




- (51) International Patent Classification:
G06F 12/02 (2006.01) *G06F 3/06* (2006.01)
- (21) International Application Number:
PCT/US2018/020152
- (22) International Filing Date:
28 February 2018 (28.02.2018)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/483,818 10 April 2017 (10.04.2017) US
15/726,313 05 October 2017 (05.10.2017) US
- (71) Applicant: WESTERN DIGITAL TECHNOLOGIES, INC. [US/US]; 5601 Great Oaks Parkway, San Jose, California 95119 (US).
- (72) Inventors: SONG, Seung-Hwan; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, California 95119 (US). DE, Arup; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, California 95119 (US). MEHRA, Pankaj; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, California 95119 (US). O'KRAFKA, Brian W.; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, California 95119 (US).
- (74) Agent: MAGEN, Burt; Vierra Magen Marcus LLP, 101 Mission Street, Suite 400, San Francisco, California 94105 (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,

(54) Title: HYBRID LOGICAL TO PHYSICAL ADDRESS TRANSLATION FOR NON-VOLATILE STORAGE DEVICES WITH INTEGRATED COMPUTE MODULE

Figure 23



(57) Abstract: A host compiles code to perform a set of one or more database operations on target and embeds an indication of whether the target data is randomly accessed data or sequentially accessed data. The compiled code is transmitted to the compute engine inside a memory system that maintains a first portion of memory for storing sequentially accessed data and a second portion of memory for storing randomly accessed data. The memory system (e.g. SSD) maintains reduced size L2P tables in volatile working memory by maintaining coarse L2P tables in the working memory for use with sequentially accessed data and maintaining fine L2P tables in the working memory for use with randomly accessed data. The compute engine uses the compiled code to perform the set of one or more database operations on the target data using the working memory.

KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

— with international search report (Art. 21(3))

HYBRID LOGICAL TO PHYSICAL ADDRESS TRANSLATION FOR NON-VOLATILE STORAGE DEVICES WITH INTEGRATED COMPUTE MODULE

BACKGROUND

[0001] Non-volatile semiconductor memory is used in solid state drives (SSD). As Internet-scale services continue to grow, real time data processing and data analytics by ad-hoc queries on large volumes of data is emerging as a critical application. Additionally, as memory density continues to scale, SSD capacities continue to scale exponentially. Current enterprise systems are ill-equipped to manage these trends as they rely on moving huge volumes of data out of the SSD and into a host system's main memory for processing. These solutions rely on storing data at one location (i.e. a storage device like a SSD) and move data to a different location (typically DRAM of the host) for computation. While this method works for some applications with limited data sizes, applications with large scale data cannot use this method because of the time wasted on transferring data and the prohibitively high cost and power consumption of including large scale (e.g. petabyte) DRAM capacity in such systems.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Like-numbered elements refer to common components in the different figures.

[0003] Figure 1 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0004] Figure 2 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0005] Figure 3 is a block diagram of one embodiment of a Front End Processor Circuit with a compute engine. The Front End Processor Circuit is part of a Controller.

[0006] Figure 4 is a block diagram of one embodiment of a Back End Processor Circuit. In some embodiments, the Back End Processor Circuit is part of a Controller.

[0007] Figure 5 is a block diagram of one embodiment of a memory package.

[0008] Figure 6 is a block diagram of one embodiment of a memory die.

[0009] Figure 7 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0010] Figure 8 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0011] Figure 9 is a block diagram of one embodiment of a Front End Processor Circuit without a compute engine. In some embodiments, the Front End Processor Circuit is part of a Controller.

[0012] Figure 10 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0013] Figure 11 is a block diagram of one embodiment of a Back End Processor Circuit with a compute engine. In some embodiments, the Back End Processor Circuit is part of a Controller.

[0014] Figure 12 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0015] Figure 13 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0016] Figure 13A is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0017] Figure 14 is a block diagram of one embodiment of a memory package with a compute engine.

[0018] Figure 15 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0019] Figure 16 is a block diagram of one embodiment of a memory die with a compute engine.

[0020] Figure 17 is a block diagram of one embodiment of a solid state drive that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations.

[0021] Figure 18 is a block diagram of one embodiment of a memory die with circuitry under the memory array.

[0022] Figure 19 is a block diagram of one embodiment of a memory die with circuitry under the memory array.

[0023] Figure 20 depicts a portion of a database.

[0024] Figure 21 is a block diagram of a block of non-volatile memory.

[0025] Figure 22 is a block diagram illustrating an implementation of a solid state drive (or other memory system) that comprises a Controller, non-volatile memory for storing data and a compute engine near the location of the data that can be used to perform common data manipulation operations, in accordance with some embodiments. The system of Figure 22 reduces the size in memory of logical address to physical address tables by having separate tables for random and sequential workload to thereby provide more memory resources to the compute engine.

[0026] Figure 23 is a flow chart describing one embodiment of a process for reducing the size in memory of logical address to physical address tables by having separate tables for random and sequential workload to thereby provide more memory resources to the compute engine inside the memory system (e.g. behind the host interface).

[0027] Figure 24 is a flow chart describing one embodiment of a process performed by a host.

[0028] Figure 25 is a flow chart describing one embodiment of a process for operating a memory system using separate logical address to physical address tables for random and sequential workload.

[0029] Figure 26 depicts fine-grained (page based) mapping that has sixteen different map entries (16 physical pages or 2 blocks in total, s sub-blocks/block, 4 pages/sub-block).

[0030] Figure 27 depicts coarse-grained mapping (sub-block based) that has four different map entries.

[0031] Figure 28 depicts coarse-grained mapping (erase block based) that has two different map entries.

DETAILED DESCRIPTION

[0032] A solid state drive (“SSD”), or other memory system, is proposed that can accelerate performance of big data applications using a hardware accelerator based approach that involves integrating compute engines inside the SSD (or other memory system) to perform common data manipulation operations such as scan, filter, aggregate, and join (and other operations). Application Programming Interfaces (“APIs”) can be exposed by the memory system (e.g., exposed by the Controller) and used by the host application to perform efficient data processing using the compute engines in the SSD or other memory system.

[0033] As the compute operations move closer to the data, the available compute-to-data bandwidth increases significantly. For example, compared to conventional systems where compute is located on the host side of the I/O interface (and outside the SSD), integrating the compute engine within a SSD can provide 1.5X-3.75X higher bandwidth, even as NAND architecture and I/O interfaces continue to scale. Similarly, even closer integration with the bits -- by moving the compute engine within a memory die -- can provide an additional 2.5X-13X increase in bandwidth available to the compute engine.

[0034] A compute engine is a hardware circuit that can be one or more of the following:

1. Hardware accelerator for data access and manipulation operations such as filtering, aggregation, compression/decompression, sorting, grouping, and joining tabular or other forms of data;
2. General purpose CPU core such as ARM/ARC/Intel Atom; and/or
3. Reconfigurable hardware circuitry such as FPGAs.

[0035] Unlike other components in an SSD that primarily facilitate reading/writing data from/to the SSD, a compute engine (which is separate from a controller) manipulates the contents of the data. For example, it can search a stream of data looking for a specific key (a set of bytes) or can sort the data in some order as specified by the host. The concept of compute engines includes but is not limited to compression/decompression engines sometimes available in SSDs.

[0036] Typical memory systems include a controller, non-volatile storage connected to the controller, and local high speed volatile memory (e.g., DRAM) connected to and local to the controller. The controller can use the local high speed volatile memory (also referred to as local memory or working memory) as a read buffer, as a write buffer, as a scratch pad and to store logical address to physical address translation tables (“L2P tables”).

[0037] In many systems, the non-volatile storage is addressed internally to the memory system using physical addresses associated with the one or more memory die. However, the host system will use logical addresses to address the various memory locations. This enables the host to assign data to consecutive logical addresses, while the memory system is free to store the data as it wishes among the locations of the one or more memory die. To enable this system, the controller typically performs address translation between the logical addresses used by the host and the physical addresses used by the memory die. One example implementation is to maintain tables (i.e. the L2P tables mentioned above) that identify the current translation between logical addresses (such as logical block addresses, known as LBA’s) and physical addresses (such as physical block addresses, known as PBA’s). An entry in the L2P table may include an identification of a logical address and corresponding physical address.

[0038] Although logical address to physical address tables (or L2P tables) include the word “tables” they need not literally be tables. Rather, the logical address to physical address tables (or L2P tables) can be any type of data structure.

[0039] In some examples, the memory space of a memory system is so large that the local memory cannot hold all of the L2P tables. In such a case, the entire set of L2P tables are stored in the non-volatile storage and a subset of the L2P tables are cached (L2P cache) in the local memory. The bigger the L2P cache, the higher performance of the memory system. In some memory systems, the L2P cache may utilize up to 95% of the local memory. Using such a large portion of the local memory is not a problem during normal operation of the memory system (e.g., during programming, erasing and reading). However, when the local compute engine (discussed above) has been engaged to perform data manipulation operations, that local compute engine needs to use a portion of the local memory as working memory to support the data manipulation operations. If 95% of the local memory is reserved for the L2P cache, then there will not be enough free memory for the compute engine and the compute engine's performance will suffer. Therefore, there is a need to reduce the amount of space (capacity) in the local memory that is reserved for the L2P cache when the local memory is needed by a compute engine, without degrading performance of the memory system.

[0040] To address this need, a non-volatile storage system is proposed that maintains reduced size logical address to physical address tables in volatile working memory by maintaining coarse logical address to physical address tables in the working memory for use with storing sequentially accessed data and storing fine logical address to physical address tables in the working memory for use with storing randomly accessed data. By maintaining reduced size logical address to physical address tables, there will be more room in the working memory for the local compute engine to use during its performance of database operations, thereby, increasing the performance (speed) of the compute engine.

[0041] One embodiment of the memory system (e.g., SSD) is connectable to (or otherwise in communication with) a host. The memory system comprises non-volatile memory (e.g., one or more memory packages including one or more memory die), a controller in communication with the non-volatile memory, working memory connected to the controller and configured to store coarse logical address to physical address tables for use with storing sequentially accessed data and fine logical address to physical address tables for use with storing randomly accessed data, and a compute engine positioned within the memory system and in communication with the controller. The compute engine is configured to receive compiled code from the host that includes instructions for performing one or more database operations on target data in the non-volatile memory and an indication of whether the target

data is sequentially accessed data or randomly accessed data. The compute engine performs the one or more database operations based on the code. The target data is accessed using the coarse logical address to physical address tables if the indication denotes that the target data is sequentially accessed data and the target data is accessed using the fine logical address to physical address tables if the indication denotes that the target data is randomly accessed data.

