



(51) International Patent Classification:
G06F 12/02 (2006.01)

(21) International Application Number:
PCT/US2017/035201

(22) International Filing Date:
31 May 2017 (31.05.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/198,015 30 June 2016 (30.06.2016) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).

(72) Inventors: **KHAN, Jawad B.**; 826 N. 26th Avenue, Cornelius, Oregon 97113 (US). **GOPAL, Vinodh**; 70 Townsend Blvd, Westborough, Massachusetts 01581 (US). **TRIKA, Sanjeev N.**; 4023 NW Riggs Drive, Portland, Oregon 97229 (US).

(74) Agent: **KELLETT, Glen M.**; Barnes & Thornburg LLP, 11 South Meridian Street, Indianapolis, Indiana 46204 (US).

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

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: TECHNOLOGIES FOR ADDRESSING DATA IN A MEMORY




FIG. 1

(57) **Abstract:** Technologies for addressing data in a memory include an apparatus that includes a memory and a controller. The memory is to store sub-blocks of data in a data table and a pointer table of locations of the sub-blocks in the data table. The controller is to manage the storage and lookup of data in the memory. Further, the controller is to store a sub-block pointer in the pointer table to a location of a sub-block in the data table and store a second pointer that references an entry where the sub-block pointer is stored in the pointer table.



TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

TECHNOLOGIES FOR ADDRESSING DATA IN A MEMORY

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application claims priority to U.S. Utility Patent Application Serial No. 15/198,015, entitled “TECHNOLOGIES FOR ADDRESSING DATA IN A MEMORY,” which was filed on June 30, 2016.

BACKGROUND

[0002] Search and replace operations are common on computers in many applications, including office applications and databases, as well as on servers used for web searches. Given that the entire data set to be searched is read from a storage device into memory before the set is searched by the CPU, these operations are generally bandwidth intensive, CPU intensive, and time intensive. For example, in a database search/select operation, a significant amount of data is read from the storage device to system memory before the CPU performs the search. The CPU then compares the entire data set against a value or a range of values, and only matching entries are kept while the rest are discarded.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The concepts described herein are illustrated by way of example and not by way of limitation in the accompanying figures. For simplicity and clarity of illustration, elements illustrated in the figures are not necessarily drawn to scale. Where considered appropriate, reference labels have been repeated among the figures to indicate corresponding or analogous elements.

[0004] FIG. 1 is a simplified block diagram of at least one embodiment of a data storage device for storing and efficiently locating data;

[0005] FIG. 2 is a simplified block diagram of at least one embodiment of an environment that may be established by the data storage device of FIG. 1;

[0006] FIGS. 3-4 are a simplified flow diagram of at least one embodiment of a method for writing data to enable efficient searching of the data that may be executed by the data storage device of FIGS. 1 and 2;

[0007] FIG. 5 is a simplified flow diagram of at least one embodiment of a method for searching for data that may be executed by the data storage device of FIGS. 1 and 2;

[0008] FIG. 6 is a simplified flow diagram of at least one embodiment of a method for replacing data that may be executed by the data storage device of FIGS. 1 and 2;

[0009] FIG. 7 is a simplified block diagram of at least one embodiment of the writing of data by the data storage device of FIGS. 1 and 2 in accordance with the method of FIGS. 3-4;

[0010] FIG. 8 is a simplified block diagram of at least one embodiment of the reconstruction of data by the data storage device of FIGS. 1 and 2 in response to a read instruction;

[0011] FIG. 9 is a simplified block diagram of at least one embodiment of a linked list of pointers generated by the data storage device of FIGS. 1 and 2 for use in searching for data; and

[0012] FIG. 10 is a simplified block diagram of at least one embodiment of a computing device including the data storage device of FIGS. 1 and 2.

DETAILED DESCRIPTION OF THE DRAWINGS

[0013] While the concepts of the present disclosure are susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and will be described herein in detail. It should be understood, however, that there is no intent to limit the concepts of the present disclosure to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives consistent with the present disclosure and the appended claims.

[0014] References in the specification to “one embodiment,” “an embodiment,” “an illustrative embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may or may not necessarily include that particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described. Additionally, it should be appreciated that items included in a list in the form of “at least one A, B, and C” can mean (A); (B); (C); (A and B); (A and C); (B and C); or (A, B, and C). Similarly, items listed in the form of “at least one of A, B, or C” can mean (A); (B); (C); (A and B); (A and C); (B and C); or (A, B, and C).

[0015] The disclosed embodiments may be implemented, in some cases, in hardware, firmware, software, or any combination thereof. The disclosed embodiments may also be implemented as instructions carried by or stored on a transitory or non-transitory machine-readable (e.g., computer-readable) storage medium, which may be read and executed by one or

more processors. A machine-readable storage medium may be embodied as any storage device, mechanism, or other physical structure for storing or transmitting information in a form readable by a machine (e.g., a volatile or non-volatile memory, a media disc, or other media device).

[0016] In the drawings, some structural or method features may be shown in specific arrangements and/or orderings. However, it should be appreciated that such specific arrangements and/or orderings may not be required. Rather, in some embodiments, such features may be arranged in a different manner and/or order than shown in the illustrative figures. Additionally, the inclusion of a structural or method feature in a particular figure is not meant to imply that such feature is required in all embodiments and, in some embodiments, may not be included or may be combined with other features.

[0017] As shown in FIG. 1, an illustrative data storage device 100 for storing and efficiently locating data includes a data storage controller 102 and a memory 116, which illustratively includes non-volatile memory 118 and volatile memory 120. As discussed in more detail below, in use, the data storage controller 102 is configured to store data to the memory 116 using a process that enables the data storage controller 102 to efficiently locate and potentially replace the data at a later time. In the illustrative embodiment, the data storage controller 102 is configured to partition a block of data into multiple sub-blocks and store the sub-blocks in a data table (referred to herein as a “deduplication data table”) such that every sub-block is unique. The data storage controller 102 additionally generates a pointer table that indicates where every sub-block of each block is located in the deduplication data table. For duplicate sub-blocks, multiple pointers in the pointer table point to the same sub-block stored in the deduplication data table. To enable efficient searching for blocks that contain a particular sub-block (e.g., search data), the data storage controller 102 generates and maintains a linked list of pointers within the pointer table wherein the pointers in the linked list all point to the same sub-block in the deduplication data table. The sub-blocks are stored in the deduplication data table at locations determined based on the data in the sub-blocks (e.g., based on a hash of the data). Each location in the deduplication data table has an associated pointer back to the linked list of pointers in the pointer table. Accordingly, once the data storage controller 102 has located a particular sub-block in the deduplication data table, based on a hash of the data, the data storage controller 102 may access and traverse the linked list of pointers to identify every block of data that includes that particular sub-block. Additionally, the data storage controller 102 may quickly replace the data of the sub-block and update the linked list of pointers accordingly. As a result, the data storage controller 102 enables search and replace operations

to occur much more efficiently in terms of bandwidth and processor usage as compared to conventional systems.

[0018] The data storage device 100 may be embodied as any type of device capable of storing data and performing the functions described herein. In the illustrative embodiment, the data storage device 100 is embodied as a solid state drive; however, in other embodiments, the data storage device 100 may be embodied as other storage devices such as a hard disk drive, a memory module device, a cache memory device, and/or other data storage device.

[0019] The data storage controller 102 of the data storage device 100 may be embodied as any type of control device, circuitry, or collection of hardware devices capable of writing, reading, locating, and replacing data in the non-volatile memory 118. In the illustrative embodiment, the data storage controller 102 includes a processor or processing circuitry 104, local memory 106, a host interface 108, lookup logic 110, a buffer 112, and memory control logic 114. Of course, the data storage controller 102 may include additional devices, circuits, and/or components commonly found in a drive controller of a solid state drive in other embodiments.

