

# UK Patent Application GB 2519629

(43) Date of A Publication

A

29.04.2015

(21) Application No: 1414545.2  
(22) Date of Filing: 15.08.2014  
(30) Priority Data:  
(31) 61866672 (32) 16.08.2013 (33) US  
(31) 14055336 (32) 16.10.2013 (33) US

(51) INT CL:  
G06F 12/02 (2006.01) G11C 16/08 (2006.01)  
G11C 29/04 (2006.01)  
(56) Documents Cited:  
US 8843711 B US 6374341 B  
US 20140223089 A US 20100017578 A  
US 20070143569 A

(71) Applicant(s):  
LSI Corporation  
1320 Ridder Park Drive, San Jose 95131, California,  
United States of America

(58) Field of Search:  
INT CL G06F, G11C  
Other: EPODOC, WPI, TXTE.

(72) Inventor(s):  
Earl T Cohen

(74) Agent and/or Address for Service:  
Williams Powell  
11 Staple Inn, LONDON, WC1V 7QH, United Kingdom

(54) Title of the Invention: **Variable-size flash translation layer**  
Abstract Title: **Reading variable size flash memory transition layer storing compressed data.**

(57) Disclosed is a method of using a variable-size flash transition layer in a non-volatile memory. The method starts when the non-volatile memory receives a read request to read data corresponding to a logical block address of the non-volatile memory. Next, a particular entry of a map, associated with the logical block address, is read, to obtain (i) a physical address of a particular page of the non-volatile memory, (ii) an offset in the particular page to stored compressed data and (iii) a length of the compressed data. Then the memory converts the offset and the length to (i) an address of a given read unit in the particular page and (ii) a number of the read units to be read. Finally, the data is read from the particular page starting from the given read unit to at most the number of the read units given to be read in (ii). The offset and length granularity are finer than one read unit.