[0042] Figures 1-19 describe various architectures for memory systems that can implement the technology proposed herein. For example, Figure 1 is a block diagram of one embodiment of SSD 10 that comprises a Controller (SSD Controller) 12, non-volatile memory packages 14 for storing data, DRAM/ReRAM 16, and a compute engine 22 near the location of the data that can be used to perform common data manipulation operations. Figure 1 presents a high-level design where the compute engine 22 is integrated within the SSD Controller 12. The compute engine 22 can be, for instance, an ASIC that is part of the SSD Controller system on a chip (“SoC”), or can be integrated (deeper) as a hardware circuit within the SSD controller. Figure 1 shows the SSD Controller 12, a SoC, including existing SSD Controller components that comprise FTL engines 32, ECC engines 34, and DDR memory controller 36 for controlling DRAM/ReRAM 16. Included within that SSD Controller 12 is the new proposed compute engine 22 that can be used to perform compute operations on data stored in the non-volatile memory of the memory packages. Examples of the compute operations include scanning the data, searching, filtering, sorting, aggregating data, joining data together, as well as other functions on the data. Figure 1 shows the SSD Controller 12 in communication with DRAM/ReRAM 16 and in communication with the set of one or more memory packages 14. In one embodiment, the SSD Controller 12 communicates with the memory packages (and/or memory die) using a Toggle Mode interface, which is an asynchronous interface that is able to communicate 32GB/s. An alternative embodiment could use ONFI interface (Open NAND Flash Interface), which is synchronous and makes use of a clock. DRAM/ReRAM 16 is one example of a local memory (e.g., high speed volatile working memory).

[0043] In the designs considered below, a memory package refers to a structure that contains one or more memory dies (e.g., NAND flash memory dies, ReRAM dies, both NAND and ReRAM, or other non-volatile technologies). The term memory package can also refer to managed memory – i.e. a memory package that contains one or more memory dies with an

embedded error correction code (“ECC”) engine/controller to correct errors detected during read operations to the memory.

[0044] The memory packages include one or more memory die. In one embodiment, each memory die will include its own chip enable that can be controlled by SSD Controller 12. In other embodiments, multiple memory die may share a chip enable, requiring SSD Controller 12 to use addressing to select between the memory die that share a chip enable. In one embodiment, the memory die in the memory packages 14 utilize NAND flash memory. In other embodiments, the memory package can include cross point ReRAM non-volatile memory, which is discussed below.

[0045] Figure 2 is a block diagram of one embodiment of a solid state drive 100 that comprises a controller 102, non-volatile memory 104 for storing data, DRAM/ReRAM 106 and a compute engine 114 near the location of the data that can be used to perform common data manipulation operations. The embodiment of Figure 2 includes an SSD controller 102 comprising a Front End Processor Circuit (FEP) 110 and one or more Back End Processor Circuits (BEP) 112. In one embodiment the FEP110 circuit is implemented on an ASIC. In one embodiment, each BEP circuit 112 is implemented on a separate ASIC. The ASICs for each of the BEP circuits 112 and the FEP circuit 110 are implemented on the same semiconductor such that the SSD controller 102 is manufactured as a SoC. FEP 110 and BEP 112 both include their own processors. In one embodiment, FEP110 and BEP 112 work as a master slave configuration where the FEP110 is the master and each BEP 112 is a slave. For example, FEP circuit 110 implements a flash translation layer, including performing memory management (e.g., garbage collection, wear leveling, etc.), logical to physical address translation, communication with the host, management of DRAM (local volatile memory) and management the overall operation of the SSD (or other non-volatile storage system). The BEP circuit 112 manages memory operations in the memory packages/die at the request of FEP circuit110. For example, the BEP circuit 112 can carry out the read, erase and programming processes. Additionally, the BEP circuit 112 can perform buffer management, set specific voltage levels required by the FEP circuit110, perform error correction (ECC), control the Toggle Mode interfaces to the memory packages, etc. In one embodiment, each BEP circuit 112 is responsible for its own set of memory packages. Figure 2 shows the FEP circuit110 in communication with each of the BEP circuits 112. In the implementation of Figure 2, the compute engine 114 is designed in as a hardware circuit within FEP110. The compute engine

can access high speed, high-bandwidth memory using the DDR interface to access the DRAM 106. In this implementation, the bandwidth available to the compute engine is limited by the bandwidth that connects the FEP110 to the BEP 112. DRAM 106 is one example of a local working memory.

[0046] Figure 3 is a block diagram of one embodiment of an FEP circuit with the compute engine 114 designed into the circuit. The FEP circuit of Figure 3 is one example implementation of FEP circuit110 of Figure 2. Figure 3 shows a PCIe interface 150 to communicate with the host and a host processor 152 in communication with that PCIe interface. The host processor 152 can be any type of processes known in the art that is suitable for the implementation. The host processor 152 is in communication with a network-on-chip (NOC) 154. An NOC is a communication subsystem on an integrated circuit, typically between cores in a SoC. NOC's can span synchronous and asynchronous clock domains or use unclocked asynchronous logic. NOC technology applies networking theory and methods to on-chip communications and brings notable improvements over conventional bus and crossbar interconnections. NOC improves the scalability of SoCs and the power efficiency of complex SoCs compared to other designs. The wires and the links of the NOC are shared by many signals. A high level of parallelism is achieved because all links in the NOC can operate simultaneously on different data packets. Therefore, as the complexity of integrated subsystems keep growing, an NOC provides enhanced performance (such as throughput) and scalability in comparison with previous communication architectures (e.g., dedicated point-to-point signal wires, shared buses, or segmented buses with bridges). Connected to and in communication with NOC 154 is the memory processor 156, the compute engine 114, SRAM 160 and a DRAM controller 162. The DRAM controller 162 is used to operate and communicate with the DRAM (e.g., DRAM 106). SRAM 160 is local RAM memory used by the compute engine 114 or the memory processor 156. The memory processor 156 is used to run the FEP circuit and perform the various memory operations. Also in communication with the NOC are two PCIe Interfaces 164 and 166. In the embodiment of Figure 3, the SSD controller will include two BEP circuits; therefore there are two PCIe Interfaces 164/166. Each PCIe Interface communicates with one of the BEP circuits. In other embodiments, there can be more or less than two BEP circuits; therefore, there can be more than two PCIe Interfaces. In these arrangements, the compute engine 114 is positioned (from the perspective of the host) behind the interface 150 to the host (e.g., on the memory system side of the interface to the host) and behind the API exposed by the Controller (e.g., exposed by the FEP circuit).

[0047] Figure 4 is a block diagram of one embodiment of the BEP circuit. The BEP circuit of Figure 4 is one example implementation of BEP circuit 112 of Figure 2. Figure 4 shows a PCIe Interface 200 for communicating with the FEP circuit(e.g., communicating with one of PCIe Interfaces 164 and 166 of Figure 3). PCIe Interface 200 is in communication with two NOCs 202 and 204. In one embodiment the two NOCs can be combined to one large NOC. Each NOC (202/204) is connected to SRAM (230/260), a buffer (232/262), processor (220/250), and a data path controller (222/252) via an XOR engine (224/254) and an ECC engine (226/256). The ECC engines 226/256 are used to perform error correction, as known in the art. The XOR engines 224/254 are used to XOR the data so that data can be combined and stored in a manner that can be recovered in case there is a programming error. The data path controller is connected to an interface module for communicating via four channels with memory packages. Thus, the top NOC 202 is associated with an interface 228 for four channels for communicating with memory packages and the bottom NOC 204 is associated with an interface 258 for four additional channels for communicating with memory packages. Each interface 228/258 includes four Toggle Mode interfaces (TM Interface), four buffers and four schedulers. There is one scheduler, buffer and TM Interface for each of the channels. The processor can be any standard processor known in the art. The data path controllers 222/252 can be a processor, FPGA, microprocessor or other type of controller. The XOR engines 224/254 and ECC engines 226/256 are dedicated hardware circuits, known as hardware accelerators. In other embodiments, the XOR engines 224/254 and ECC engines 226/256 can be implemented in software. The scheduler, buffer, and TM Interfaces are hardware circuits.

[0048] Figure 5 is a block diagram of one embodiment of a memory package. For example, the memory package of Figure 5 is an example implementation of a memory package included in memory packages 14 of Figure 1 or memory packages 104 of Figure 2. Figure 5 shows a plurality of memory die 292 connected to a memory bus (data lines and chip enable lines) 294. The memory bus 294 connects to a Toggle Mode Interface 296 for communicating with the TM Interface of an BEP circuit (see e.g. Figure 4). In some embodiments, the memory package can include a small controller connected to the memory bus and the TM Interface. The memory package can have one or more memory die. In one embodiment, each memory package includes eight or 16 memory die; however, other numbers of memory die can also be implemented. The technology described herein is not limited to any particular number of memory die.

[0049] Figure 6 is a functional block diagram of one embodiment of a memory die 300. Note that memory is one example implementation of memory die 292 of Figure 5. The components depicted in Figure 6 are electrical circuits. Memory die 300 includes a three dimensional memory structure 326 of memory cells (such as, for example, a 3D array of memory cells), control circuitry 310, and read/write circuits 328. In other embodiments, a two dimensional array of memory cells can be used. Memory structure 326 is addressable by word lines via a row decoder 324 and by bit lines via a column decoder 332. The read/write circuits 328 include multiple sense blocks 350 including SB1, SB2,...,SBp (sensing circuitry) and allow a page of memory cells to be read or programmed in parallel. Commands and data are transferred between to/from memory die 300 via lines 318. In one embodiment, memory die 300 includes a set of input and/or output (I/O) pins that connect to lines 318.