[0020] The processor 104 may be embodied as any type of processor capable of performing the functions described herein. For example, the processor 104 may be embodied as a single or multi-core processor(s), digital signal processor, Field Programmable Gate Arrays (FPGA), microcontroller, or other processor or processing/controlling circuit. Similarly, the local memory 106 may be embodied as any type of volatile and/or non-volatile memory or data storage capable of performing the functions described herein. In the illustrative embodiment, the local memory 106 stores firmware and/or other instructions executable by the processor 104 to perform the described functions of the data storage controller 102. In some embodiments, the processor 104 and the local memory 106 may form a portion of a System-on-a-Chip (SoC) and be incorporated, along with other components of the data storage controller 102, onto a single integrated circuit chip.

[0021] The host interface 108 may also be embodied as any type of hardware processor, processing circuitry, input/output circuitry, and/or collection of components capable of facilitating communication of the data storage device 100 with a host device or service (e.g., a host application). That is, the host interface 108 embodies or establishes an interface for accessing data stored on the data storage device 100 (e.g., stored in the memory 116). To do so, the host interface 108 may be configured to utilize any suitable communication protocol and/or technology to facilitate communications with the data storage device 100 depending on the type of data storage device. For example, the host interface 108 may be configured to communicate

with a host device or service using Serial Advanced Technology Attachment (SATA), Peripheral Component Interconnect express (PCIe), Serial Attached SCSI (SAS), Universal Serial Bus (USB), and/or other communication protocol and/or technology in some embodiments.

[0022] In the illustrative embodiment, the lookup logic 110 is embodied as dedicated circuitry and/or device configured to facilitate locating sub-blocks of data stored in the memory 116, including calculating hashes of the sub-blocks to determine their locations in the deduplication data table, comparing search data to sub-block data, as well as locating and traversing the linked list of pointers that point to a particular sub-block in the deduplication data table. The lookup logic 110 may be embodied as a co-processor, an application specific integrated circuit (ASIC), or other dedicated circuitry or device. In such embodiments, the lookup logic 110 provides a hardware accelerated implementation of the search-related operations described herein. In other embodiments, at least a portion of the lookup logic 110 may be embodied as firmware or other processor-executable instructions.

[0023] The buffer 112 of the data storage controller 102 is embodied as volatile memory used by the data storage controller 102 to temporarily store data that is being read from or written to the memory 116. The particular size of the buffer 112 may be dependent on the total storage size of the memory 116. The memory control logic 114 is illustrative embodied as hardware circuitry and/or one or more devices configured to control the read/write access to data at particular storage locations of memory 116.

[0024] The non-volatile memory 118 may be embodied as any type of data storage capable of storing data in a persistent manner (even if power is interrupted to non-volatile memory 118). For example, in the illustrative embodiment, the non-volatile memory 118 is embodied as one or more non-volatile memory devices. The non-volatile memory devices of the non-volatile memory 118 are illustratively embodied as byte-addressable, write-in-place non-volatile memory devices. However, in other embodiments, the non-volatile memory 118 may be embodied as any combination of memory devices that use chalcogenide phase change material (e.g., chalcogenide glass), three-dimensional (3D) crosspoint memory, or other types of byte-addressable, write-in-place non-volatile memory, ferroelectric random-access memory (FeTRAM), nanowire-based non-volatile memory, phase change memory (PCM), memory that incorporates memristor technology, Magnetoresistive random-access memory (MRAM) or Spin Transfer Torque (STT)-MRAM.

[0025] The volatile memory 120 may be embodied as any type of data storage capable of storing data while power is supplied to the volatile memory 120. For example, in the

illustrative embodiment, the volatile memory 120 is embodied one or more volatile memory devices, and is periodically referred to hereinafter as volatile memory 120 with the understanding that the volatile memory 120 may be embodied as other types of non-persistent data storage in other embodiments. The volatile memory devices of the volatile memory 120 are illustratively embodied as dynamic random-access memory (DRAM) devices, but may be embodied as other types of volatile memory devices and/or memory technologies capable of storing data while power is supplied to volatile memory 120.

[0026] Referring now to FIG. 2, in use, the data storage device 100 may establish an environment 200. The illustrative environment 200 includes a data management module 220 and an interface module 270. Each of the modules and other components of the environment 200 may be embodied as firmware, software, hardware, or a combination thereof. For example the various modules, logic, and other components of the environment 200 may form a portion of, or otherwise be established by, the data storage controller 102 or other hardware components of the data storage device 100. As such, in some embodiments, any one or more of the modules of the environment 200 may be embodied as a circuit or collection of electrical devices (e.g., a data management circuit 220, an interface circuit 270, etc.). In the illustrative embodiment, the environment 200 includes the data pointer table 202 which includes a set of pointers (also referred to herein as “sub-block pointers”) to sub-blocks in the deduplication data table. The environment 200 additionally includes data 204, which includes the deduplication data table and the sub-blocks organized therein. Further, the environment 200 includes hashes 206 which are calculated from the content of an input set of data, such as a sub-block. As described in more detail herein, the sub-blocks are stored at locations in the data pointer table based at least in part on a hash 206 of the sub-block. This scheme is referred to herein as hash-based indexing. Additionally, the illustrative environment 200 includes reference counts 208 which are values that indicate how many pointers in the data pointer table 202 refer to a particular sub-block in the deduplication data table 204 at any given time. Additionally, the illustrative environment 200 includes pointers 210, which include pointers from the deduplication data table 204 to the data pointer table 202, as well as forward and backward pointers of linked lists that track which pointers within the data pointer table 202 are associated with a particular sub-block stored in the deduplication data table 204. Each of the data pointer table 202, the data 204, the hashes 206, the reference counts 208, and the pointers 210 may be accessed by the various modules and/or sub-modules of the data storage controller 102.

[0027] The data management module 220, which may be embodied as hardware, firmware, software, virtualized hardware, emulated architecture, and/or a combination thereof

as discussed above, is configured to manage the writing, reading, searching, and replacing of data in the memory 116. To do so, the data management module 220 includes a pointer table management module 230, a duplicate analysis module 240, a reference count management module 250, and a search management module 260. In the illustrative embodiment, the pointer table management module 230 is configured to store pointers in the data pointer table 202 to physical addresses within the deduplication data table 204 where each sub-block of a block is stored. In response to a read request, the pointer table management module 230 is configured to traverse the pointers in the data pointer table associated with a given block to be read from the memory 116 to identify the locations (i.e., physical addresses) in the deduplication data table 204 where every sub-block of the data block is stored. Accordingly, the data block may be reassembled in response to the read request.

[0028] The duplicate analysis module 240 is configured to determine whether a given sub-block is a duplicate of (e.g., contains the same data) another sub-block that has already been stored in the deduplication data table 204. To do so, in the illustrative embodiment, the duplicate analysis module 240 includes a hash generator module 242. The hash generator module 242 is configured to analyze an input data set, such as a data sub-block and generate a hash such as one or more values representative of the data in the sub-block. Given that, in the illustrative embodiment, the deduplication data table 204 uses hash-based indexing, the duplicate analysis module 240 is configured to use the hash generated by the hash generator module 242 to locate a sub-block stored in the deduplication data table 204 at a location that is based on the hash, and compare the actual data of the stored sub-block to the input data set (i.e., an input sub-block to be written to the deduplication data table 204), rather than examining every sub-block stored in the entire deduplication data table 204 to identify a duplicate. More specifically, in hash-based indexing, at least a portion of the computed hash acts as an address range (i.e., a pointer) and this address range has several locations (i.e., entries) in which that data sub-block may be stored in the deduplication data table 204. In at least some embodiments, each stored address range (i.e., pointer) may include bits that indicate which entry in the address range contains the exact data of the stored data sub-block. If all entries become full, the last entry may point to an extra location that contains the data of the data sub-block. When determining whether a particular sub-block to be written to the non-volatile memory is a duplicate of a previously-stored sub-block, the duplicate analysis module 240 may be configured to access the data sub-blocks, if any, stored in respective entries in the physical address range indicated by the hash and perform a bit-by-bit or byte-by-byte comparison of the particular sub-block to each previously-stored sub-block.