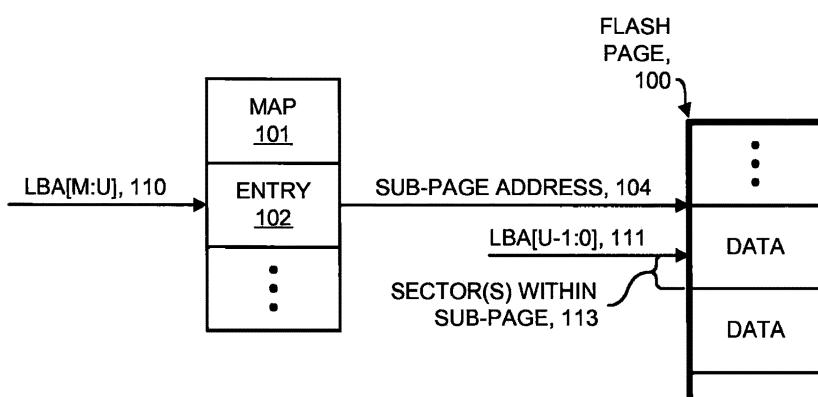
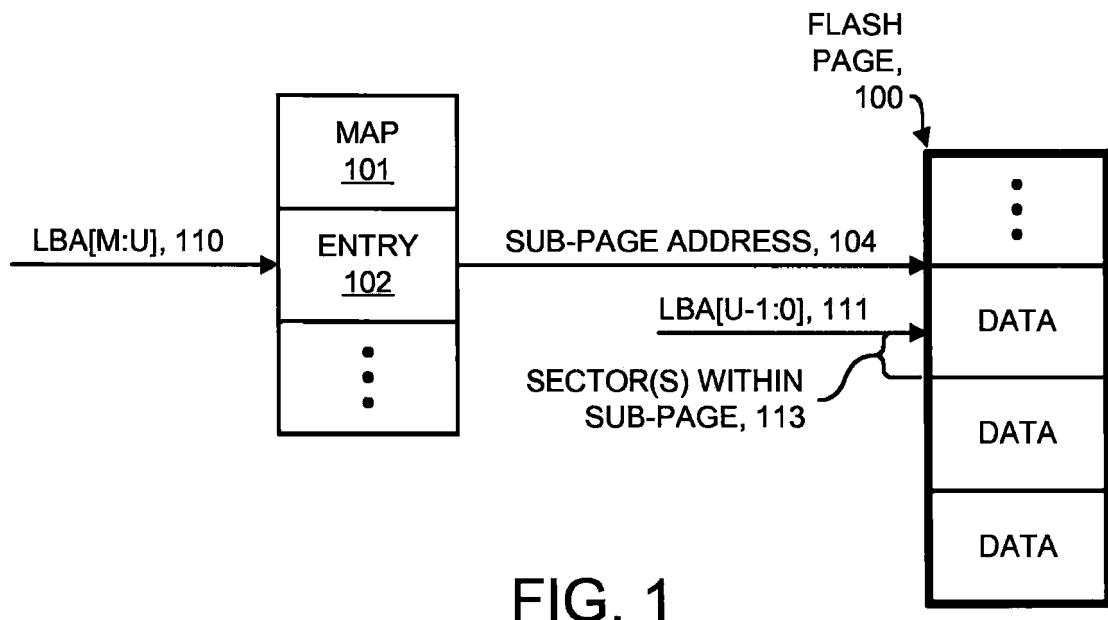
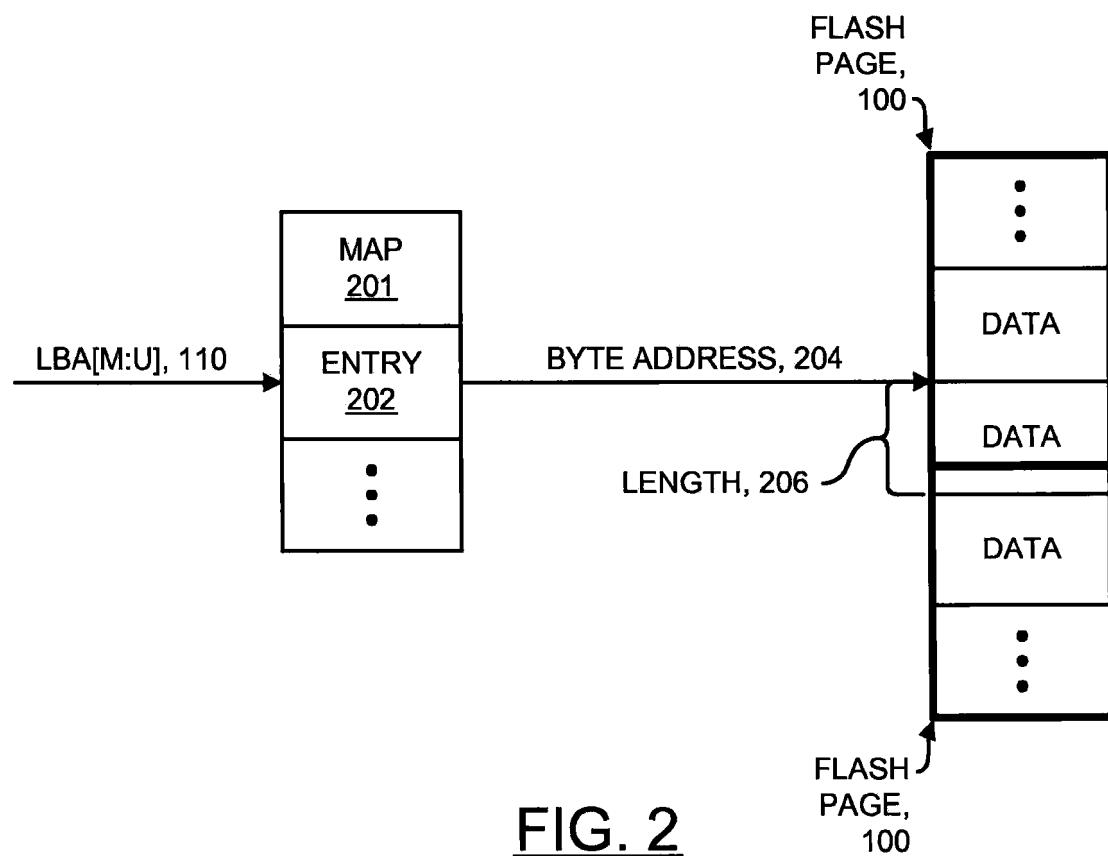
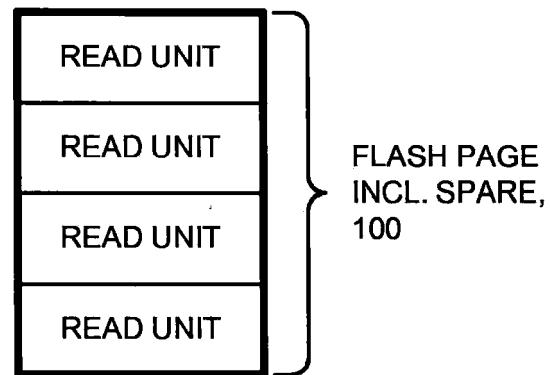
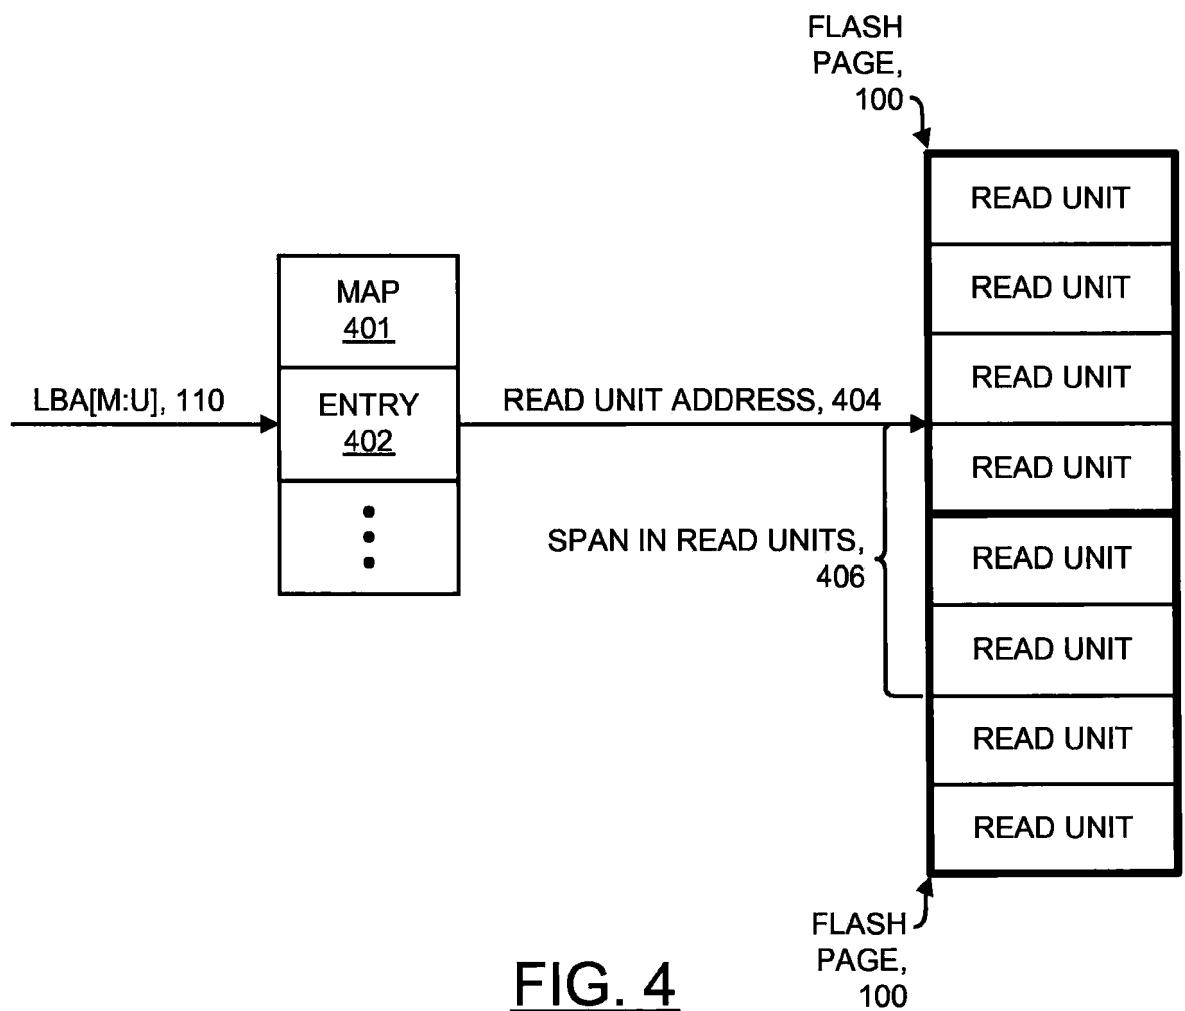