[0050] Memory structure 326 may comprise one or more arrays of memory cells including a 3D array. The memory structure may comprise a monolithic three dimensional memory structure in which multiple memory levels are formed above (and not in) a single substrate, such as a wafer, with no intervening substrates. The memory structure may comprise any type of non-volatile memory that is monolithically formed in one or more physical levels of arrays of memory cells having an active area disposed above a silicon substrate. The memory structure may be in a non-volatile memory device having circuitry associated with the operation of the memory cells, whether the associated circuitry is above or within the substrate.

[0051] Control circuitry 310 cooperates with the read/write circuits 328 to perform memory operations (e.g., erase, program, read, and others) on memory structure 326, and includes a state machine 312, an on-chip address decoder 314, a power control module 316 and a temperature detection circuit 318. The state machine 312 provides die-level control of memory operations. Temperature detection circuit 318 (which is on memory die 300) is configured to detect temperature at the memory structure 326, and can be any suitable temperature detection circuit known in the art. In one embodiment, state machine 312 is programmable by software. In other embodiments, state machine 312 does not use software and is completely implemented in hardware (e.g., electrical circuits). In one embodiment, control circuitry 310 includes registers, ROM fuses and other storage devices for storing default values such as base voltages and other parameters.

[0052] The on-chip address decoder 314 provides an address interface between addresses used by a host or controller to the hardware address used by the decoders 324 and

332. Power control module 316 controls the power and voltages supplied to the word lines and bit lines during memory operations. It can include drivers for word line layers (discussed below) in a 3D configuration, select transistors (e.g., SGS and SGD transistors, described below) and source lines. Power control module 316 may include charge pumps for creating voltages. The sense blocks include bit line drivers.

[0053] Multiple memory elements in memory structure 326 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 flash memory) typically contain memory elements connected in series. A NAND string is an example of a set of series-connected memory cells and select gate transistors that can be used to implement memory structure 326 as a three-dimensional memory structure.

[0054] A NAND flash memory array may be configured so that the array is composed of multiple NAND strings of which a NAND string is composed of multiple memory cells 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 exemplary, and memory cells may be otherwise configured.

[0055] The memory cells may be arranged in the single memory device 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, or in structures not considered arrays.

[0056] A three dimensional memory array is arranged so that memory cells 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 z direction is substantially perpendicular and the x and y directions are substantially parallel to the major surface of the substrate).

[0057] 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 cells. The vertical columns may be arranged in a two dimensional configuration, e.g., in an x-y plane, resulting

in a three dimensional arrangement of memory cells, with memory cells on multiple vertically stacked memory planes. Other configurations of memory elements in three dimensions can also constitute a three dimensional memory array.

[0058] By way of non-limiting example, in a three dimensional NAND memory array, the memory elements may be coupled together to form vertical NAND strings with charge-trapping material that traverse across multiple horizontal memory device levels. One example of a three dimensional NAND memory array that can be used to implement memory structure 126 can be found in U.S. Patent 9,343,156, incorporated herein by reference in its entirety.

[0059] 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.

[0060] A person of ordinary skill in the art will recognize that the technology described herein is not limited to a single specific memory structure, but covers many relevant memory structures within the spirit and scope of the technology as described herein and as understood by one of ordinary skill in the art.

[0061] Although an example memory system is a three dimensional memory structure that includes vertical NAND strings with charge-trapping material, other (2D and 3D) memory structures can also be used with the technology described herein. For example, floating gate memories (e.g., NAND-type and NOR-type flash memory), ReRAM memories, magnetoresistive memory (e.g., MRAM), and phase change memory (e.g., PCRAM) can also be used.

[0062] One example of a ReRAM memory includes reversible resistance-switching elements arranged in cross point arrays accessed by X lines and Y lines (e.g., word lines and bit lines). One example of a three dimensional memory array that can be used to implement memory structure 126 can be found in U.S. Patent Application 2016/0133836, "High Endurance Non-Volatile Storage," incorporated herein by reference in its entirety.

[0063] In another embodiment, the memory cells may include conductive bridge memory elements. A conductive bridge memory element may also be referred to as a programmable metallization cell. A conductive bridge memory element may be used as a state

change element based on the physical relocation of ions within a solid electrolyte. In some cases, a conductive bridge memory element may include two solid metal electrodes, one relatively inert (e.g., tungsten) and the other electrochemically active (e.g., silver or copper), with a thin film of the solid electrolyte between the two electrodes. As temperature increases, the mobility of the ions also increases causing the programming threshold for the conductive bridge memory cell to decrease. Thus, the conductive bridge memory element may have a wide range of programming thresholds over temperature.

[0064] Magnetoresistive memory (MRAM) stores data by magnetic storage elements. The elements are formed from two ferromagnetic plates, each of which can hold a magnetization, separated by a thin insulating layer. One of the two plates is a permanent magnet set to a particular polarity; the other plate's magnetization can be changed to match that of an external field to store memory. This configuration is known as a spin valve and is the simplest structure for an MRAM bit. A memory device is built from a grid of such memory cells. In one embodiment for programming, each memory cell lies between a pair of write lines arranged at right angles to each other, parallel to the cell, one above and one below the cell. When current is passed through them, an induced magnetic field is created.

[0065] Phase change memory (PCRAM) exploits the unique behavior of chalcogenide glass. One embodiment uses a GeTe - Sb₂Te₃ super lattice to achieve non-thermal phase changes by simply changing the co-ordination state of the Germanium atoms with a laser pulse (or light pulse from another source). Therefore, the doses of programming are laser pulses. The memory cells can be inhibited by blocking the memory cells from receiving the light. Note that the use of “pulse” in this document does not require a square pulse, but includes a (continuous or non-continuous) vibration or burst of sound, current, voltage light, or other wave.

[0066] Figure 7 is a block diagram of one embodiment of a solid state drive 400 that comprises a controller 402, non-volatile memory packages 404 for storing data, DRAM/ReRAM 406, and a compute engine 412 near the location for that data that can be used to perform common data manipulation operations. Controller 402 includes FEP circuit 410. In the embodiment of Figure 7, compute engine 412 is integrated within FEP circuit 410 and the one or more BEP circuits 422 are now incorporated within the memory packages 404. In this implementation, the SSD controller contains only one ASIC, for the FEP circuit. That is, the SSD controller 402 is in communication with the memory packages 404, where each memory

package includes multiple memory die 420 and one or more BEP circuits 422. One example embodiment of memory die 420 is depicted in Figure 6. One example of BEP circuit 422 is depicted in Figure 4. One example of FEP circuit 410 with an integrated compute engine 412 is depicted in Figure 3. DRAM/ReRAM 406 is an example of a local memory (e.g., local high speed working volatile memory).

[0067] Figure 8 is a block diagram of one embodiment of a solid state drive 450 that comprises a controller 460, non-volatile memory packages 454 for storing data, DRAM/ReRAM 456, and a compute engine 464 near the location of the data that can be used to perform common data manipulation operations. In the embodiment of Figure 8 the compute engine 464 is a standalone ASIC (application specific integrated circuit) that is integrated with the SSD controller 460 as a SoC. In this implementation, controller 460 includes a FEP circuit 460 in communication with one or more BEP circuits 462. Compute engine 464 is outside of and connected to FEP circuit 462, connected to the BEP circuit and connected to the high speed DRAM memory with separate interfaces. The bandwidth available to the compute engine 464 is lower than or equal to the bandwidth of the embodiment of Figure 2. This implementation is preferred when the development of the FEP circuit 462 and the compute engine 464 needs to be kept separate. One example of BEP circuit 422 is depicted in Figure 4. One example of memory packages 454 is depicted in Figure 5. DRAM/ReRAM 456 is an example of a local memory (e.g., local high speed volatile working memory).

[0068] Figure 9 is a block diagram of one embodiment of a FEP circuit without a compute engine, that is suitable for the embodiment of Figure 8 (e.g., FEP circuit 460). Figure 9 shows all the components of Figure 3, but without the compute engine. That is, Figure 9 depicts PCIe interface 150, host processor 152, NOC 154, memory processor 156, SRAM 160, DRAM controller 162, and PCIe Interfaces 164 and 166. In the embodiment of Figure 9, the SSD controller will include two BEP circuits; therefore there are two PCIe Interfaces. Each PCIe Interface communicates with one of the BEP circuits. In other embodiments, there can be more or less than two BEP circuits; therefore, there can be more or less than two PCIe Interfaces.

[0069] Figure 10 is a block diagram of one embodiment of a solid state drive 600 that comprises a controller 602, non-volatile memory packages 604 for storing data, DRAM/ReRAM 606, and compute engine 616 near the location of the data that can be used to perform common data manipulation operations. Controller 602 includes a FEP circuit 612

connected to one or more BEP circuits 614. In this embodiment a compute engine 616 is integrated with an BEP circuit 614. That is, the compute engine 616 is implemented in the ASIC for the BEP circuit 614. The bandwidth available to the compute engine is now determined by the number of toggle mode channels present in each BEP circuit and the bandwidth of the toggle mode channels. The BEP circuit 614 may also contain an optional interface 620 to connect to the DRAM/ReRAM chip. A direct interface to the high speed memory provides the compute engine 616 with fast access to the memory to store temporary working data. In the absence of a direct interface, temporary working data is streamed via the interface that connects the BEP circuits to the FEP circuit. One example of FEP circuit 612 is depicted in Figure 9. One example of memory packages 604 is depicted in Figure 5. DRAM/ReRAM 606 is an example of a local memory (e.g., local high speed volatile working memory).

[0070] Figure 11 is a block diagram of one embodiment of an BEP circuit that includes a compute engine. The embodiment of the BEP circuit of Figure 11 is appropriate for use in the embodiment of Figure 10 (e.g., as an BEP circuit 614). The components of Figure 11 are the same as the components of Figure 4, but further includes a compute engine 702 connected to the top NOC 202 and a second compute engine 704 connected to the bottom NOC 204. In another embodiment, one compute engine can connect to both NOCs. In another embodiment, the two NOCs are connected together and the combined NOC will connect to one, two or multiple compute engines. In the embodiment of Figure 11, there is one compute engine for each set of four channels. In other embodiments, the channels grouped together can include more or less than four channels.