[0029] The reference count management module 250 is configured to selectively set, increase, or decrease the reference counts 208 associated with the sub-blocks stored in the deduplication data table 204. The reference counts 208 indicate the number of pointers in the data pointer table 202 that presently point to the sub-block. For example, when a data sub-block is initially stored in a data table 218 and the data sub-block is not a duplicate of another data sub-block, the reference count management module 250 is configured to set the reference count for that sub-block to one, which indicates that only one pointer in the data pointer table 202 points to that sub-block. If a later-stored sub-block is a duplicate of the earlier-stored data sub-block, the reference count management module 250 is configured to increment the reference count 208 for that already-stored sub-block in the deduplication data table 204. Similarly, if the number of pointers in the data pointer table 202 decreases, the reference count management module 250 is configured to decrease the reference count for that sub-block correspondingly. A reference count of zero indicates that the corresponding location in the deduplication data table 204 is presently unused and is available for storage of a new data sub-block.

[0030] The search management module 260 is configured to generate a linked list of pointers in the data pointer table that refer to a particular sub-block in the data pointer table. In the illustrative embodiment, the search management module 260 is configured to maintain a separate linked list for each sub-block stored in the deduplication data table 204. Further, the search management module is configured to store a pointer in association with each sub-block in the data deduplication table that refers back to the beginning of the linked list (i.e., the first pointer in the data pointer table 202) for that sub-block. In the illustrative embodiment, the search management module 260 is configured to update the corresponding linked list each time a pointer is added to or removed from the data pointer table 202. Further, the search management module 260 is configured to obtain search data (e.g., an input sub-block), determine whether a matching sub-block exists in the deduplication data table 204, such as by utilizing the duplicate analysis module 240 perform an analysis of whether a duplicate exists (i.e., a matching sub-block) of the input search data, and if a match (i.e., a duplicate) is found, traverse the corresponding linked list of pointers in the data pointer table 202 to identify every data block that includes the matching sub-block. In the illustrative embodiment, the search management module 260 is also configured to identify replacement data associated with a replace request, write the replacement data to the data pointer table as a new sub-block, unless it is a duplicate of an existing sub-block, and update the pointers in the data pointer table and the

linked list to be based on the replacement data (i.e., the new sub-block) instead of the search data (i.e., the original sub-block).

[0031] It should be appreciated that each of the pointer table management module 230, the duplicate analysis module 240, the hash generator module 242, the reference count management module 250, and the search management module 260 may be separately embodied as hardware, firmware, software, virtualized hardware, emulated architecture, and/or a combination thereof. For example, the pointer table management module 230 may be embodied as a hardware component, while the duplicate analysis module 240, the hash generator module 242, the reference count management module 250, and the search management module 260 are embodied as virtualized hardware components or as some other combination of hardware, firmware, software, virtualized hardware, emulated architecture, and/or a combination thereof.

[0032] The interface module 270, which may be embodied as hardware, firmware, software, virtualized hardware, emulated architecture, and/or a combination thereof as discussed above, is configured to handle various instructions, including but not limited to, data storage instructions, data read instructions, data trimming instructions, data search instructions, and data replacement instructions received from a host 280, which may be embodied as an application, service, and/or other device. In some embodiments, the interface module 270 may be configured to handle other instructions as well, including self-monitoring, analysis and reporting technology (“SMART”) instructions, and other instructions defined in the non-volatile memory express (“NVMe”) specification. To handle the various instructions, the interface module 270 is configured to identify a received instruction and any data and/or parameters associated with the instruction, and transmit those items to the data management module 220. In the illustrative embodiment, in response to a search request, the interface module 270 is configured to provide, to the host 280, identifiers of the blocks that contain the matching sub-block. The interface module 270 may additionally be configured to provide acknowledgements and/or success or failure messages to the host 280 in response to requests.

[0033] Referring now to FIG. 3, in use, the data storage controller 102 of the data storage device 100 may execute a method 300 for performing a write operation that enables efficient searching of data. The method begins at block 302 in which the data storage controller 102 determines whether a write instruction has been received from the host 280. If so, the method 300 advances to block 304 in which the data storage controller 102 obtains a sub-block of a data block to be written to the memory 116, such as by reading a 64 byte set of data from a memory address associated with the write request. Otherwise, the method 300 loops back to

block 302 in which the data storage controller 102 continues to wait for a write request from the host 280.

[0034] In block 306, the data storage controller 102 generates a hash of the sub-block. Subsequently, in block 308, the data storage controller 102 searches the deduplication data table 204 for a match of the present sub-block (i.e., the sub-block obtained in block 304). In doing so, in the illustrative embodiment, the data storage controller 102 searches the deduplication data table 204 for a match based on the hash generated in block 306. As described above, the hash is indicative of the underlying data in the sub-block and influences the location where the sub-block is stored in the data deduplication table. Accordingly, in the illustrative embodiment, the data storage controller 102 accesses one or more sub-blocks stored at locations in the deduplication data table 204 that are associated with the hash and compares the actual data of each sub-block to the present sub-block to determine whether a match exists, as indicated in block 312. In block 314, the data storage controller 102 determines whether a match was found. If not, the method 300 advances to block 316, in which the data storage controller 102 writes the present sub-block to a location in the deduplication data table 204. In doing so, as indicated in block 318, the data storage controller 102 may write the present sub-block to a location in the deduplication data table 204 that is based on the hash of the present sub-block, such as by using the hash-based indexing scheme described above.

[0035] In block 320, the data storage controller 102 generates a pointer to the location in the deduplication data table 204 where the present sub-block was written. Subsequently, in block 322, the data storage controller 102 writes the pointer to the data pointer table, such as at a position in the data pointer table corresponding to the present sub-block of the data block being written. In block 324, the data storage controller 102 sets the reference count associated with the sub-block to one, to indicate that one pointer in the data pointer table 202 references the sub-block. In the illustrative embodiment, in block 326, the data storage controller 102 generates a back pointer that points to the entry (i.e., pointer) in the data pointer table that refers to the location in the deduplication data table 204 where the sub-block was written. The data storage controller 102 additionally writes the back pointer in the deduplication data table 204 in association with the sub-block that was written to the deduplication data table 204. In doing so, the data storage controller 102 may store the back pointer as an additional field in the deduplication data table 204 or in a separate data structure in the memory 116. By generating and storing the back pointer to the entry in the data pointer table, the data storage controller 102 can easily locate the beginning of a linked list of pointers in the data pointer table that refer to the sub-block in the deduplication data table 204. While only one pointer in the data pointer

table 202 presently points to the sub-block, the following description of the operations performed by the data storage controller 102 when a match is found in block 314 explain the formation of a linked list of multiple pointers in the data pointer table that refer to the same sub-block.