FIG. 1

GB 2519629 A

**FIG. 1****FIG. 2**

**FIG. 3****FIG. 4**

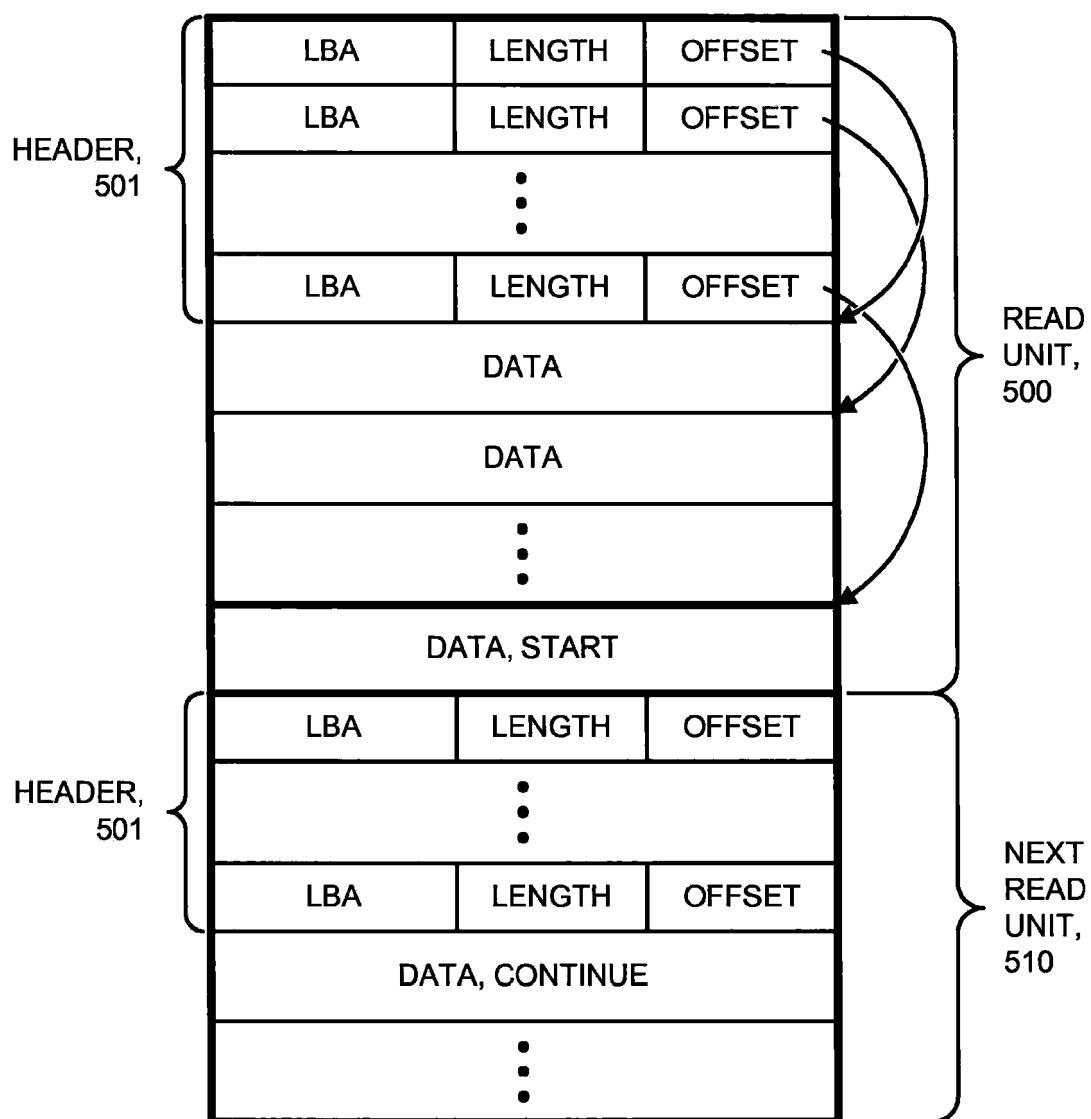
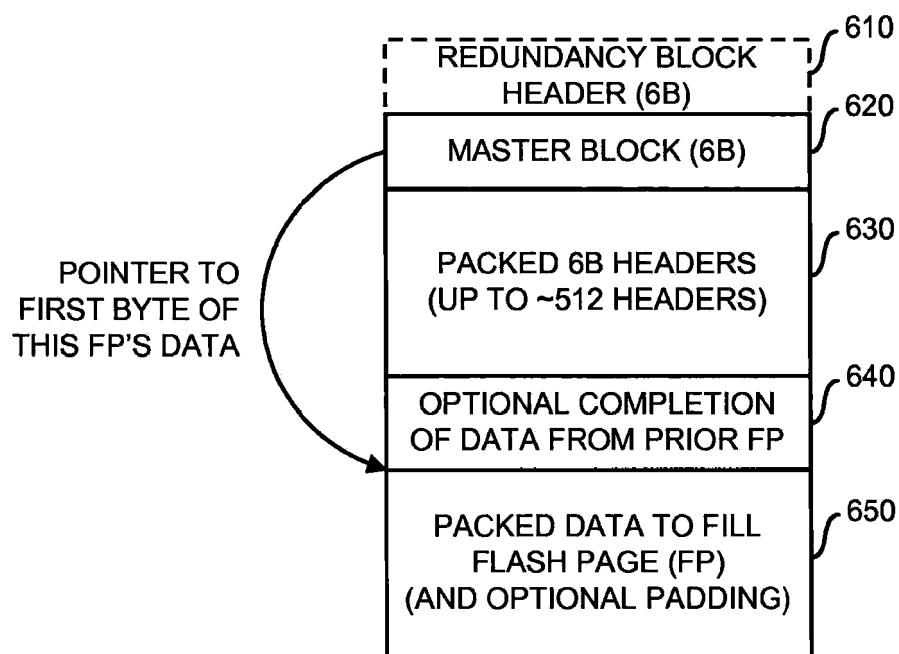