[0071] Figure 12 is a block diagram of one embodiment of a solid state drive 800 that comprises a controller 802, non-volatile memory packages 804 for storing data, DRAM/ReRAM 806 and a compute engine 824 near the location of the data that can be used to perform common data manipulation operations. Controller 802 includes FEP circuit 820 connected to one or more BEP circuits 822. In the embodiment of Figure 12, compute engine 824 is a standalone ASIC that is connected directly to the toggle mode (TM) channels from the BEP circuits. In such implementations, the compute engine should optionally include an ECC engine in order to decode and correct data read from the flash memory (or other type of nonvolatile memory in the memory packages) before being processed by the compute engine. The compute engine 824 can also be connected to the high speed, high-bandwidth DRAM

memory 806 through a standard DDR interface to the DRAM/ReRAM chip and to FEP circuit 820. One example of FEP circuit 820 is depicted in Figure 9. One example of memory packages 804 is depicted in Figure 5. One example of BEP circuit 822 is depicted in Figure 4. DRAM/ReRAM 806 is an example of a local memory (e.g., local high speed volatile memory).

[0072] The embodiments discussed above show various implementations of integrating the compute engine with the controller. In a different set of implementations, the compute engine can be integrated with the memory package, referred to as memory package level integration.

[0073] Figure 13 is a block diagram of one embodiment of a solid state drive 850 that includes memory package level integration, comprising a controller 852, non-volatile memory packages 854 for storing data, DRAM/ReRAM 856 and a compute engine 862 near the location of the data that can be used to perform common data manipulation operations. Controller 852 includes FEP circuit 858 connected to one or more BEP circuits 860. The one or more BEP circuits 860 connect to the non-volatile memory packages 854. One example of FEP circuit 858 is depicted in Figure 9. One example of BEP circuit 860 is depicted in Figure 4. DRAM/ReRAM 856 is an example of a local memory (e.g., local high speed volatile working memory). In the embodiment depicted in Figure 13, the compute engine is integrated with each memory package. A memory package, which typically includes multiple memory die (e.g., NAND non-volatile memory or other type of non-volatile memory), is now modified to include the compute engine ASIC (also known as a compute engine core) within the memory package. In one embodiment, the memory package should also include an ECC engine (or at least the decoder portion of the ECC engine) to decode code words read from the memory and to correct the data read from the non-volatile memory die before being processed by the compute engine. Thus, compute engine 862 includes an ECC engine. In other embodiments, the compute engine can operate on data that has not been subjected to ECC decoding. The memory package can optionally include high-speed memory like DRAM to support the compute engine with access to temporary working data. As the data management operations are within the memory package, the bandwidth available to the compute engine can be much higher than the toggle mode (TM) bandwidth available outside of the memory package.

[0074] Figure 13A is a block diagram of one embodiment of a solid state drive 880 that includes controller 882, non-volatile memory packages 884 for storing data, and DRAM/ReRAM 886. Controller 882 includes FEP circuit 888 connected to one or more BEP

circuits 890. The one or more BEP circuits 890 connect to the non-volatile memory packages 884. One example of FEP circuit 888 is depicted in Figure 9. One example of BEP circuit 890 is depicted in Figure 4. DRAM/ReRAM 886 is an example of a local memory (e.g., local high speed volatile working memory).

[0075] The embodiment depicted in Figure 13A includes multiple (or distributed) compute engines, such that compute engine 892 is positioned in controller 882 and a set of compute engines (with built-in ECC engine) 894 are positioned in non-volatile memory packages 884. For example, compute engine 892 is a standalone ASIC that is connected directly to the toggle mode (TM) channels from the BEP circuits (the interface between the BEP circuits and the memory packages/die). Compute engine 892 can also be connected to the high speed, high-bandwidth DRAM memory 886 through a standard DDR interface to the DRAM/ReRAM chip and to FEP circuit 888. Compute engine 894 is integrated with each memory package. In one embodiment, the memory package also includes an ECC engine (or at least the decoder portion of the ECC engine) to decode code words read from the memory and to correct the data read from the non-volatile memory die before being processed by the compute engine. Thus, compute engine 894 includes an ECC engine. In other embodiments, the compute engine can operate on data that has not been subjected to ECC decoding. The memory package can optionally include high-speed memory like DRAM to support the compute engine with access to temporary working data. As some data manipulation operations are within the memory package, the bandwidth available to the compute engine can be much higher than the toggle mode (TM) bandwidth available outside of the memory package. In some embodiments, the compute engines 892 and 894 will split up the work performed on the data. For example, code from the hosts can program the system to perform some operations on compute engine 892 and other operations on compute engine 894. For instance, the compute engine 894 could perform error correction coding (ECC) function along with simple application level tests, and the compute engine 892 could be executing a flash translation layer (FTL) optimized for sequential or indexed-sequential workloads, along with more complex filtering, sorting and grouping functions at the application query level.

[0076] Figure 14 is a block diagram of one embodiment of a memory package that includes a compute engine. The embodiment of 14 can be used to implement one of the memory packages 854 in Figure 13 or memory packages 884 of Figure 13A. The memory package of Figure 14 includes a plurality of memory die 904 connected to a memory bus 906

(analogous to the memory bus of Figure 5). Memory bus 906 is connected to a TM interface 908 for communicating with an BEP circuit. Additionally, Figure 14 shows a compute engine 910 connected to the memory bus and to an ECC engine 912. The ECC engine 912 is also connected to memory bus 906. Memory read from a memory die can be subjected to ECC decoding (including fixing errors) and then presented to the compute engine 910 to perform any of the compute operations discussed herein.

[0077] Figure 15 is a block diagram of one embodiment of a solid state drive 950 that comprises a controller 952, non-volatile memory packages 956 for storing data, DRAM/ReRAM 954, and a compute engine near the location of that data that can be used to perform common data manipulation operations. Controller 952 includes FEP circuit 960 connected to one or more BEP circuits 962. The one or more BEP circuits 962 connect to the non-volatile memory packages 956. One example of FEP circuit 960 is depicted in Figure 9. One example of BEP circuit 962 is depicted in Figure 4. DRAM/ReRAM 954 is an example of a local memory (e.g., local high speed volatile working memory).

[0078] The embodiment of Figure 15 implements memory package level integration. For example, each memory package includes multiple memory die and a compute engine 970 integrated within each memory die 972. In one embodiment, the compute engine will include an ECC engine to decode (including correcting) data read from the memory die. The ECC engine can be part of the compute engine or separate from the compute engine but otherwise included in the memory die.

[0079] For purposes of this document, the concept of the working memory (or local memory or DRAM discussed above) being connected to the controller includes the working memory being inside the controller. A compute engine that is in communication with the controller can be inside or separate from the controller.

[0080] Figure 16 is a block diagram of one embodiment of a memory die 1000 that includes a compute engine. For example, the memory die 1000 is an example implementation of memory die 972 of Figure 15. The embodiment of Figure 16 includes the elements of the embodiment of Figure 6. For example, memory die 1000 includes a three dimensional memory structure 326 of memory cells (such as, for example, a 3D array of memory cells), control circuitry 310, read/write circuits 328, row decoder 324 and column decoder 332. Control circuitry 310 includes state machine 312, on-chip address decoder 314, a power control module

316 and a temperature detection circuit 318. Additionally, in the embodiment of Figure 16, control circuitry 310 further includes ECC engine 1017 and compute engine 1019. Data read from the memory structure 326 is decoded using ECC engine 1017 and provided to compute engine 1019 for performing various compute operations, as discussed herein.

[0081] While the embodiments discussed above show the SSD controller to be implemented as a two ASIC solution containing a BEP ASIC and a FEP ASIC, it is also possible that the SSD controller is implemented with more or less than two ASICs. In that case, the design space can be expanded to place the compute engine within any one or more of the ASICs. Additionally, the compute engine can be placed outside of the ASICs. In other embodiments, the SSD controller can include different architectures, other than the FEP/BEP architecture. Even in the other architectures, the SSD controller can still be configured to include a compute engine inside one of the ASICs or circuits or modules. Additionally, a compute engine can be added to SSDs that are not implemented using ASICs, but implemented using other hardware.

[0082] The embodiment of Figure 15 includes integrating the compute engine within the memory die (such as a NAND memory die or ReRAM memory die). Figure 17 is a block diagram providing additional details for implementing an embodiment of the system of Figure 15. Specifically, Figure 17 shows a host in communication with a SSD 1100 (implemented on a printed circuit board) that includes a Big NVM controller 1102 and a Small NVM controller 1114. The Big NVM controller 1102 is in communication with DRAM 1104 and memory package 1106.

[0083] In one embodiment, memory package 1106 includes several memory die 1110, optional DRAM (or MRAM/RRAM/PCM/eDRAM) 1112, and Small NVM Controller 1114. Each of the memory die 1110 has an on die compute engine (CE). In one embodiment the on die compute engine is implemented using CMOS technology on the top surface of a substrate and under the monolithic three-dimensional memory array. Potentially, eDRAM/STT-MRAM/PCM as well as SRAM can be integrated. The on die compute engine (CE) can perform some of the data manipulation operations.

[0084] In one embodiment, Small NVM Controller 1114 includes a compute engine (CE) that can perform some of the data manipulation operations. Small NVM Controller 1114 can communicate with the internal memory dies and external chips (i.e. Big NVM controller

and DRAM in Figure 17). Optional DRAM 1112 is used for the Small NVM Controller 1114 to store working data sets. By off-loading computation from the Big NVM Controller 1102 to Small NVM controller with a compute engine (CE) 1114 and the simple CE of the memory die 1110, the external DRAM requirement and communication overhead can be reduced.

[0085] Figure 17 shows that each of Big NVM Controller 1102, DRAM 1104, memory die 1110, DRAM 1112 and Small NVM Controller 1114 can be implemented on separate silicon die in three different packages mounted on one printed circuit board. Thus, Figure 17 provides a big and small NVM controller architecture. The Big NVM Controller 1102 interfaces with the host and DRAM. The Small NVM Controller 1114 can be inside any of the memory packages. The Small NVM Controller 1114 includes a computational engine with optional DRAM and manages multiple NVM channels. A mapping table can be stored in the optional DRAM (or MRAM/PRAM).

