.

6. The method of Claim 2 further comprising using information displayed on the graph to determine a capacitor size.
7. The method of Claim 1, further comprising displaying a graph of power consumption over the time period.
8. The method of Claim 1, wherein the storage system comprises a three-dimensional memory.
9. The method of Claim 1, wherein the storage system is embedded in the host.
10. The method of Claim 1, wherein the storage system is removably connected to the host.
11. A method comprising:
 - performing the following in a computing device:
 - receiving information regarding activity of a host after the host initiates a process to flush commands to a storage system;
 - receiving information regarding activity of the storage system after the host initiates the process to flush commands to the storage system; and
 - simultaneously displaying the information regarding the activity of the host and the information regarding the activity of the storage system.

12. The method of Claim 11, wherein the simultaneous displaying shows periods of time in which there is host activity without storage system activity after the host flushes commands to the storage system.
13. The method of Claim 11, wherein the information regarding the activity of the host and the information regarding the activity of the storage system are displayed on a graph.
14. The method of Claim 13, further comprising displaying a plot of power consumption on the graph.
15. The method of Claim 13, further comprising displaying indicators on the graph of when the host initiates and ends the process to flush commands to the storage system.
16. The method of Claim 11, further comprising determining a size of a write buffer based on the information that is simultaneously displayed.
17. The method of Claim 11, further comprising determining a capacitor size based on the information that is simultaneously displayed.
18. A computing device comprising:
means for receiving information concerning host operations of a host performed over a time period;

means for receiving information concerning storage system operations of a storage system performed over the time period; and

means for simultaneously displaying both the host operations and the storage system operations over the time period.

19. The computing device of Claim 18, further comprising means for displaying a graph showing power consumption over the time period.

20. The computing device of Claim 18, further comprising means for displaying indicators of a start and stop of a flush operation.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2A




FIG. 2B




FIG. 3




FIG. 4




FIG. 5




FIG. 6








FIG. 7






FIG. 9A




FIG. 9B




FIG. 9C




FIG. 9D




FIG. 9E




FIG. 9F




FIG. 9G




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/067445

A. CLASSIFICATION OF SUBJECT MATTER IPC (20200101) G06F 3/05, G06F 3/06 CPC (20180501) G06F 3/05, G06F 3/0604, G06F 3/0673 According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) IPC (20200101) G06F 3/06 CPC (20130101) G06F 3/0604, G06F 3/0653 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		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2018129448 A1 DUZLY YACOV[IL]; SOBOL EYAL[IL]; SHAKED TAL[IIL]; HOD LIAT[IL]; GILAD OMER[IL]; INNA ZEVULUN EINAT[IL] 10 May 2018 (2018/05/10) whole document	1-20
A	US 9128995 B1 SPLUNK INC[US]; SPLUNK INC[US] + (SPLUNK INC, ; SPLUNK, INC) 08 Sep 2015 (2015/09/08) whole document	1-20
A	US 2015227459 A1 SONY CORP[JP] + (SONY CORPORATION, ; Sony Corporation) 13 Aug 2015 (2015/08/13) 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 05 Apr 2020		Date of mailing of the international search report 05 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 ZOZULYA Irina Telephone No. 972-73-3927253

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2019/067445

Patent document cited search report			Publication date	Patent family member(s)			Publication Date
US	2018129448	A1	10 May 2018	US	2018129448	A1	10 May 2018
				US	10296260	B2	21 May 2019
				CN	109690468	A	26 Apr 2019
				DE	112017005637	T5	25 Jul 2019
				US	2019227735	A1	25 Jul 2019
				US	10564888	B2	18 Feb 2020
				WO	2018089097	A1	17 May 2018
US	9128995	B1	08 Sep 2015	US	9128995	B1	08 Sep 2015
				US	9130832	B1	08 Sep 2015
				US	9130860	B1	08 Sep 2015
				US	9146954	B1	29 Sep 2015
				US	9146962	B1	29 Sep 2015
				US	9158811	B1	13 Oct 2015
				US	9208463	B1	08 Dec 2015
				US	9210056	B1	08 Dec 2015
				US	9245057	B1	26 Jan 2016
				US	9286413	B1	15 Mar 2016
				US	9294361	B1	22 Mar 2016
				US	2016105330	A1	14 Apr 2016
				US	2016105325	A1	14 Apr 2016
				US	9491059	B2	08 Nov 2016
				US	2016226720	A1	04 Aug 2016
				US	9521047	B2	13 Dec 2016
				US	2016103907	A1	14 Apr 2016
				US	9584374	B2	28 Feb 2017
				US	2016105335	A1	14 Apr 2016
				US	9590877	B2	07 Mar 2017
				US	2016147380	A1	26 May 2016
				US	9596146	B2	14 Mar 2017