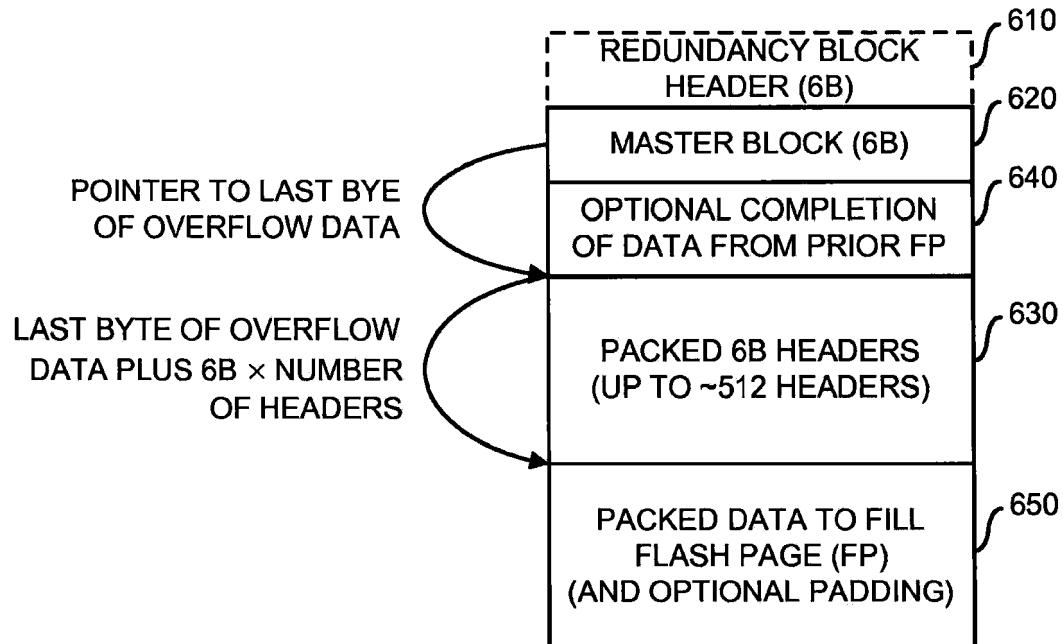
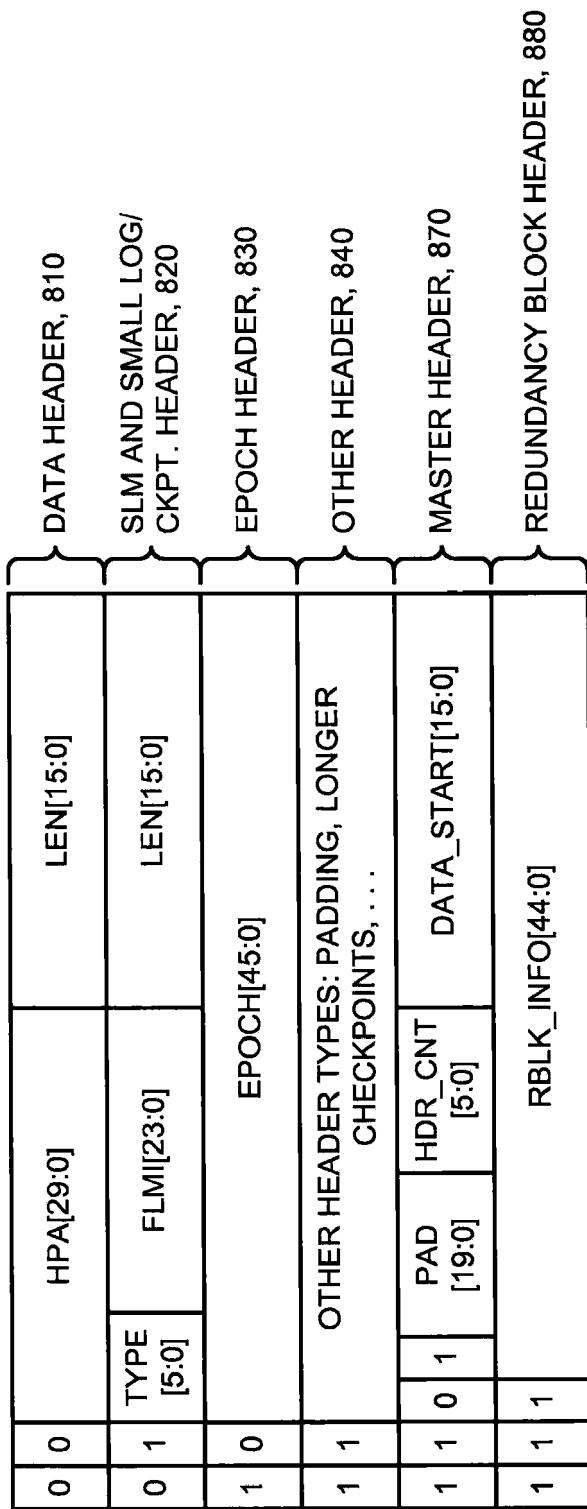
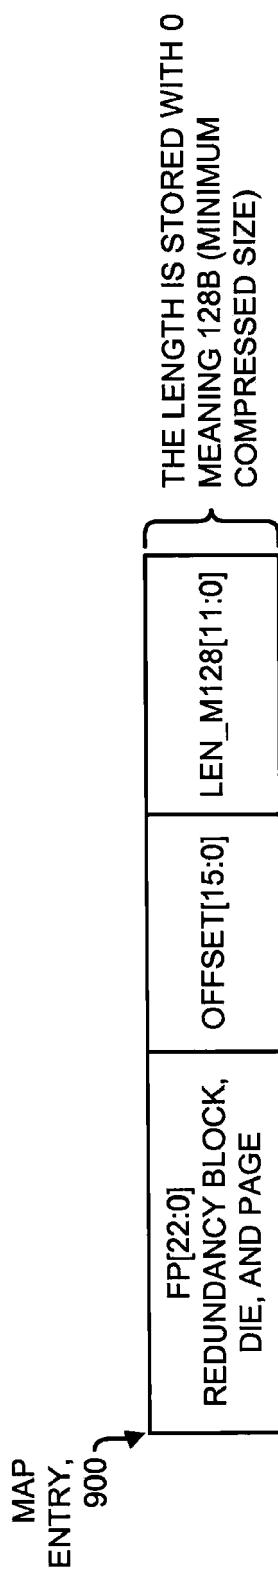
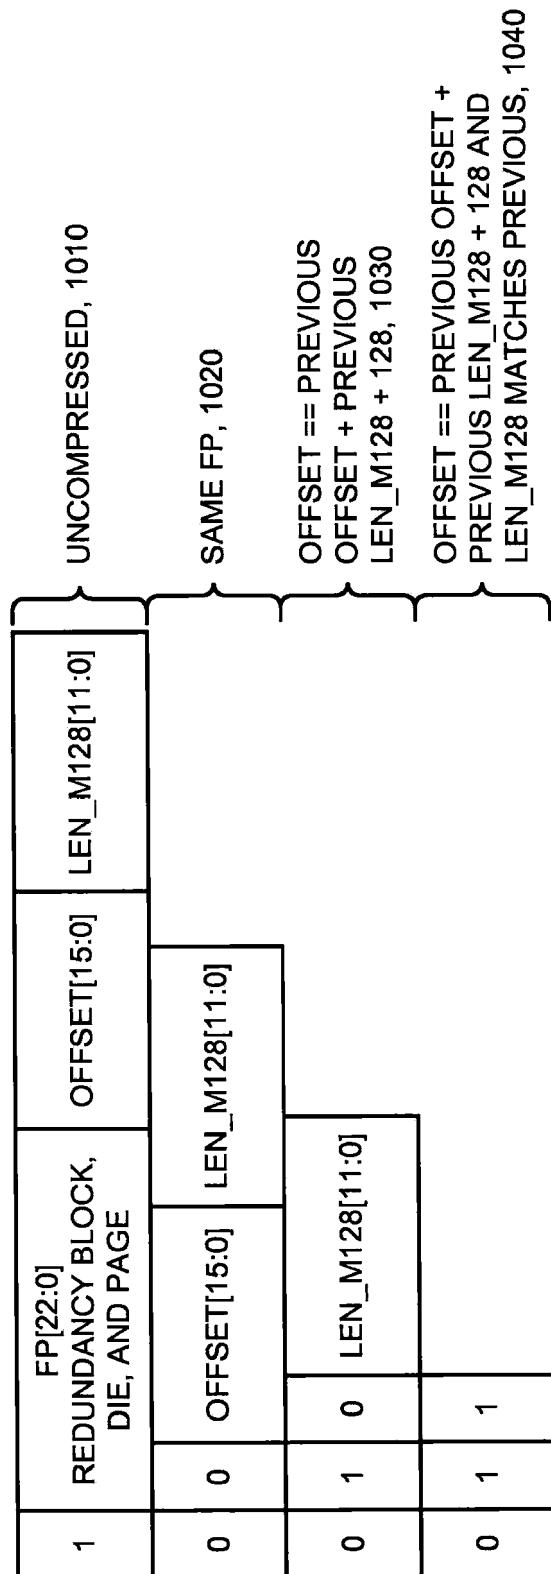