[0086] Figure 18 is a block diagram of one embodiment of a memory die 1200 with circuitry under the memory array. Figure 18 shows a monolithic three-dimensional memory structure 1202 with multiple layers. Underneath the memory structure 1202 is circuitry 1204 that is implemented on the top surface of the substrate 1206 and under the memory array 1202. In one embodiment, the circuitry 1204 is implemented using CMOS technology. For example, simple computational logic can be integrated in the CMOS logic under the memory array 1202 potentially with eDRAM/STT-MRAM/PCM as well as SRAM/latches. Simple circuitry logic (i.e., randomizer, ID generator, PUF, or AES) and simple error management logic (i.e., error location map or a simple error avoiding algorithm such as read reference optimizer) as well as ECC can be integrated in the CMOS logic under the memory array 1202 as examples of the compute engine discussed above. This improves latency and performance by eliminating data transfer overhead from the memory die to the separate controller die. An FPGA could be integrated, supporting multiple configurations with a single system on a chip as an aforementioned compute engine. An FPGA can be integrated, supporting multiple configurations within a system on a chip.

[0087] Additionally, other functions can be integrated as an aforementioned compute engine. For example, a CPU or parallel computational engine can be integrated as an aforementioned compute engine. An SIMD engine ("GPU"), neural network, DSP engine (e.g., image/audio processing), digital logic operation (multiplication, addition, subtraction, XOR, etc.), data mining (apriori, k-means, pagerank, decision tree) or pattern matching (i.e.,

Hamming distance calculation), FPGA fabric supporting multiple configurations in the memory die, high speed I/O circuits with memory equalizers, circuits for optical or capacitor/inductive coupling based on interconnections can also be used. In one embodiment, the compute engine needs to be able to work with encrypted data when AES is bypassed for specific applications. In some embodiments, the compute engine may need to work with erroneous data when ECC is bypassed for specific applications

[0088] Figure 19 is a block diagram of one embodiment of a memory die 1300 with circuitry 1304 under the memory array 1302 for using the non-volatile memory die 1300 as a non-volatile-FPGA. The memory die 1300 will include a three-dimensional monolithic memory array 1302. Implemented on the top surface of the substrate 1306, and under the memory array 1302, will be CMOS logic 1304 that implements a FPGA to be used as a compute engine (per the discussion above). This system will use the memory array 1302 (NAND or other type of non-volatile memory) as configuration storage for the reconfigurable logic 1304 of the FPGA. That is, configuration data stored in memory array 1302 is used to configure the FPGA's. This will make the FPGA non-volatile. This will allow for fast boot up compared to conventional FPGAs, which require a reading of configuration data from a discrete nonvolatile memory device to the volatile FPGA cell array. When the FPGA (hardware accelerator/compute engine) is not needed, the configuration storage (the memory array) can be used as just normal non-volatile storage, saving idle power.

[0089] In many systems, the complete set of L2P tables are stored in the non-volatile memory (e.g., the Memory Packages), with a subset being cached within local working memory (e.g., see DRAM [or ReRam or SRAM] of Figures 1, 2, 7, 8, 10, 12, 13, 13A, 15, and 17, or another buffer for the Controller). When there is a cache miss because the needed L2P table (or portion of table) is not within the cache of the local working memory, then the needed L2P table (or portion of table) is retrieved from the non-volatile memory and loaded into the cache of the local working memory. The L2P tables consume a significant amount of space in the local working memory.

[0090] In certain embodiments, the compute engine discussed above with respect to Figures 1-19 will also use the local working memory to stored data while performing its computational operations as described above. The use of the local working memory by the compute engine may reduce the space available for the cache storing the L2P tables, thereby, causing a greater "miss" rate for the cache and degrading performance. Limiting the amount

of local working memory available to the compute engine will reduce the performance of the compute engine

[0091] To alleviate the performance issues related to too much demand on the local working memory, it is proposed to reduce the amount of space needed to store L2P tables in the working memory without significantly increasing the “miss” rate by maintaining coarse logical address to physical address tables in the working memory for use with sequentially accessed data and maintaining fine logical address to physical address tables in the working memory for use with randomly accessed data. Certain embodiments disclosed herein take advantage of special access patterns of database applications (or other applications) to use a lower granularity of L2P tables (e.g., erase block level entry L2P tables) for data with known sequential access patterns. Other intermediate levels of granularity such as sub-erase block or word line can also be used with these known sequential access patterns. Page level entries are continued to be used for data with known random access patterns.

[0092] Figure 20 depicts a portion of an example database and illustrates an example of randomly accessed data and sequentially accessed data. Figure 20 shows two levels of branch nodes and one level of leaf nodes. The branch nodes act as indexes to lower level nodes. The leaf nodes can include many data entries (e.g., in this case names). The first level of branch nodes includes nodes B, F, H and Sam. Node B points to nodes Ba, C and Cr of the second level of branch nodes. Node F points to nodes Fe, Kar and Lu of the second level of branch nodes. Node N points to nodes Na, Pa and Ph of the second level of branch nodes. Node Sam points to nodes Samm, St and V of the second level of branch nodes. Branch node Kar points to leaf nodes Krl, Kathy, Ken and Kim. Branch node Lu points to leaf nodes Luis, Mark, Mary, Mike and Manny. Branch node Na points to leaf nodes Nan, Nancy, Nicky, Nicole and Norm. Branch node Pa points to leaf nodes Pablo, Paul, Paula and Peter. Branch node Ph points to leaf nodes Phil, Pierre, Rachel, Rajiv and Raoul.

[0093] In the example of Figure 20, the branch nodes are accesses randomly and, therefore, are determined to be randomly accessed data (for a random workload). The leaf nodes are stored in consecutive addresses (i.e. in alphabetical order) in the non-volatile memory as sequentially accessed data; therefore, they can be written into the non-volatile memory together as such data is serialized using a write buffer between the host and the non-volatile memory for higher performance.

[0094] Log-structured file systems optimized for the non-volatile memory can support aggregation of multiple random write operations into a single sequential write operation. Then, it typically requires sequential reading of the selected leaf blocks for finding a specific node (e.g., finding a specific last name) when it needs to be searched. HDFS (Hadoop File System) and analytics databases are additional examples that require sequential access to the non-volatile memory. Thus, the SSD (or other device) workload for the database application can be divided into random and sequential workloads.

[0095] For example, the non-volatile memory can be divided into (at least) two sections. A first section (or portion) of memory (e.g., a first set of memory cells) is configured to store sequentially accessed data (e.g., leaf nodes). A second section (or portion) of memory (e.g., a second set of memory cells) is configured to store randomly accessed data (e.g., meta data, pointers, branch nodes, etc.). The first section/portion is physically separate from the second section/portion. In one example implementation, a first set of memory die for a package or system comprise the first section and a second set of memory die for the package or system comprise the second section. The different sections can also be on the same memory die. L2P tables for data stored in the second section (i.e., store randomly accessed data) will include entries for each page of data and are referred to fine L2P tables (or fine logical address to physical address tables). L2P tables for data stored in the first section (i.e., sequentially accessed data) will include entries at the block level and are referred to a coarse L2P tables (or coarse logical address to physical address tables). In one example of a coarse L2P table, there is one L2P entry per block, where the block is the unit of erase. In one example, each L2P entry includes a logical address used by the host and a physical address in the non-volatile memory. In other embodiments, coarse L2P tables for data stored in the second section (i.e., sequential data) will include entries at the word line level or other sub-block level.

[0096] Because the L2P tables for data stored in the first section (i.e., sequentially accessed data) will include entries at the block level, there will be many less L2P entries. For example, Figure 21 is a block diagram of an example block of non-volatile memory that includes sixty four word lines (WL0-WL63), with each word line having three pages (Page 0, Page 1 and Page 2). Thus, the example block of Figure 21 includes 192 pages. Other blocks can have more or less than sixty four word lines and more or less than three pages per word line. If the example block of Figure 21 is used to store random data than the corresponding L2P table would include 192 L2P entries. However, if the example block of Figure 21 is used

to store sequential data than the corresponding L2P table would include 1 L2P entry, which saves 191 L2P entries' worth of space. As a typical system can have thousands or millions of blocks, the savings for only storing 1 L2P entry per block (that is, the savings for using coarse L2P tables instead of fine L2P tables) is very significant. The reduction in memory space needed by coarse L2P tables opens up room in the local working memory for the local compute engine, without increasing the "miss" rate for the cache of L2P tables in the local working memory. Because the local compute engine will have more local (volatile) working memory (e.g., DRAM) to work with, computational performance will increase.

[0097] Figure 22 is a block diagram of a block of non-volatile memory system (e.g., a SSD or other device) that reduces the amount of space needed to store L2P tables in the working memory without significantly increasing the "miss" rate by maintaining coarse L2P tables in the working memory for use with sequentially accessed data and maintaining fine L2P tables in the working memory for use with randomly accessed data. The block diagram of Figure 22, while most similar to Figure 1, can be used to generalize and applies to any of the embodiments of Figures 1-2, 7-8, 10, 12, 13, 13A, 15, and 17-19. That is, the discussion herein about the technology for reducing the amount of space needed to store L2P tables in the working memory by maintaining coarse L2P tables and fine L2P tables applies to and can be implemented by all of the embodiments of Figures 1-2, 7-8, 10, 12, 13, 13A, 15, and 17-19.

[0098] Figure 22 depicts SSD (or other device) 1400 that includes controller 1402, non-volatile memory 1404 (e.g., Memory Packages) and volatile working memory (e.g., DRAM) 1406. Controller 1402 includes, among other things, a host interface 1420 and a local compute engine 1422 (see more complete discussion above of the compute engine). In other embodiments that incorporate the proposed technology for reducing the amount of space needed to store L2P tables in the working memory, as described above, the compute engine can be outside of the controller, including in a memory package (see e.g., the embodiments of Figures 13, 13A, 15 and 17-19) and/or in a memory die (see e.g., the embodiments of Figures 15 and 17-19).

[0099] Non-volatile memory 1404 include an area 1430 for storing a complete set of L2P tables, section/portion 1432 configured to store randomly accessed data and a section/portion 1434 configured to store sequentially accessed data. Section 1432 is physically separate from the second section 1434. In other embodiments, the randomly accessed data and sequentially accessed data can be stored in a common section. In some embodiments, the

coarse L2P tables (for sequentially accessed data) are stored in a same area as the fine L2P tables (for randomly accessed data). In other embodiments, the coarse L2P tables (for sequentially accessed data) are stored separately from the fine L2P tables (for randomly accessed data). In one example, the fine L2P tables (for randomly accessed data) are stored in the section 1432 and the coarse L2P tables (for sequentially accessed data) are stored in the section 1434.