[0036] Referring back to block 314, if the data storage controller 102 determines that a match has been found, the method 300 advances to block 332 of FIG. 4. In block 332, the data storage controller 102 generates a pointer to the location of the matching sub-block that already exists in the deduplication data table 204. Further, the data storage controller 102 writes the pointer to an entry in the data pointer table 202, as indicated in block 334. In the illustrative embodiment, the data storage controller 102 writes the pointer at a location in the data pointer table 202 associated with the logical position of the present sub-block within the data block that is being written to the memory 116. For example, if the present sub-block is the first sub-block of the data block, the data storage controller 102 writes the pointer in the data pointer table 202 in an entry for the first sub-block of that data block. Further, the data storage controller 102 increments the reference count for the present sub-block in the deduplication data table 204, as indicated in block 336. In block 338, the data storage controller 102 reads an existing back pointer stored in association with the matching sub-block (i.e., the present sub-block) in the data pointer table. As described above, the back pointer provides the location in the data pointer table 202 of the first pointer that refers to the sub-block in the deduplication data table 204. The first pointer may simultaneously represent both the beginning and the end of the linked list if no other pointers in the data pointer table 202 have been added to the linked list. Otherwise, the data storage controller 102 may identify another back pointer that is stored in association with the first pointer in the linked list. The additional back pointer points to a subsequent pointer stored in the data pointer table 202 that also refers to the same sub-block. The data storage controller may traverse multiple such back pointers, each stored in association with a different pointer in the data pointer table 202 until the data storage controller 102 reaches the end of the linked list (i.e., no additional back pointers exist in the linked list). This process of traversing the linked list of back pointers is indicated in block 340.

[0037] In block 342, the data storage controller 102 generates a new back pointer that points to the pointer written to the data pointer table 202 in block 334. Further, as indicated in block 344, the data storage controller 102 extends the linked list by writing the generated back pointer to the end of the linked list in the data pointer table 202. Accordingly, the pointer stored during block 334 becomes the end of the linked list. As indicated in block 346, in the illustrative embodiment, the data storage controller 102 additionally writes a forward pointer

that points to the location of the previous end of the linked list. The data storage controller 102 writes the forward pointer in association with the pointer stored during block 334. In other words, in the illustrative embodiment, the linked list is a doubly-linked list, which enables the data storage controller 102 to easily traverse the linked list in either direction.

[0038] Referring again to FIG. 3, regardless of whether the data storage controller 102 found a match or did not find a match in block 314, the data storage controller 102 determines, in block 330, whether more sub-blocks are to be written to the memory 116. For example, if the present sub-block was only the first sub-block of a larger block, the data storage controller 102 may determine that additional sub-blocks are to be stored. If so, the method 300 loops back to block 304 to obtain the next sub-block to write to the memory 116. Otherwise, the method 300 loops back to block 302 in which the data storage controller 102 awaits another write request.

[0039] Referring now to FIG. 5, the data storage controller 102 of the data storage device 100 may execute a method 500 for performing a search operation. The search operation takes advantage of the data structures generated during the write operation described above to quickly find every pointer in the data pointer table 202 that references a particular sub-block in the deduplication data table 204. The method 500 begins with block 502, in which the data storage controller 102 determines whether a search request has been received, such as from the host 280. After the data storage controller 102 receives a search request, the method 300 advances to block 504 in which the data storage controller 102 obtains search data from the search request. In the illustrative embodiment, the search data is a portion of a file, such as a sub-block. In block 506, the data storage controller 102 generates a hash of the search data. This operation is similar to the process of generating a hash of a sub-block to be written to the memory 116, as described with reference to block 306 of FIG. 3. In block 508, the data storage controller 102 searches the deduplication data table 204 for a match of the present search data. In doing so, in the illustrative embodiment, the data storage controller 102 searches the deduplication data table 204 based on the hash generated in block 506. Further, in the illustrative embodiment, the data storage controller 102 compares the actual data of any sub-blocks that initially appear as a potential match to the search data to determine whether a match exists (e.g., whether the compared data is equal). This process is similar to blocks 308 through 312 of the method 300, described with reference to FIG. 3.

[0040] In block 514, the data storage controller 102 determines whether a match has been found. If not, the data storage controller 102 returns no search results in response to the search request, as indicated in block 516. For example, the data storage controller 102 may

return an indication, such as a numeric code, indicative of no search results. Referring back to block 514, if the data storage controller 102 instead determines that a match was found, the method advances to block 518. In block 518, the data storage controller 102 obtains a back pointer associated with the matching sub-block in the deduplication data table 204. In block 520, the data storage controller 102 traverses the linked list in the data pointer table 202, starting with the pointer in the data pointer table 202 that was pointed to by the back pointer obtained in block 518. In situations in which only one pointer in the data pointer table 202 references the sub-block, traversing the linked list entails locating that one pointer in the data pointer table 202 that was referenced by the back pointer stored in association with the sub-block in the deduplication data table 204. In other situations, in which multiple pointers in the data pointer table 202 reference the sub-block, the data storage controller 102 traverses the linked list by following the back pointers associated with each pointer in the linked list. As described above, each entry (i.e., pointer) in the data pointer table 202 is associated with a sub-block within a larger data block. Each data block has a logical address associated with it. In block 522, the data storage controller 102 returns the logical addresses of the blocks that contain entries (e.g., pointers) in the data pointer table 202 that are within the linked list that was traversed in block 520. In the illustrative embodiment, these logical addresses are the search results.

[0041] Referring now to FIG. 6, in use, the data storage controller 102 of the data storage device 100 may execute a method 600 for performing a replace operation. The method 600 begins with block 602 in which the data storage controller 102 determines whether a replace request has been received, such as from the host 280. After a replace request has been received, the method 600 advances to block 604 in which the data storage controller 102 obtains search data and replacement data. In block 606, the data storage controller 102 performs a search based on the search data. In doing so, in the illustrative embodiment, the data storage controller 102 performs the method 500, described with reference to FIG. 5. In block 608, the data storage controller 102 determines whether one or more search results were found in the search. If not, the method 600 loops back to block 602, in which the data storage controller 102 awaits another replace request. However, if the data storage controller 102 determines at block 608 that at least one search result was found, the data storage controller 102 may present one of the search results to the user, such as through a graphical user interface (not shown) to enable the user to review the surrounding content and determine whether to replace the search data with the replacement data for that search result.

[0042] In block 612, the data storage controller 102 determines whether to proceed with replacing the search data with the replacement data. In the illustrative embodiment, the data storage controller 102 receives a response from the user, such as through the graphical user interface, of whether to proceed with replacing the search data with the replacement data in the search result. If the data storage controller 102 determines not to proceed with replacing the search data with the replacement data, the method 600 loops back to block 608 in which the data storage controller 102 determines whether another search result was found. Otherwise, the method 600 advances to block 614, in which the data storage controller 102 determines a hash value based on the replacement data. In block 616, the data storage controller 102 writes the replacement data to the deduplication data table 204 at a location that is based on the hash that was determined in block 614. As should be understood from the foregoing description, the search data and the replacement data are treated as sub-blocks in the above process. Further, given that the content of a sub-block affects its hash, which in turn affects the location where the sub-block is stored in the deduplication data table 204, the replacement data written to the deduplication data table 204 is likely stored in a different location in the deduplication data table 204 than the original search data. Accordingly, in block 618, the data storage controller 102 updates the corresponding pointers in the data pointer table to refer to the location of the replacement data in the deduplication data table 204, rather than the location of the search data and updates the linked list to reflect that the present entry in the data pointer table no longer points to the search data in the deduplication data table 204. Afterwards, the method 600 loops back to block 608 to determine whether an additional search result was found.