FIG. 5

FIG. 6FIG. 7

FIG. 8FIG. 9

FIG. 10

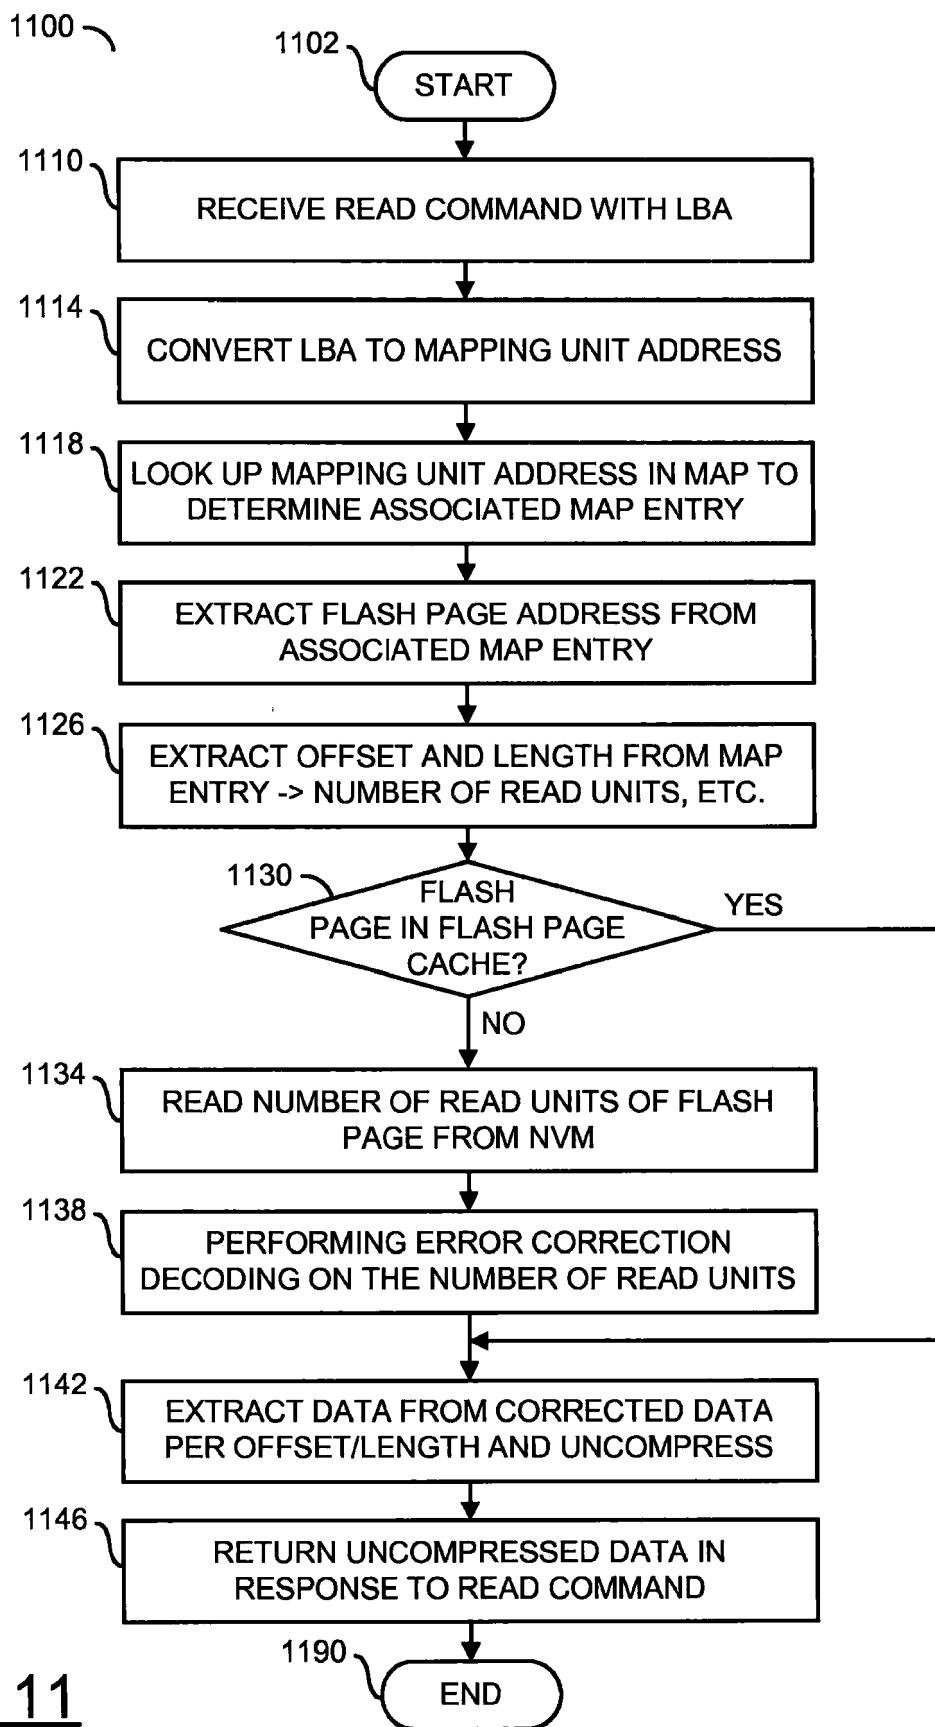
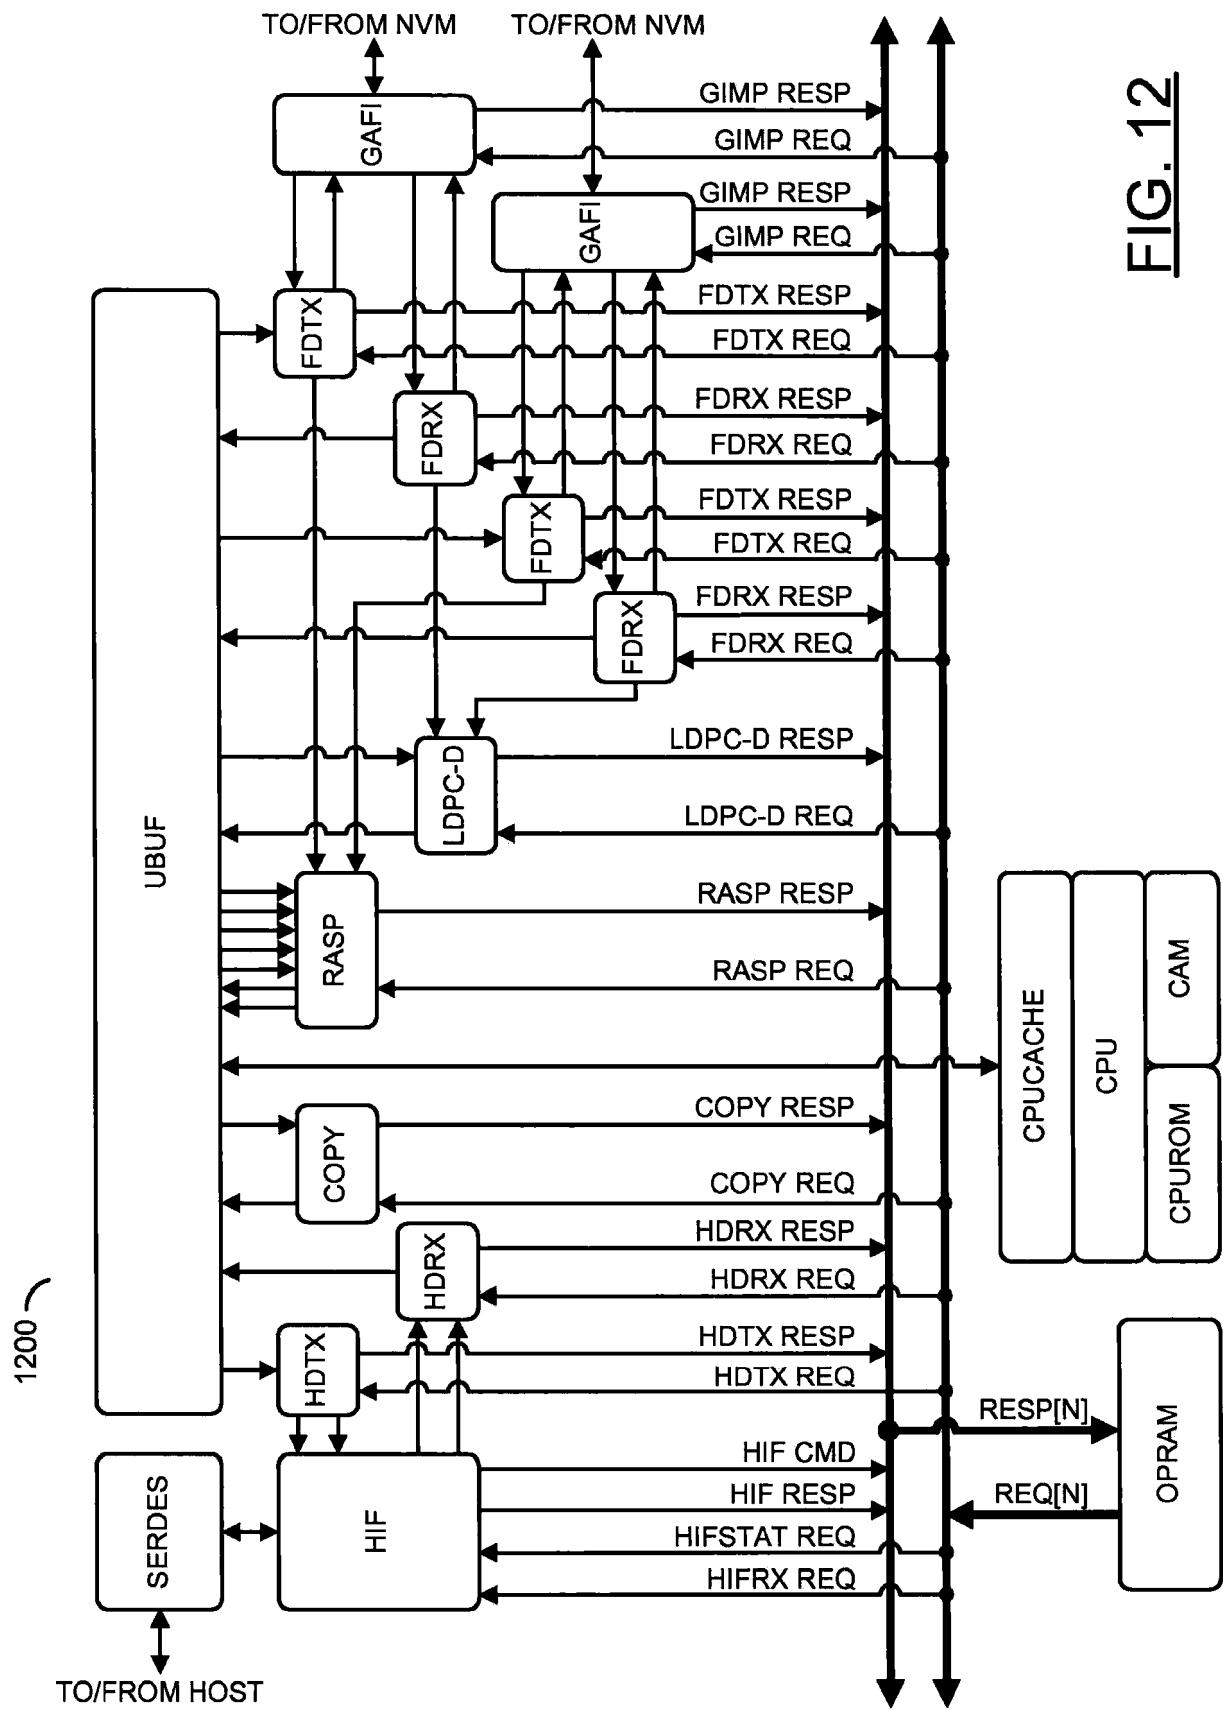