[00100] In one embodiment, volatile working memory 1406 includes sequential L2P cache 1440, random L2P cache 1442, working data for compute engine 1444 and miscellaneous region 1446. Sequential L2P cache 1440 is a cache that stores a subset of the coarse L2P tables. Random L2P cache 1442 is a cache that stores a subset of the fine L2P tables. Working data for compute engine 1444 is memory space used by the local compute engine to perform data computations, as discussed above. Miscellaneous region 1446 is used by controller 1402 for other purposes. Because the coarse L2P tables stored in Sequential L2P cache 1440 are much smaller than the fine L2P tables, there is more room in volatile working memory 1406 for a bigger working data for compute engine 1444, which results in higher performance for the local compute engine 1422.

[00101] In one example implementation, compute engine 1422 is configured to use the coarse logical address to physical address tables of the sequential L2P cache 1440 to perform address translation between logical and physical addresses for the leaf nodes for one or more database data structures stored in the portion 1434 of non-volatile memory 1404 and compute engine 1422 is configured to use the fine logical address to physical address tables of random L2P cache 1442 to perform address translation between logical and physical addresses for the branch nodes for one or more database data structures stored in portion 1432 of non-volatile memory 1404.

[00102] As discussed above, the memory system treats sequentially accessed data different than randomly accessed data. In one embodiment, the memory system (based on the operations being performed with the data) determines whether the data is sequentially accessed data or randomly accessed data. In another embodiment, the host determines whether the data is sequentially accessed data or randomly accessed data. Figure 22 depicts a host 1460, in communication with memory system 1400 (via host interface 1420 of controller 1402), that determines whether the data is sequentially accessed data or randomly accessed data. Host 1460 includes a Just-In-Time Compiler 1462 and a memory interface 1464. Host 1460

determines whether target data being accessed by a set of one or more database operations is randomly accessed data or sequentially accessed data. Just-In-Time Compiler 1462 compiles code to perform the set of one or more database operations on a local compute engine 1422 located inside memory system 1400 and embeds in the compiled code an indication of whether the target data is randomly accessed data or sequentially accessed data. Memory interface 1464 is used to communicate with host interface 1420 of memory system 1400, including transmitting the compiled code with the embedded indication to controller 1402 and/or compute engine 1422 (via host interface 1420). The target data can be accessed by the compute engine performing the address translation or the compute engine requesting (or working with) the controller to perform the address translation. Thus, the compute engine can perform memory reads (or writes) or the compute engine can work with the controller to perform a memory read (or write).

[00103] Figure 23 is a flow chart describing one embodiment of a process for reducing the size in memory of L2P tables by having separate tables for random and sequential workload to thereby provide more memory resources to a local compute engine. In one embodiment, the process of Figure 23 is performed by memory system 1400 of Figure 22 (representing any of the embodiments of Figures 1-2, 7-8, 10, 12, 13, 15, and 17-19). In step 1502, the controller (or other entity) maintains reduced size L2P tables in volatile working memory by storing fine L2P tables in the working memory for use with randomly accessed data and coarse L2P tables in the working memory for use with sequentially accessed data. In step 1504, the controller performs memory operations using the reduced size L2P tables. For example, memory operations (e.g., program, erase, read) for sequentially accessed data are performed using coarse L2P tables and memory operations for randomly accessed data are performed using fine L2P tables. In step 1506, the compute engine (inside the controller or outside the controller, or same die as controller or different die than controller) uses the volatile working memory, including extra space due to reduced size L2P tables, to store working data while performing one or more database operations. Steps 1502-1506 can be performed sequentially or concurrently.

[00104] Figure 24 is a flow chart describing one embodiment of a process performed by a host when operating a memory system (e.g., such as the memory systems of Figures 1-2, 7-8, 10, 12, 13, 15, 17-19 and 22) performing the process of Figure 23. In step 1602, the host determines whether target data being accessed by a set of one or more database operations is

randomly accessed data or sequentially accessed data. In step 1604, the host (e.g., Just-In-Time Compiler 1462) compiles code to perform the set of one or more database operations on a compute engine located inside a memory system. The code is compiled for the specific compute engine chosen to perform the task. In step 1606, the host (e.g., Just-In-Time Compiler 1462) embeds in the compiled code an indication of whether the target data is randomly accessed data or sequentially accessed data. In step 1608, the host (via memory interface 1464) transmits the compiled code with the embedded indication to the compute engine (e.g. via host interface 1420 of controller 1402 or any of the other controller architectures described above).

[00105] Figure 25 is a flow chart describing one embodiment of a process for operating a memory using separate logical address to physical address tables for random and sequential workload. In one embodiment, the process of Figure 25 is performed by the memory system (any of the embodiments of Figures 1-2, 7-8, 10, 12, 13, 15, 17-19 and 22) in response to step 1608 of Figure 24. In one embodiment, the system of Figure 22 (representing any of the embodiments of Figures 1-2, 7-8, 10, 12, 13, 15, and 17-19) will perform the process of Figure 25 on the database of Figure 20.

[00106] In step 1702 the memory system receives the compiled code (see step 1608) from the host. For example, the compiled code is received at host interface 1420 and forwarded to another portion of the controller and/or to the compute engine. The compiled code includes instructions for performing one or more database operations on target data in the non-volatile memory. The compiled code includes one or more logical addresses for the target data that need to be translated to physical addresses in the non-volatile memory. The compiled code also includes an indication of whether the target data is sequentially accessed data or randomly accessed data. The indication discussed herein can be a flag or variable that has one value for sequentially accesses data and another value for randomly accessed data. Alternatively, the indication can be the logical address used by the host, where a first set of logical addresses are reserved for and indicate sequential data and a second set of logical addresses are reserved for and indicate random data.

[00107] In step 1704, the memory system determines whether the target data is sequentially accesses data or randomly accessed data based on reading the indication from the compiled code. This can be performed by the compute engine and/or the controller.

[00108] In another embodiment, the memory system classifies the target data as random or sequential and stores randomly accessed data separately from sequentially accessed data. Therefore, based on the data received or addressed the memory system determines whether the job is for random or sequential data. In another embodiment, the memory system can determine whether the job is for randomly accessed data or sequentially accessed data based on the instructions from the host. For example, if the host is requesting sequential reads then the system assumes the job is for sequentially accessed data and if the host is requesting random reads then the system assumes the job is for randomly accessed data. In another embodiment, the memory system can determine whether the job is for randomly accessed data or sequentially accessed data by utilizing a data structure related to the database. In some embodiments, the instructions from the host are instructions for the local compute engine to perform computations on a set of data. If those instructions are accessing data previously classified as sequentially accessed data then the system assumes the job is for sequentially accessed data and if those are accessing data previously classified as randomly accessed data then the system assumes the job is for randomly accessed data. Alternatively, if those instructions require a scan operation then the system assumes the job is for sequentially accessed data and if those instructions require a seek operation then the system assumes the job is for randomly accessed data. Similarly, if the instructions cause a sequential writing of data then the operation then the system assumes the job is for sequentially accessed data and if the instructions cause a single write then the operation then the system assumes the job is for randomly accessed data.

[00109] Looking back at step 1704, if the controller and/or compute engine determines that the job is for randomly accessed data, then in step 1708 the controller and/or compute engine attempts to access the L2P entry for the address received in step 1702 in the appropriate L2P table in the volatile working memory (e.g., DRAM). Since the job is for randomly accessed data, the system attempts to access the L2P entry for a page address in a fine L2P table for random data in the cache of fine L2P tables, such as random L2P cache 1442. Although all L2P tables are stored in the non-volatile memory, only a subset of L2P tables are stored in the relevant cache. Therefore, when the system attempts to access the L2P entry in the random L2P cache it is possible that the L2P entry is in the cache (“hit”) or is not in the cache (“miss”). If the L2P entry for the page address received in step 1702 is not in the cache so that there is a cache miss (step 1710), then in step 1712 the system fetches the L2P table that includes the page address received in step 1702 from non-volatile memory (e.g., from area 1430 of non-volatile memory 1404). In step 1714, the fetched L2P table is stored in the cache

(e.g., in random L2P cache 1742). In step 1716, the L2P entry for the page address received in step 1702 that is in the fetched L2P table is used by the controller and/or compute engine to translate the page address received in step 1702 from a logical address to a physical page address in the non-volatile memory. In step 1718, one or more operations are performed on the random data associated with the physical page address from step 1716 in response to the instructions received from the host in step 1702. Step 1718 includes the compute engine performing the one or more database operations based on the compiled code, including performing one or more memory operations on randomly accessed target data where the target data is accessed using the fine logical address to physical address tables (see step 1716) since the indication denotes that the target data is randomly accessed data. Examples of database operations include scanning the data, searching, filtering, sorting, aggregating data, joining data together, as well as other functions on the data. When performing step 1718, the compute engine will use the local working memory as a scratch pad or for other use.

[00110] Note that if in step 1710 it is determined that there was a cache hit because the relevant L2P entry was resident in the random L2P cache, then the process proceeds to step 1716 to translate the page address received in step 1702 from the logical address to the physical page address in the non-volatile memory, as described above.

[00111] If, in step 1704, the controller and/or compute engine determines that the job is for sequentially accessed data, then in step 1728 the system attempts to access the L2P entry for a block address (or sub-block or word line or other) in a coarse L2P table for sequential data in the cache of coarse L2P tables, such as sequential L2P cache 3140. If the L2P entry for the page address received in step 1702 is not in the cache so that there is a cache miss (step 1730), then in step 1732 the system fetches the L2P table that includes the block address (or sub-block or word line or other) from non-volatile memory (e.g., from area 1430 of non-volatile memory 1404). In step 1734, the fetched L2P table is stored in the cache (e.g., in sequential L2P cache 1440). In step 1736, the L2P entry for the block address received in step 1702 that is in the fetched L2P table is used to translate the block address received in step 1702 from a logical block address to a physical block address in the non-volatile memory. If the address received in step 1702 is not a block level address, then it is converted into a block address. Note that steps 1716 and 1726 comprise performing address translation for a logical address in the compiled code using coarse L2P tables or fine L2P tables based on the embedded indication.