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2019/067445

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	US	2016103889 A1	14 Apr 2016
	US	9614736 B2	04 Apr 2017
	US	2016103890 A1	14 Apr 2016
	US	9747351 B2	29 Aug 2017
	US	2016103862 A1	14 Apr 2016
	US	9753961 B2	05 Sep 2017
	US	2016103878 A1	14 Apr 2016
	US	9755912 B2	05 Sep 2017
	US	2016103909 A1	14 Apr 2016
	US	9755913 B2	05 Sep 2017
	US	2016103559 A1	14 Apr 2016
	US	9760240 B2	12 Sep 2017
	US	2016103891 A1	14 Apr 2016
	US	9760613 B2	12 Sep 2017
	US	2017034014 A1	02 Feb 2017
	US	9762455 B2	12 Sep 2017
	US	2016105334 A1	14 Apr 2016
	US	9838280 B2	05 Dec 2017
	US	2016103887 A1	14 Apr 2016
	US	9864797 B2	09 Jan 2018
	US	2017141969 A1	18 May 2017
	US	9960970 B2	01 May 2018
	US	2017019487 A1	19 Jan 2017
	US	9967351 B2	08 May 2018
	US	2016105338 A1	14 Apr 2016
	US	9985863 B2	29 May 2018
	US	2016216855 A1	28 Jul 2016
	US	10152561 B2	11 Dec 2018
	US	2017046374 A1	16 Feb 2017
	US	10193775 B2	29 Jan 2019
	US	2017017368 A1	19 Jan 2017

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2019/067445

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	US 10198155	B2	05 Feb 2019
	US 2017046127	A1	16 Feb 2017
	US 10209956	B2	19 Feb 2019
	US 2016104076	A1	14 Apr 2016
	US 10235638	B2	19 Mar 2019
	US 10305758	B1	28 May 2019
	US 2017339024	A1	23 Nov 2017
	US 10331742	B2	25 Jun 2019
	US 2017339029	A1	23 Nov 2017
	US 10333799	B2	25 Jun 2019
	US 2018041402	A1	08 Feb 2018
	US 10380189	B2	13 Aug 2019
	US 2017147681	A1	25 May 2017
	US 10417108	B2	17 Sep 2019
	US 2017083572	A1	23 Mar 2017
	US 10417225	B2	17 Sep 2019
	US 2016103888	A1	14 Apr 2016
	US 10447555	B2	15 Oct 2019
	US 2016103883	A1	14 Apr 2016
	US 10474680	B2	12 Nov 2019
	US 2017329462	A1	16 Nov 2017
	US 10503348	B2	10 Dec 2019
	US 2017322985	A1	09 Nov 2017
	US 10503745	B2	10 Dec 2019
	US 2017337252	A1	23 Nov 2017
	US 10503746	B2	10 Dec 2019
	US 10505825	B1	10 Dec 2019
	US 10515096	B1	24 Dec 2019
	US 2017351720	A1	07 Dec 2017
	US 10521409	B2	31 Dec 2019
	US 2016366036	A1	15 Dec 2016

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2019/067445

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	US 10536353	B2	14 Jan 2020
	US 2018248960	A1	30 Aug 2018
	US 10547695	B2	28 Jan 2020
	US 2018075134	A1	15 Mar 2018
	US 10565241	B2	18 Feb 2020
	US 2018052912	A1	22 Feb 2018
	US 10572518	B2	25 Feb 2020
	US 2018241660	A1	23 Aug 2018
	US 10572541	B2	25 Feb 2020
	US 2016103838	A1	14 Apr 2016
	US 10592093	B2	17 Mar 2020
	US 2016103908	A1	14 Apr 2016
	US 2016103918	A1	14 Apr 2016
	US 2016104090	A1	14 Apr 2016
	US 2016104091	A1	14 Apr 2016
	US 2016104093	A1	14 Apr 2016
	US 2016105329	A1	14 Apr 2016
	US 2016132575	A1	12 May 2016
	US 2016292611	A1	06 Oct 2016
	US 2016294606	A1	06 Oct 2016
	US 2017004433	A1	05 Jan 2017
	US 2017140071	A1	18 May 2017
	US 2017155558	A1	01 Jun 2017
	US 2017195195	A1	06 Jul 2017
	US 2018024901	A1	25 Jan 2018
	US 2018234307	A1	16 Aug 2018
	US 2018349482	A1	06 Dec 2018
	US 2018365309	A1	20 Dec 2018
	US 2019095478	A1	28 Mar 2019
	US 2019147363	A1	16 May 2019
	US 2020019555	A1	16 Jan 2020

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2019/067445

Patent document cited search report			Publication date	Patent family member(s)			Publication Date
US	2015227459	A1	13 Aug 2015	US	2015227459	A1	13 Aug 2015
				US	9836392	B2	05 Dec 2017
				CN	103324443	A	25 Sep 2013
				CN	103324443	B	05 Jan 2018
				JP	2013225368	A	31 Oct 2013
				JP	5853899	B2	09 Feb 2016
				US	2013254498	A1	26 Sep 2013
				US	9058162	B2	16 Jun 2015