[0043] Referring now to FIG. 7, a block diagram 700 illustrates an embodiment of the writing of data by the data storage device 100. In the illustrative embodiment, the data storage controller 102 obtains a data block 710 to be written to the memory 116. The illustrative data block 710 includes a plurality of data sub-blocks. In the illustrative embodiment, the data block 710 includes 64 sub-blocks, each of which is 64 bytes. In other embodiments, the number of sub-blocks in a data block and the number of bytes of data in a sub-block may differ. The data pointer table 202, located in the memory 116, includes a set of entries (i.e., pointers) 722, 724, 726, each of which corresponds to a respective one of the sub-blocks of the data block 710. The data pointer table 202 additionally includes other sets of entries (i.e., pointers) 730, 740, 750 corresponding to other data blocks. Each entry (i.e., pointer) in the data pointer table 202 references a physical address 760 (i.e., a location) in the deduplication data table 204 where a sub-block is stored. Additionally, the data storage controller 102 sets and/or increments the reference count 208 associated with each sub-block stored in the deduplication data table 204.

[0044] Referring now to FIG. 8, a block diagram 800 illustrates an embodiment of the reconstruction of data by the data storage device 100, such as in response to a read request from the host 280. In order to read a data block (e.g., data block 710), the data storage controller 102 accesses the set 720 of pointers in the data pointer table 202 associated with the data block to be read. The data storage controller 102 sequentially reads each pointer in the set 720, reads the sub-block stored at the location referenced by the pointers, and combines the read sub-blocks to form the data block 710.

[0045] Referring now to FIG. 9, a block diagram 900 illustrates an embodiment of a linked list of pointers that may be generated by the data storage controller 102 when writing data and traversed by the data storage controller 102 when searching for data. A sub-block 910 in the deduplication data table 204 is associated with a back pointer 912 that provides the address of a pointer 920 in the data pointer table 202. The pointer 920 references the sub-block 910. A back pointer 922 stored in association with the pointer 920 references another pointer 930 in the data pointer table. The pointer 930 references the same sub-block 910 as the pointer 920. A back pointer 932 stored in association with the pointer 930 references another pointer 940 in the data pointer table 202. The pointer 940 also references the sub-block 910. Accordingly, in the data pointer table 202, three pointers 920, 930, and 940 reference the sub-block 910 in the deduplication data table 204. In the illustrative embodiment, a null back pointer 942 (i.e., a pointer having a value of zero or null) is stored in association with the pointer 940 to indicate the end of the linked list. Accordingly, in the illustrative embodiment, in response to a search request, the data storage controller 102 may locate the sub-block 910 in the deduplication data table 204, follow the back pointer 912 to the pointer 920, then traverse the back pointers 922 and 932 to identify the pointers in the data pointer table, and their associated data blocks, that include the search data (i.e., the sub-block 910). Additionally, in the illustrative embodiment, the linked list is a double-linked list, such that each of the pointers 940, 930, and 920 may be traversed in reverse order by following forward pointers 944 and 934. This is advantageous in enabling the data storage controller 102 to more quickly traverse the list in either direction and make changes, such as deletions and/or substitutions of pointers in the linked list.

[0046] Referring now to FIG. 10, in some embodiments, the data storage device 100 may be incorporated in, or form a portion of, a computing device 1000. The computing device 1000 may be embodied as any type of computing device in which the data storage device 100 may be used. For example, the computing device 1000 may be embodied as a smart phone, a tablet computer, a notebook, a laptop computer, a netbook, an Ultrabook™, a wearable

computing device, a pair of smart glasses, a head-mounted computing device, a cellular phone, a desktop computer, a smart device, a personal digital assistant, a mobile Internet device, a server, a data storage device, and/or any other computing/communication device. As shown in FIG. 10, the illustrative computing device 1000 includes a processor 1010, an input/output (“I/O”) subsystem 1012, and a main memory 1014. Of course, the computing device 1000 may include other or additional components, such as those commonly found in a typical computing device (e.g., various input/output devices and/or other components), in other embodiments. Additionally, in some embodiments, one or more of the illustrative components may be incorporated in, or otherwise form a portion of, another component. For example, the memory 1014, or portions thereof, may be incorporated in the processor 1010 in some embodiments.

[0047] The processor 1010 may be embodied as any type of processor capable of performing the functions described herein. For example, the processor 1010 may be embodied as a single or multi-core processor(s), digital signal processor, microcontroller, or other processor or processing/controlling circuit. Similarly, the memory 1014 may be embodied as any type of volatile or non-volatile memory or data storage capable of performing the functions described herein. In operation, the memory 1014 may store various data and software used during operation of the computing device 1000 such as operating systems, applications, programs, libraries, and drivers. The memory 1014 is communicatively coupled to the processor 1010 via the I/O subsystem 1012, which may be embodied as circuitry and/or components to facilitate input/output operations with the processor 1010, the memory 1014, and other components of the computing device 1000. For example, the I/O subsystem 1012 may be embodied as, or otherwise include, memory controller hubs, input/output control hubs, firmware devices, communication links (i.e., point-to-point links, bus links, wires, cables, light guides, printed circuit board traces, etc.) and/or other components and subsystems to facilitate the input/output operations.

[0048] As shown in FIG. 10, the data storage device 100 may be incorporated in, or form a portion of, one or more other components of the computing device 1000. For example, the data storage device 100 may be embodied as, or otherwise be included in, the main memory 1014. Additionally or alternatively, the data storage device 100 may be embodied as, or otherwise included in, a solid state drive 1020 of the computing device 1000. Further, in some embodiments, the data storage device 100 may be embodied as, or otherwise included in, a hard disk drive 1030 of the computing device 1000. Of course, in other embodiments, the data storage device 100 may be included in or form a portion of other components of the computing device 1000.

[0049] Reference to memory devices can apply to different memory types, and in particular, any memory that has a bank group architecture. Memory devices generally refer to volatile memory technologies. Volatile memory is memory whose state (and therefore the data stored on it) is indeterminate if power is interrupted to the device. Nonvolatile memory refers to memory whose state is determinate even if power is interrupted to the device. Dynamic volatile memory requires refreshing the data stored in the device to maintain state. One example of dynamic volatile memory includes DRAM (dynamic random access memory), or some variant such as synchronous DRAM (SDRAM). A memory subsystem as described herein may be compatible with a number of memory technologies, such as DDR4 (DDR version 4, initial specification published in September 2012 by JEDEC), DDR4E (in development by JEDEC), LPDDR4 (LOW POWER DOUBLE DATA RATE (LPDDR) version 4, JESD209-4, originally published by JEDEC in August 2014), WIO2 (Wide I/O 2 (WideIO2), JESD229-2, originally published by JEDEC in August 2014), HBM (HIGH BANDWIDTH MEMORY DRAM, JESD235, originally published by JEDEC in October 2013), DDR5 (DDR version 5, currently in discussion by JEDEC), LPDDR5 (currently in discussion by JEDEC), HBM2 (HBM version 2), currently in discussion by JEDEC), and/or others, and technologies based on derivatives or extensions of such specifications.

[0050] In addition to, or alternatively to, volatile memory, in one embodiment, reference to memory devices can refer to a nonvolatile memory device whose state is determinate even if power is interrupted to the device.

EXAMPLES

[0051] Illustrative examples of the technologies disclosed herein are provided below. An embodiment of the technologies may include any one or more, and any combination of, the examples described below.

[0052] Example 1 includes an apparatus comprising a memory to store (i) sub-blocks of data in a data table and (ii) a pointer table of locations of the sub-blocks in the data table; and a controller to manage the storage and lookup of data in the memory, wherein the controller is to store a sub-block pointer in the pointer table to a location of a sub-block in the data table; and store a second pointer that references an entry where the sub-block pointer is stored in the pointer table.