[00112] In step 1738, the compute engine performs the one or more database operations based on the compiled code, including performing one or more memory operations on sequentially accessed target data where the target data was accessed using the coarse logical address to physical address tables (see step 1736) since the indication denotes that the target data is sequentially accessed data. When performing step 1736, the compute engine will use the local working memory as a scratch pad or for other use.

[00113] Note that if in step 1730 it is determined that there was a cache hit because the relevant L2P entry was resident in the sequential L2P cache, then the process proceeds to step 1736 to translate the page address received in step 1702 from the logical address to the physical block address (or sub-block or word line or other) in the non-volatile memory, as described above.

[00114] In the discussion of step 1732 above, it was assumed that the non-volatile memory stored a full set of fine L2P tables and a full set of coarse L2P tables. In another embodiment, the non-volatile memory stores a full set of fine L2P tables and no (or a subset of) coarse L2P tables. In this latter embodiment, when the system needs to fetch a coarse L2P table in step 3332 it will instead fetch one or more corresponding fine L2P tables and use the fetched one or more corresponding fine L2P tables to generate the needed coarse L2P table on demand (i.e. on the fly). Thus, in one example implementation, the controller is configured to create the coarse logical address to physical address tables for the working memory on demand from the fine logical address to physical address mapping tables in the non-volatile memory.

[00115] While the discussion above has focused on a database example, in certain embodiments the principle can be generalized to different host workloads and applications. One embodiment reorganizes the embedded DRAM L2P tables to efficiently support both small block (e.g., 4 KB) writes (e.g., random data operations) and large block writes (100 KB or more) (e.g., sequential data operations). For workloads with a relatively large amount of data requiring large block writes (e.g., sequential data writes), this significantly reduces the amount of DRAM required to hold the L2P cache. One embodiment assumes a workload with these properties: small block (4 KB or so) reads and writes to a relatively small portion of the overall logical address space; and large block (100 KB or more) writes to the remainder of the logical address space, with both large and small random reads. Applications with these properties allow for much more DRAM-efficient mappings for the large block writes: a single mapping entry can represent a much larger region of storage. If the large block writes are

performed contiguously to physical storage, the same mapping entry allows small block reads using simple offsets into the large contiguously written block. Examples of applications that exhibit this behavior include: HDFS (Hadoop File System); Databases using Log-Structured Merge (LSM) trees (including, e.g., Cassandra, LevelDB, and RocksDB); Log-structured file systems, such as ZFS and its derivatives; storage software implementations that perform write serialization for optimized RAID and improved flash performance; and scan-heavy analytics applications.

[00116] Figures 26-28 depicts L2P tables that can support sequentially accessed data and randomly accessed data. Such tables that supports both small and large block writes can be organized in a variety of ways. These include partitioning for sequentially and randomly written data. Figures 26-28 provide more description with examples.

[00117] Figure 26 depicts an embodiment of fine-grained (page based) mapping in a persistent memory that has sixteen different map entries (16 physical pages or 2 blocks in total, 8 sub-blocks/block, 4 pages/sub-block). The fine-grained (page based) L2P tables of Figure 26 are for randomly accessed data. Figure 26 shows blocks divided into two sub-blocks. Each sub-block has a set pages. Each page represents data stored in memory cells connected to a common word line. Two types of mapping are depicted: direct mapping and hash mapping. With direct mapping, logical page addresses are mapped to physical page addresses in sequence. With hash mapping, a hash table is used to map a logical page address to any of the physical page addresses, in any order.

[00118] Figure 27 depicts an embodiment of coarse-grained mapping (sub-block based) that has four different map entries. The logical block address can be mapped to physical block address for sequentially accessed data. Figure 27 also teaches to use offsets for the logical block addresses and the physical block addresses to access individual pages so that sequentially accessed data can be accessed randomly. Therefore, the controller and/or the compute engine can be configured to perform a random read of data written as sequentially accessed data using a logical block address and an offset. Two types of mapping are depicted: direct mapping and hash mapping. With direct mapping, logical block addresses are mapped to physical block addresses in sequence. With hash mapping, a hash table is used to map a logical block address to any of the physical block addresses, in any order.

[00119] Figure 28 depicts another embodiment of coarse-grained mapping (erase block based) that has two different map entries. The logical block address can be mapped to physical block address for sequentially accessed data. Figure 28 also teaches to use sub-block offsets and page offsets for the logical page addresses and the physical page addresses to access individual pages so that sequentially accessed data can be accessed randomly. Two types of mapping are depicted: direct mapping and hash mapping.

[00120] In one embodiment, fine grained L2P tables (e.g., page address based) are stored in a persistent memory (i.e. NAND flash, ReRAM, PCM, etc.). Logical page addresses can be directly mapped to physical page address, or it can be mapped to certain page addresses through a hash function. In one embodiment, when the table is brought to local working memory (e.g., the L2P cache) for the sequentially written data, the fine grained table is reduced into the coarse grained table which is either of sub-block based or erase block based (or other granularity coarser than page level, as discussed above). For this sequentially written data, logical page addresses can be randomly read through sub-block offset and page offset values in the logical page address with the coarse grained table that stores sub-block or erase block address translation table. As illustrated in Figures 26-28, for the case that has 16 physical pages with 2 sub-blocks/block and 4 pages/sub-block, compared to the fine-grained mapping based on page that has 16 different table (also known as a map) entries, coarse-grained mapping based on sub-block or erase block significantly reduce the number of table entries (e.g. 4 for sub-block mapping, 2 for erase-block mapping).

[00121] An extent table contains an ordered set of entries mapping variable sized, contiguous regions of logical addresses to physical locations in storage. This would allow the mappings for an arbitrary mix of various sized writes to be kept in a single table. Extent tables of this type can be implemented as a pool of fixed-size records managed as a tree (binary, red-black, B+, etc.) or skip list. Such an organization makes it easy to "page" records between DRAM (the local working memory) and suitable for storage if the entire table cannot fit in DRAM. Other ways to support multiple data sizes in an L2P table should be apparent to someone skilled in the art.

[00122] One embodiment includes an apparatus, comprising a memory system. The memory system includes non-volatile memory, a controller in communication with the non-volatile memory, working memory connected to the controller and configured to store a coarse logical address to physical address table for use with storing sequentially accessed data related

to database operations and a fine logical address to physical address table for use with storing randomly accessed data related to database operations, and a compute engine positioned within the memory system and in communication with the controller. The compute engine is configured to: receive code from a host, the code including instructions for performing one or more database operations on target data in the non-volatile memory and an indication of whether the target data is sequentially accessed data or randomly accessed data, perform the one or more database operations based on the code, and the target data being accessed using the coarse logical address to physical address table in response to the indication denoting that the target data is sequentially accessed data, perform the one or more database operations based on the code, and the target data being accessed using the fine logical address to physical address table in response to the indication denoting that the target data is randomly accessed data.

[00123] In one example implementation, the code received from the host is compiled code that includes the indication of whether the target data is sequentially accessed data or randomly accessed data, the controller includes a host interface, the compute engine is behind the host interface, the host interface receives the compiled code from the host and forwards the compiled code to the compute engine, and the compute engine is configured to use the controller to perform address translation between logical and physical addresses for the target data based on the coarse logical address to physical address tables and the fine logical address to physical address tables in the working memory.

[00124] In one example implementation, the non-volatile memory includes a memory package that is separate from and connected to the controller and the memory package includes one or more non-volatile memory dies and the compute engine. Additionally, the compute engine is positioned on a first memory die of the one of the non-volatile memory dies; the code received from the host is compiled code that is received by the compute engine on the first memory die; and the compute engine executes the compiled code to access the target data on the first memory die based on the coarse logical address to physical address tables in the working memory for sequentially accessed and the fine logical address to physical address tables in the working memory for data randomly accessed data.

[00125] One embodiment includes a method, comprising: maintaining, in a memory system, reduced size logical to physical (L2P) tables in volatile working memory by maintaining coarse L2P tables in the working memory for use with sequentially accessed data and maintaining fine L2P tables in the working memory for use with randomly accessed data;

receiving, by the memory system from a host, compiled code instructing a compute engine within the memory system to perform a set of one or more database operations on target data in the memory system, wherein the compiled code is embedded with an indication of whether the target data is randomly accessed data or sequentially accessed data; using, by the compute engine inside the memory system, the compiled code to perform the set of one or more database operations on the target data using the working memory, including using space in the working memory made available due to reduced size L2P tables, to store working data for the set of one or more database operations; and performing address translation for a logical address in the compiled code using coarse L2P tables or fine L2P tables based on the embedded indication.

[00126] One embodiment includes a memory system, comprising: non-volatile memory including a first set of memory cells configured to store branch nodes of a database as randomly accessed data and a second set of memory cells configured to store leaf nodes of the database as sequentially accessed data; working memory; and a controller in communication with the non-volatile memory and the working memory. The controller includes a host interface. The controller is configured to: store a first logical address to physical address table in the working memory for use with storing randomly accessed data in the first set of memory cells and a second logical address to physical address table in the working memory for use with storing sequentially accessed data in the second set of memory cells, the first logical address to physical address table providing address translation at a lower level of granularity than the second logical address to physical address table; receive a request from a host via the host interface, the request including an indication of whether target data of the database to be accessed by the request is sequentially accessed data or randomly accessed data; access the target data from the first set of memory cells using the first logical address to physical address table in response to the request including an indication that the target data is randomly accessed data, and access the target data from the second set of memory cells using the second logical address to physical address table in response to the request including an indication that the target data is sequentially accessed data; and a compute engine that is behind the host interface and is in communication with the controller and the working memory, the compute engine being configured to perform data manipulation operations on sequentially accessed data stored in the second set of memory cells.

[00127] For purposes of this document, reference in the specification to “an embodiment,” “one embodiment,” “some embodiments,” or “another embodiment” may be used to describe different embodiments or the same embodiment.

[00128] For purposes of this document, a connection may be a direct connection or an indirect connection (e.g., via one or more other parts). In some cases, when an element is referred to as being connected or coupled to another element, the element may be directly connected to the other element or indirectly connected to the other element via intervening elements. When an element is referred to as being directly connected to another element, then there are no intervening elements between the element and the other element. Two devices are “in communication” if they are directly or indirectly connected so that they can communicate electronic signals between them.