FIG. 11



**VARIABLE-SIZE FLASH TRANSLATION LAYER**

This application relates to U.S. Provisional Application No. 61/888,681, filed October 9, 2013, U.S. Provisional Application 5 No. 61/866,672, filed August 16, 2013, and U.S. Provisional Application No. 61/755,169, filed January 22, 2013, each of which are hereby incorporated by reference in their entirety.

This application relates to U.S. Serial No. 13/053,175, filed March 21, 2011, which relates to U.S. Provisional Application 10 No. 61/316,373, filed March 22, 2010, each of which are hereby incorporated by reference in their entirety.

This application also relates to International Application PCT/US2012/058583, with an International Filing Date of October 4, 2012, which claims the benefit of U.S. Provisional Application 15 No. 61/543,707, filed October 5, 2011, each of which are incorporated by reference in their entirety.

**Field of the Invention**

The invention relates to computing host and input/output 20 device technology generally and, more particularly, to a method

and/or apparatus for implementing a variable-size flash transition layer.

### **Background**

5       Conventional solid state drives store a fixed, integer number of host logical blocks in each page of nonvolatile memory. Storage efficiency issues arise when either a user data size or a usable size of each page of the nonvolatile memory is not fixed. Architectures for variable size flash transition  
10      layers in the solid state drives are hardware intense. Page headers are used to identify where the user data is stored among multiple read units within the pages of the solid state drive, and extracting the data involves first reading and parsing the page headers.

15

### **Summary**

According to an aspect of the present invention, there is provided a method for using a variable-size flash transition layer. Step (A) receives a read request to read data  
20      corresponding to a logical block address from a non-volatile memory. Step (B) reads a particular entry of a map to obtain (i) a physical address of a particular page of the nonvolatile memory, (ii) an offset in the particular page to compressed data previously stored and (iii) a length of the compressed  
25      data. The particular

entry is associated with the logical block address. Step (C) converts the offset and the length to (i) an address of a given read unit in the particular page and (ii) a number of the read units to be read. Step (D) reads from the particular 5 page at most the number of the read units starting from the given read unit. An offset and length granularity are finer than one read unit.

**Brief Description of the Figures**

10 Embodiments of the invention will now be described, by way of example only, referring to the following detailed description and the appended claims and drawings in which:

15 FIG. 1 is an illustration of selected details of an embodiment of mapping of a logical block address to fixed-sized regions within a flash page;

FIG. 2 is an illustration of selected details of an embodiment of mapping of a logical block address to a variable sized region that optionally spans flash pages;

20 FIG. 3 is an illustration of an embodiment of a flash page comprising an integer number of read units;

FIG. 4 is an illustration of selected details of an embodiment of mapping of a logical block address to a variable sized region spanning one or more read units;

FIG. 5 is an illustration of selected details of an embodiment of a read unit comprising headers and data;

FIG. 6 is an illustration of selected details of an embodiment of a flash page comprising headers and data;

5 FIG. 7 is an illustration of selected details of an embodiment of a flash page comprising headers and data in accordance with an embodiment of the invention;

FIG. 8 is an illustration of selected details of an embodiment of various types of headers;

10 FIG. 9 is an illustration of selected details of an embodiment of a map entry;

FIG. 10 is an illustration of selected details of an embodiment of various compressed map entries;

15 FIG. 11 is a flow diagram of reading a non-volatile memory; and

FIG. 12 is an illustration of selected details of an embodiment of a solid-state drive controller.