[0053] Example 2 includes the subject matter of Example 1, and wherein the sub-block is a first sub-block, the entry is a first entry, and the controller is further to determine whether a second sub-block to be written to the data table matches the first sub-block; and store, in

response to a determination that the second sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a second entry and (ii) a third pointer in association with the first entry, wherein the third pointer is to reference the second entry.

[0054] Example 3 includes the subject matter of any of Examples 1 and 2, and wherein the controller is further to store, in response to the determination that the second sub-block matches the first sub-block, a fourth pointer in association with the second entry, wherein the fourth pointer references the first entry.

[0055] Example 4 includes the subject matter of any of Examples 1-3, and wherein the controller is further to obtain a third sub-block to be written to the data table; determine whether the third sub-block matches the first sub-block; and store, in response to a determination that the third sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a third entry and (ii) a fifth pointer in association with the second entry, wherein the fifth pointer is to reference the third entry, and (iii) a sixth pointer in association with the third entry, wherein the sixth pointer is to reference the second entry.

[0056] Example 5 includes the subject matter of any of Examples 1-4, and wherein the controller is further to generate a hash of the obtained sub-block; and search the memory for a match of the obtained sub-block, and wherein to store the sub-block at a location in the data table comprises to store the sub-block at a location in the data table in response to a determination that the memory does not include a match for the sub-block.

[0057] Example 6 includes the subject matter of any of Examples 1-5, and wherein to search the memory for a match of the obtained sub-block comprises to identify one or more locations in the data table associated with the hash; and compare data stored at each of the one or more locations to the sub-block to determine whether the data stored at each location matches the data of the sub-block.

[0058] Example 7 includes the subject matter of any of Examples 1-6, and wherein the controller is further to obtain search data to locate in the memory; search the data table for a sub-block that matches the search data; determine, in response to an identification of a sub-block that matches the search data, a third pointer to an entry in the pointer table, wherein a sub-block pointer to the sub-block is stored; determine a logical address of a block associated with the sub-block pointer; and add the logical address to a search result set.

[0059] Example 8 includes the subject matter of any of Examples 1-7, and wherein the entry is a first entry and the controller is further to identify a fourth pointer stored in association with the first entry in the pointer table, wherein the fourth pointer is to reference a second entry in the pointer table wherein a second sub-block pointer to the sub-block is stored; determine a

second logical address of a second block, wherein the second block is associated with the second sub-block pointer; and add the second logical address to the search result set.

[0060] Example 9 includes the subject matter of any of Examples 1-8, and wherein the controller is further to obtain search data; identify a sub-block in the data table, wherein the sub-block matches the search data; identify a sub-block pointer in the pointer table, wherein the sub-block pointer is to point to a location of the sub-block in the data table; and traverse a linked list of additional sub-block pointers in the pointer table, wherein each sub-block pointer in the linked list is to point to the sub-block.

[0061] Example 10 includes the subject matter of any of Examples 1-9, and wherein the linked list is a doubly-linked linked list of pointers.

[0062] Example 11 includes the subject matter of any of Examples 1-10, and wherein the controller is further to obtain replacement data; generate a hash as a function of the replacement data; and write the replacement data to the data table at a replacement location determined as a function of the hash.

[0063] Example 12 includes the subject matter of any of Examples 1-11, and wherein the controller is further to traverse the linked list and update the sub-block pointers to point to the replacement location.

[0064] Example 13 includes the subject matter of any of Examples 1-12, and wherein the controller is further to set a reference count associated with a location in the data table in response to storage of the sub-block.

[0065] Example 14 includes the subject matter of any of Examples 1-13, and wherein to obtain the sub-block of the block comprises to obtain a 64 byte sub-block of a 4 kilobyte block.

[0066] Example 15 includes the subject matter of any of Examples 1-14, and wherein to store a sub-block pointer in the pointer table comprises to store a 64 bit sub-block pointer.

[0067] Example 16 includes the subject matter of any of Examples 1-15, and wherein the memory includes a plurality of non-volatile memory devices.

[0068] Example 17 includes the subject matter of any of Examples 1-16, and wherein the memory includes a plurality of byte-addressable write-in-place memory devices.

[0069] Example 18 includes the subject matter of any of Examples 1-17, and furthering including one or more of at least one processor communicatively coupled to the memory, a network interface communicatively coupled to a processor, a display communicatively coupled to a processor, or a battery coupled to the apparatus.

[0070] Example 19 includes a method comprising storing, by a controller of an apparatus, a sub-block pointer in a pointer table to a location of the sub-block in a data table;

and storing, by the controller, a second pointer that references an entry where the sub-block pointer is stored in the pointer table.

[0071] Example 20 includes the subject matter of Example 19, and wherein the sub-block is a first sub-block and the entry is a first entry, the method furthering including determining, by the controller, whether a second sub-block to be written to the data table matches the first sub-block; and storing, by the controller and in response to a determination that the second sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a second entry and (ii) a third pointer in association with the first entry, wherein the third pointer is to reference the second entry.

[0072] Example 21 includes the subject matter of any of Examples 19 and 20, and furthering including storing, by the controller and in response to the determination that the second sub-block matches the first sub-block, a fourth pointer in association with the second entry, wherein the fourth pointer references the first entry.

[0073] Example 22 includes the subject matter of any of Examples 19-21, and furthering including obtaining, by the controller, a third sub-block to be written to the data table; determining, by the controller, whether the third sub-block matches the first sub-block; and storing, by the controller and in response to a determination that the third sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a third entry and (ii) a fifth pointer in association with the second entry, wherein the fifth pointer is to reference the third entry, and (iii) a sixth pointer in association with the third entry, wherein the sixth pointer is to reference the second entry.

[0074] Example 23 includes the subject matter of any of Examples 19-22, and furthering including generating, by the controller, a hash of the obtained sub-block; and searching, by the controller, the memory for a match of the obtained sub-block, wherein storing the sub-block at a location in the data table comprises storing the sub-block at a location in the data table in response to a determination that the memory does not include a match for the sub-block.

[0075] Example 24 includes the subject matter of any of Examples 19-23, and wherein searching the memory for a match of the obtained sub-block comprises identifying one or more locations in the data table associated with the hash; and comparing data stored at each of the one or more locations to the sub-block to determine whether the data stored at each location matches the data of the sub-block.

[0076] Example 25 includes the subject matter of any of Examples 19-24, and furthering including obtaining, by the controller, search data to locate in the memory; searching,

by the controller, the data table for a sub-block that matches the search data; determining, by the controller and in response to an identification of a sub-block that matches the search data, a third pointer to an entry in the pointer table, wherein a sub-block pointer to the sub-block is stored; determining, by the controller, a logical address of a block associated with the sub-block pointer; and adding, by the controller, the logical address to a search result set.

[0077] Example 26 includes the subject matter of any of Examples 19-25, and wherein the entry is a first entry, the method furthering including identifying, by the controller, a fourth pointer stored in association with the first entry in the pointer table, wherein the fourth pointer is to reference a second entry in the pointer table wherein a second sub-block pointer to the sub-block is stored; determining, by the controller, a second logical address of a second block, wherein the second block is associated with the second sub-block pointer; and adding, by the controller, the second logical address to the search result set.

[0078] Example 27 includes the subject matter of any of Examples 19-26, and furthering including obtaining, by the controller, search data; identifying, by the controller, a sub-block in the data table, wherein the sub-block matches the search data; identifying, by the controller, a sub-block pointer in the pointer table, wherein the sub-block pointer is to point to a location of the sub-block in the data table; and traversing, by the controller, a linked list of additional sub-block pointers in the pointer table, wherein each sub-block pointer in the linked list is to point to the sub-block.