[00129] For purposes of this document, the term “based on” may be read as “based at least in part on.”

[00130] For purposes of this document, without additional context, use of numerical terms such as a “first” object, a “second” object, and a “third” object may not imply an ordering of objects, but may instead be used for identification purposes to identify different objects.

[00131] For purposes of this document, the term “set” of objects may refer to a “set” of one or more of the objects.

[00132] The foregoing detailed description has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. The described embodiments were chosen in order to best explain the principles of the proposed technology and its practical application, to thereby enable others skilled in the art to best utilize it in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope be defined by the claims appended hereto.

CLAIMS

What is claimed is:

1. An apparatus, comprising:
 - a memory system, comprising:
 - non-volatile memory;
 - a controller in communication with the non-volatile memory;
 - working memory connected to the controller and configured to store a coarse logical address to physical address table for use with storing sequentially accessed data related to database operations and a fine logical address to physical address table for use with storing randomly accessed data related to database operations; and
 - a compute engine positioned within the memory system and in communication with the controller,
 - wherein the compute engine is configured to:
 - receive code from a host, the code including instructions for performing one or more database operations on target data in the non-volatile memory and an indication of whether the target data is sequentially accessed data or randomly accessed data,
 - perform the one or more database operations based on the code, and the target data being accessed using the coarse logical address to physical address table in response to the indication denoting that the target data is sequentially accessed data, and
 - perform the one or more database operations based on the code, and the target data being accessed using the fine logical address to physical address table in response to the indication denoting that the target data is randomly accessed data.
2. The apparatus of claim 1, wherein:
 - the code received from the host is compiled code that includes the indication of whether the target data is sequentially accessed data or randomly accessed data.
 - the controller includes a host interface;
 - the compute engine is behind the host interface;
 - the host interface is configured to receive the compiled code from the host and forward the compiled code to the compute engine;

the compute engine is further configured to use the controller to perform address translation between logical and physical addresses for the target data based on the coarse logical address to physical address table and the fine logical address to physical address table in the working memory.

3. The apparatus of claims 1 or 2, wherein:

the non-volatile memory includes a first portion of non-volatile memory configured to store sequentially accessed data comprising leaf node data for one or more database data structures and a second portion of non-volatile memory configured to store randomly accessed data comprising branch node data for one or more database data structures;

the controller is further configured to use the coarse logical address to physical address table to perform address translation between logical and physical addresses for sequentially accessed data stored in the first portion of non-volatile memory; and

the controller is further configured to use the fine logical address to physical address table to perform address translation between logical and physical addresses for randomly accessed data stored in the second portion of non-volatile memory.

4. The apparatus of claims 1 or 2, wherein:

the non-volatile memory includes a first portion of non-volatile memory configured to store sequentially accessed data comprising leaf node data for one or more database data structures and a second portion of non-volatile memory configured to store randomly accessed data comprising branch node data for one or more database data structures;

the compute engine is further configured to use the coarse logical address to physical address table to perform address translation between logical and physical addresses for sequentially accessed data stored in the first portion of non-volatile memory; and

the compute engine is further configured to use the fine logical address to physical address table to perform address translation between logical and physical addresses for randomly accessed data stored in the second portion of non-volatile memory.

5. The apparatus of claim 4, wherein:

the compute engine is further configured to perform a random read of data written as sequentially accessed data using a logical block address and an offset.

6. The apparatus of claims 1 or 2, wherein:

the non-volatile memory includes a first portion of non-volatile memory configured to store leaf nodes of a database as sequentially accessed data and a second portion of non-volatile memory configured to store branch nodes of the database randomly accessed data;

the compute engine is further configured to use the coarse logical address to physical address table to perform address translation between logical and physical addresses for the leaf nodes stored in the first portion of non-volatile memory; and

the compute engine is further configured to use the fine logical address to physical address table to perform address translation between logical and physical addresses for the branch nodes stored in the second portion of non-volatile memory.

7. The apparatus of claim 1, wherein:

the working memory is volatile memory;

the compute engine is further configured to use the volatile working memory to store working data for the compute engine while performing the one or more database operations;

the non-volatile memory includes a memory package that is separate from and connected to the controller; and

the memory package includes one or more non-volatile memory dies and the compute engine.

8. The apparatus of claim 7, wherein:

the compute engine is positioned on a first memory die of the one of the non-volatile memory dies;

the code received from the host is compiled code that is received by the compute engine on the first memory die; and

the compute engine is configured to execute the compiled code to access the target data on the first memory die based on

the coarse logical address to physical address table in the working memory for sequentially accessed data, and

the fine logical address to physical address table in the working memory for data randomly accessed data.

9. The apparatus of any of claims 1-8, wherein:

the controller includes a host interface, the compute engine is behind the host interface.

the compute engine is further configured to use the working memory to store working data for the compute engine while performing the one or more database operations;

the code received from the host is compiled code that includes the indication of whether the target data is sequentially accessed data or randomly accessed data;

the non-volatile memory includes a first portion of non-volatile memory configured to store sequentially accessed data and a second portion of non-volatile memory configured to store randomly accessed data;

the compute engine is further configured to use the coarse logical address to physical address table to perform address translation between logical and physical addresses for sequentially accessed data stored in the first portion of non-volatile memory; and

the compute engine is further configured to use the fine logical address to physical address table to perform address translation between logical and physical addresses for randomly accessed data stored in the second portion of non-volatile memory.

10. The apparatus of any of claims 1-9, wherein:

the host comprises a host computer in communication with the memory system, the apparatus comprises the host computer, the host computer includes a Just-In-Time Compiler, the host computer is configured to:

determine whether the target data is sequentially accessed data or randomly accessed data,

create the code by using the Just-In-Time Compiler to compile instructions for the compute engine to perform the one or more database operations, and

embed the indication of whether the target data is sequentially accessed data or randomly accessed data.

11. A method, comprising:

maintaining, in a memory system, reduced size logical to physical (L2P) tables in volatile working memory by maintaining coarse L2P tables in the working memory for use with sequentially accessed data and maintaining fine L2P tables in the working memory for use with randomly accessed data;

receiving, by the memory system from a host, compiled code instructing a compute engine within the memory system to perform a set of one or more database operations on target data in the memory system, wherein the compiled code is embedded with an indication of whether the target data is randomly accessed data or sequentially accessed data;

using, by the compute engine inside the memory system, the compiled code to perform the set of one or more database operations on the target data using the working memory, including using space in the working memory made available due to reduced size L2P tables, to store working data for the set of one or more database operations; and

performing address translation for a logical address in the compiled code using coarse L2P tables or fine L2P tables based on the embedded indication.

12. The method of claim 11, wherein:

the embedded indication of whether the target data is sequentially or random accessed data is inferred from a type of the one or more database operations specified by the code.

13. The method of claims 11 or 12, wherein:

the performing address translation comprises translating logical block addresses to physical block addresses for sequentially accessed data and comprises translating logical page addresses to physical page addresses for randomly accessed data.

14. The method of claims 11, 12 or 13, wherein:

the using the compiled code to perform the set of one or more database operations on the target data comprises accessing a first subset of the target data in a first set of blocks of non-volatile memory dedicated to storing sequential data and accessing a second subset of the target data in a second set of blocks of the non-volatile memory dedicated to storing random data, the first set of blocks are different than the second set of blocks.

15. The method of claims 11, 12 or 13, wherein:

the using the compiled code to perform the set of one or more database operations on the target data comprises accessing leaf node data of a database in a first set of blocks of non-volatile memory dedicated to storing sequential data and accessing branch node data of the database in a second set of blocks of the non-volatile memory dedicated to storing random data, the first set of blocks are different than the second set of blocks.

Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8




Figure 9




Figure 10




Figure 11




Figure 12




Figure 13




Figure 13A




Figure 14




Figure 15





Figure 16






Figure 17

Figure 18




Figure 19




Figure 20




Figure 21




Figure 23




Figure 22




Figure 24





Figure 25



Fine-Grained (Page Based) Mapping in a Persistent Memory that Has 16 Different Map Entries (16 physical pages or 2 block in total, 2 sub blocks/block, 4 pages/sub-block)

Direct Mapping




FIG. 26

Logical Sub-Block Address	Physical Sub-Block Address
15	15
14	14
13	13
12	12
11	11
10	10
9	9
8	8
7	7
6	6
5	5
4	4
3	3
2	2
1	1
0	0

FIG. 26 (Cont.)

Hash Mapping




FIG. 26 (Cont.)

Logical Sub-Block Address	Physical Sub-Block Address
15	7
14	6
13	5
12	4
11	3
10	2
9	1
8	0
7	15
6	14
5	13
4	12
3	11
2	10
1	9
0	8

FIG. 26 (Cont.)

Coarse-Grained Mapping (Sub-block Based) **that Has 4 Different Map Entries** **Direct Mapping**




FIG. 27





FIG. 27 (Cont.)

Figure 28

Coarse-Grained Mapping (Erase Block Based) that Has 2 Different Map Entries



INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/020152

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F12/02 G06F3/06
ADD.

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

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)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/143029 A1 (SIVASANKARAN VIJAY [IN] ET AL) 21 May 2015 (2015-05-21) paragraph [0016] - paragraph [0023]; claims 9,10; figures 1-6 paragraph [0031] - paragraph [0037] -----	1-15
X	US 2008/183955 A1 (YANG CHENG-CHIH [TW] ET AL) 31 July 2008 (2008-07-31) page 11 - page 18; claims 1-6; figures 1-7 paragraph [0030] - paragraph [0042] -----	1,11
A	US 2015/143028 A1 (JUNG HOE SEUNG [KR] ET AL) 21 May 2015 (2015-05-21) paragraph [0055] - paragraph [0070] paragraph [0104] - paragraph [0112] -----	1-15



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

31 May 2018

Date of mailing of the international search report

11/06/2018

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Toader, Elena Lidia

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/020152

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015143029	A1	21-05-2015	NONE

US 2008183955	A1	31-07-2008	TW 200832440 A 01-08-2008
		US 2008183955 A1	31-07-2008

US 2015143028	A1	21-05-2015	CN 104657279 A 27-05-2015
		KR 20150057068 A	28-05-2015
		TW 201520764 A	01-06-2015
		US 2015143028 A1	21-05-2015