[0079] Example 28 includes the subject matter of any of Examples 19-27, and wherein the linked list is a doubly-linked linked list of pointers.

[0080] Example 29 includes the subject matter of any of Examples 19-28, and furthering including obtaining, by the controller, replacement data; generating, by the controller, a hash as a function of the replacement data; and writing, by the controller, the replacement data to the data table at a replacement location determined as a function of the hash.

[0081] Example 30 includes the subject matter of any of Examples 19-29, and furthering including traversing, by the controller, the linked list; and updating, by the controller, the sub-block pointers to point to the replacement location.

[0082] Example 31 includes the subject matter of any of Examples 19-30, and furthering including setting, by the controller, a reference count associated with a location in the data table in response to storage of the sub-block.

[0083] Example 32 includes the subject matter of any of Examples 19-31, and wherein obtaining the sub-block of the block comprises obtaining a 64 byte sub-block of a 4 kilobyte block.

[0084] Example 33 includes the subject matter of any of Examples 19-32, and wherein storing a sub-block pointer in the pointer table comprises storing a 64 bit sub-block pointer.

[0085] Example 34 includes one or more machine-readable storage media comprising a plurality of instructions stored thereon that, when executed, cause an apparatus to perform the method of any of Examples 19-33.

[0086] Example 35 includes the subject matter of Example 34, and an apparatus comprising means for storing a sub-block pointer in a pointer table to a location of the sub-block in a data table; and means for storing a second pointer that references an entry where the sub-block pointer is stored in the pointer table.

[0087] Example 36 includes the subject matter of any of Examples 34 and 35, and wherein the sub-block is a first sub-block and the entry is a first entry, the apparatus further including means for determining whether a second sub-block to be written to the data table matches the first sub-block; and means for storing, in response to a determination that the second sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a second entry and (ii) a third pointer in association with the first entry, wherein the third pointer is to reference the second entry.

[0088] Example 37 includes the subject matter of any of Examples 34-36, and further including means for storing, in response to the determination that the second sub-block matches the first sub-block, a fourth pointer in association with the second entry, wherein the fourth pointer references the first entry.

[0089] Example 38 includes the subject matter of any of Examples 34-37, and further including means for obtaining a third sub-block to be written to the data table; means for determining whether the third sub-block matches the first sub-block; and means for storing, in response to a determination that the third sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a third entry and (ii) a fifth pointer in association with the second entry, wherein the fifth pointer is to reference the third entry, and (iii) a sixth pointer in association with the third entry, wherein the sixth pointer is to reference the second entry.

[0090] Example 39 includes the subject matter of any of Examples 34-38, and further including means for generating a hash of the obtained sub-block; and means for searching the memory for a match of the obtained sub-block, wherein the means for storing the sub-block at a location in the data table comprises means for storing the sub-block at a location

in the data table in response to a determination that the memory does not include a match for the sub-block.

[0091] Example 40 includes the subject matter of any of Examples 34-39, and wherein the means for searching the memory for a match of the obtained sub-block comprises means for identifying one or more locations in the data table associated with the hash; and means for comparing data stored at each of the one or more locations to the sub-block to determine whether the data stored at each location matches the data of the sub-block.

[0092] Example 41 includes the subject matter of any of Examples 34-40, and furthering including means for obtaining search data to locate in the memory; means for searching the data table for a sub-block that matches the search data; means for determining, in response to an identification of a sub-block that matches the search data, a third pointer to an entry in the pointer table, wherein a sub-block pointer to the sub-block is stored; means for determining a logical address of a block associated with the sub-block pointer; and means for adding the logical address to a search result set.

[0093] Example 42 includes the subject matter of any of Examples 34-41, and wherein the entry is a first entry, the apparatus furthering including means for identifying a fourth pointer stored in association with the first entry in the pointer table, wherein the fourth pointer is to reference a second entry in the pointer table wherein a second sub-block pointer to the sub-block is stored; means for determining a second logical address of a second block, wherein the second block is associated with the second sub-block pointer; and means for adding the second logical address to the search result set.

[0094] Example 43 includes the subject matter of any of Examples 34-42, and furthering including means for obtaining search data; means for identifying a sub-block in the data table, wherein the sub-block matches the search data; means for identifying a sub-block pointer in the pointer table, wherein the sub-block pointer is to point to a location of the sub-block in the data table; and means for traversing a linked list of additional sub-block pointers in the pointer table, wherein each sub-block pointer in the linked list is to point to the sub-block.

[0095] Example 44 includes the subject matter of any of Examples 34-43, and wherein the linked list is a doubly-linked linked list of pointers.

[0096] Example 45 includes the subject matter of any of Examples 34-44, and furthering including means for obtaining replacement data; means for generating a hash as a function of the replacement data; and means for writing the replacement data to the data table at a replacement location determined as a function of the hash.

[0097] Example 46 includes the subject matter of any of Examples 34-45, and furthering including means for traversing the linked list; and means for updating the sub-block pointers to point to the replacement location.

[0098] Example 47 includes the subject matter of any of Examples 34-46, and furthering including means for setting a reference count associated with a location in the data table in response to storage of the sub-block.

[0099] Example 48 includes the subject matter of any of Examples 34-47, and wherein the means for obtaining the sub-block of the block comprises means for obtaining a 64 byte sub-block of a 4 kilobyte block.

[00100] Example 49 includes the subject matter of any of Examples 34-48, and wherein the means for storing a sub-block pointer in the pointer table comprises means for storing a 64 bit sub-block pointer.

WHAT IS CLAIMED IS:

1. An apparatus comprising:
 - a memory to store (i) sub-blocks of data in a data table and (ii) a pointer table of locations of the sub-blocks in the data table; and
 - a controller to manage the storage and lookup of data in the memory, wherein the controller is to:
 - store a sub-block pointer in the pointer table to a location of a sub-block in the data table; and
 - store a second pointer that references an entry where the sub-block pointer is stored in the pointer table.
2. The apparatus of claim 1, wherein the sub-block is a first sub-block, the entry is a first entry, and the controller is further to:
 - determine whether a second sub-block to be written to the data table matches the first sub-block; and
 - store, in response to a determination that the second sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a second entry and (ii) a third pointer in association with the first entry, wherein the third pointer is to reference the second entry.
3. The apparatus of claim 2, wherein the controller is further to store, in response to the determination that the second sub-block matches the first sub-block, a fourth pointer in association with the second entry, wherein the fourth pointer references the first entry.
4. The apparatus of claim 3, wherein the controller is further to:
 - obtain a third sub-block to be written to the data table;
 - determine whether the third sub-block matches the first sub-block; and
 - store, in response to a determination that the third sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a third entry and (ii) a fifth pointer in association with the second entry, wherein the fifth pointer is to reference the third entry, and (iii) a sixth pointer in association with the third entry, wherein the sixth pointer is to reference the second entry.
5. The apparatus of claim 1, wherein the controller is further to:
 - generate a hash of the obtained sub-block; and

search the memory for a match of the obtained sub-block, and
wherein to store the sub-block at a location in the data table comprises to store the sub-block at a location in the data table in response to a determination that the memory does not include a match for the sub-block.

6. The apparatus of claim 5, wherein to search the memory for a match of the obtained sub-block comprises to:

identify one or more locations in the data table associated with the hash; and
compare data stored at each of the one or more locations to the sub-block to determine whether the data stored at each location matches the data of the sub-block.

7. The apparatus of claim 1, wherein the controller is further to:
obtain search data to locate in the memory;
search the data table for a sub-block that matches the search data;
determine, in response to an identification of a sub-block that matches the search data, a third pointer to an entry in the pointer table, wherein a sub-block pointer to the sub-block is stored;

determine a logical address of a block associated with the sub-block pointer; and
add the logical address to a search result set.

8. The apparatus of claim 7, wherein the entry is a first entry and the controller is further to:

identify a fourth pointer stored in association with the first entry in the pointer table, wherein the fourth pointer is to reference a second entry in the pointer table wherein a second sub-block pointer to the sub-block is stored;

determine a second logical address of a second block, wherein the second block is associated with the second sub-block pointer; and

add the second logical address to the search result set.

9. The apparatus of claim 1, wherein the controller is further to:
obtain search data;
identify a sub-block in the data table, wherein the sub-block matches the search data;

identify a sub-block pointer in the pointer table, wherein the sub-block pointer is to point to a location of the sub-block in the data table; and

traverse a linked list of additional sub-block pointers in the pointer table, wherein each sub-block pointer in the linked list is to point to the sub-block.

10. The apparatus of claim 9, wherein the linked list is a doubly-linked linked list of pointers.

11. The apparatus of claim 9, wherein the controller is further to:
obtain replacement data;
generate a hash as a function of the replacement data; and
write the replacement data to the data table at a replacement location determined as a function of the hash.

12. The apparatus of claim 11, wherein the controller is further to traverse the linked list and update the sub-block pointers to point to the replacement location.

13. A method comprising:
storing, by a controller of an apparatus, a sub-block pointer in a pointer table to a location of the sub-block in a data table; and
storing, by the controller, a second pointer that references an entry where the sub-block pointer is stored in the pointer table.

14. The method of claim 13, wherein the sub-block is a first sub-block and the entry is a first entry, the method further comprising:

determining, by the controller, whether a second sub-block to be written to the data table matches the first sub-block; and

storing, by the controller and in response to a determination that the second sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a second entry and (ii) a third pointer in association with the first entry, wherein the third pointer is to reference the second entry.

15. The method of claim 14, further comprising storing, by the controller and in response to the determination that the second sub-block matches the first sub-block, a fourth

pointer in association with the second entry, wherein the fourth pointer references the first entry.

16. The method of claim 15, further comprising:
obtaining, by the controller, a third sub-block to be written to the data table;
determining, by the controller, whether the third sub-block matches the first sub-block; and
storing, by the controller and in response to a determination that the third sub-block matches the first sub-block, (i) the sub-block pointer in the pointer table in a third entry and (ii) a fifth pointer in association with the second entry, wherein the fifth pointer is to reference the third entry, and (iii) a sixth pointer in association with the third entry, wherein the sixth pointer is to reference the second entry.

17. The method of claim 13, further comprising:
generating, by the controller, a hash of the obtained sub-block; and
searching, by the controller, the memory for a match of the obtained sub-block, wherein storing the sub-block at a location in the data table comprises storing the sub-block at a location in the data table in response to a determination that the memory does not include a match for the sub-block.

18. The method of claim 17, wherein searching the memory for a match of the obtained sub-block comprises:
identifying one or more locations in the data table associated with the hash; and
comparing data stored at each of the one or more locations to the sub-block to determine whether the data stored at each location matches the data of the sub-block.

19. The method of claim 13, further comprising:
obtaining, by the controller, search data to locate in the memory;
searching, by the controller, the data table for a sub-block that matches the search data;
determining, by the controller and in response to an identification of a sub-block that matches the search data, a third pointer to an entry in the pointer table, wherein a sub-block pointer to the sub-block is stored;

determining, by the controller, a logical address of a block associated with the sub-block pointer; and

adding, by the controller, the logical address to a search result set.

20. The method of claim 19, wherein the entry is a first entry, the method further comprising:

identifying, by the controller, a fourth pointer stored in association with the first entry in the pointer table, wherein the fourth pointer is to reference a second entry in the pointer table wherein a second sub-block pointer to the sub-block is stored;

determining, by the controller, a second logical address of a second block, wherein the second block is associated with the second sub-block pointer; and

adding, by the controller, the second logical address to the search result set.

21. The method of claim 13, further comprising:

obtaining, by the controller, search data;

identifying, by the controller, a sub-block in the data table, wherein the sub-block matches the search data;

identifying, by the controller, a sub-block pointer in the pointer table, wherein the sub-block pointer is to point to a location of the sub-block in the data table; and

traversing, by the controller, a linked list of additional sub-block pointers in the pointer table, wherein each sub-block pointer in the linked list is to point to the sub-block.

22. The method of claim 21, wherein the linked list is a doubly-linked linked list of pointers.

23. The method of claim 21, further comprising:

obtaining, by the controller, replacement data;

generating, by the controller, a hash as a function of the replacement data; and

writing, by the controller, the replacement data to the data table at a replacement location determined as a function of the hash.

24. One or more machine-readable storage media comprising a plurality of instructions stored thereon that, when executed, cause an apparatus to perform the method of any of claims 13-23.

25. An apparatus comprising means for performing the method of any of claims 13-23.




FIG. 1





FIG. 2

3/10






FIG. 4




FIG. 5

6/10




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

A. CLASSIFICATION OF SUBJECT MATTER**G06F 12/02(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 12/02; G06F 12/06; G06F 12/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: pointer table, data table, reference, entry, match, traverse, linked list

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014-0237163 A1 (LSI CORPORATION) 21 August 2014 See paragraphs [0015], [0021]-[0024], [0030]-[0033]; claims 1, 4, 8; and figures 1-6.	1,5-6,13,17-18 ,24-25
Y		2-4,7-12,14-16 ,19-23
Y	US 7380059 B2 (DAVID ALAN BURTON) 27 May 2008 See column 7, line 36 - column 8, line 30; and figures 5-7.	2-4,14-16
Y	US 2006-0168392 A1 (HO-SUNG HWANG) 27 July 2006 See paragraph [0041]; and figures 3, 5.	7-8,19-20
Y	US 6912641 B2 (TIANLONG CHEN et al.) 28 June 2005 See column 7, lines 7-12; and figure 3.	9-12,21-23
A	US 8725982 B2 (AZUL SYSTEMS, INC.) 13 May 2014 See column 3, line 41 - column 4, line 54; and figures 3-4.	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

24 August 2017 (24.08.2017)

Date of mailing of the international search report

25 August 2017 (25.08.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/035201

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0237163 A1	21/08/2014	US 9189409 B2	17/11/2015
US 7380059 B2	27/05/2008	US 2004-0230737 A1	18/11/2004
		US 2006-0265568 A1	23/11/2006
		US 2007-0260846 A1	08/11/2007
		US 7124243 B2	17/10/2006
US 2006-0168392 A1	27/07/2006	CN 100419714 C	17/09/2008
		CN 1466060 A	07/01/2004
		JP 2004-013895 A	15/01/2004
		KR 10-0453053 B1	15/10/2004
		KR 10-2003-0095438 A	24/12/2003
		US 2003-0229753 A1	11/12/2003
US 6912641 B2	28/06/2005	EP 1629386 A1	01/03/2006
		EP 1629386 A4	13/12/2006
		US 2004-0221134 A1	04/11/2004
		US 2005-0283584 A1	22/12/2005
		US 7185165 B2	27/02/2007
		WO 2004-099996 A1	18/11/2004
US 8725982 B2	13/05/2014	US 2010-0180090 A1	15/07/2010
		US 2013-0311741 A1	21/11/2013
		US 7647458 B1	12/01/2010
		US 8046544 B2	25/10/2011
		US 8452938 B1	28/05/2013