#### Detailed Description of the Embodiments

20 Selected embodiments seek to provide a variable-size flash transition layer that may (i) support a wide range of data sizes (ii) create headers with a tiling process,

(iii) parse headers with an un-tiling process, (iv) place all headers at a start of each page, (v) place all data after the headers in each page, (vi) use offsets and headers not aligned to read unit boundaries and/or (vii) be 5 implemented as one or more integrated circuits and/or associated firmware.

Embodiments of the present invention may be implemented in numerous ways, for example, as a process, an article of manufacture, an apparatus, a system, a composition of matter, and a computer readable medium such as a computer readable storage medium (e.g., media in an optical and/or magnetic mass storage device such as a disk, an integrated circuit having nonvolatile storage such as flash storage), or a 10 computer network wherein program instructions are sent over optical or electronic communication links. The detailed description provides an exposition of one or more embodiments of the invention that enable improvements in cost, profitability, performance, efficiency, and utility of use in 15 the field identified above. The detailed description includes an introduction to facilitate understanding of the remainder of the detailed description. The introduction includes 20 example embodiments of one or more of systems, methods, articles of manufacture, and computer readable media in 25 accordance with concepts described herein. As is

discussed in more detail, the invention encompasses all possible modifications and variations within the scope of the issued claims.

Flash translation layers (e.g., FTLs) map logical block addresses (e.g., LBAs) in a logical block address space (such as used by a host to perform input/output operations to an input/output device) to physical locations in a nonvolatile memory (e.g., NVM), such as a NAND flash nonvolatile memory. The mapping operates on aligned units of one or more logical blocks, termed a mapping unit, such that each mapping unit has a corresponding physical location where data of the mapping unit is stored (including the possibility of a NULL physical location if the mapping unit has never been written or is trimmed). For example, with 4 kilobyte (e.g., KB) mapping units, eight contiguous (and typically eight-sector aligned) Serial Advanced Technology Attachment (e.g., SATA) 512 byte sectors are mapped as a single unit. Generally, a translation table, such as a map, has an entry per mapping unit to store a respective translation from the logical block address associated with the mapping unit to a physical address in the nonvolatile memory and/or other control information.

Nonvolatile memories, such as NAND flash, provide a writeable (or programmable) unit called a flash page. A flash page comprises a number of user (non-error correction code) data bytes

and an amount of spare space for meta-data and error correction coding (e.g., ECC), and is generally a smallest writable unit of the nonvolatile memory. Typical flash page sizes are 8KB or 16KB or 32KB of user data, whereas typical mapping unit sizes are 4KB or 5 8KB. (While the term "user" data is used with respect to flash pages, some flash pages store "system" data such as map data and/or checkpoint data. User data is intended to refer in general to non-ECC portions of a flash page.) The flash pages are organized into blocks, typically 128, 256, or 512 flash pages per block. A 10 block is the minimum-sized unit that can be erased, and a flash page is erased before the page is able to be (re)written.

Referring to FIG. 1, an illustration of selected details of an embodiment of mapping of a logical block address to fixed-sized regions within a flash page is shown. Traditional flash 15 transition layers assume a number of user data bytes in a flash page (e.g., Flash Page 100) is a power-of-two (and/or a multiple of the sector size), and divide the flash page into an integer number of mapping units (each shown as DATA in FIG. 1). For example, with 16KB of user data per flash page and 4KB mapping units, each flash 20 page contains four mapping units, and the flash transition layer maps an address of each mapping unit (e.g., LBA[M:U] 110) to a respective flash page and one of the four mapping units within the

respective flash page. That is, each map entry contains respective fields such as:

flash\_page\_address[n-1:0], mapping\_unit\_within\_flash\_page[k-1:0]  
where the flash\_page\_address refers to a unique flash page in the  
5 nonvolatile memory, and the mapping\_unit\_within\_flash\_page refers  
to one of  $2^k$  mapping-unit-size portions of each flash page (k being  
fixed for the entire nonvolatile memory). Sub-page address 104 is  
a combination of flash\_page\_address and mapping\_unit\_within\_flash\_  
page. For sector-based addressing, lower-order bits of the logical  
10 block address (e.g., LBA[U-1:0] 111) specify a sub-portion, such as  
a number of sectors (e.g., sector(s) within sub-page 113) within  
the mapping unit.

Referring to FIG. 2, an illustration of selected details  
of an embodiment of mapping of a logical block address to a  
15 variable-sized region that optionally spans flash pages is shown.  
Variable-size flash translation layers (e.g., VFTLs) conceptually  
map an address of a mapping unit (e.g., LBA[M:U] 110) to a  
variable-sized region of one or more flash pages (because, for  
example, data of the mapping unit is compressed prior to being  
20 stored in flash, and/or because, in another example, the mapping  
units are written by the host as variable-sized pieces, such as for  
an object store). However, providing a complete byte address 204

and byte data length 206 in each map entry makes the map entries large when compared with traditional flash transition layers.

Variable-size flash translation layers are used in some solid-state disks (e.g., SSDs). The solid-state disk systems have

5 generally been designed for higher-end client and/or enterprise applications where random access performance constraints are a driving factor in the overall system design. To configure a variable-size flash translation layer for a low-end and/or mobile environment, changes may be implemented to configure for sequential

10 performance as the driving factor. Embodiments of the invention provide one or more ways of organizing user data and VFTL meta-data to enable less expensive and more efficient low-end and mobile nonvolatile memory systems where sequential read performance is a dominant constraint.

15 Referring to FIG. 3, an illustration of an embodiment of a flash page comprising an integer number of read units is shown. In some embodiments, variable-size flash translation layers perform the mapping from addresses of mapping units to physical addresses by mapping to an Epage (e.g., ECC page) address, also termed a 20 "read unit" address. An Epage (or read unit) is the minimum amount of data that can be read from the nonvolatile memory and corrected by the error correction code used to protect contents of the

nonvolatile memory. That is, each read unit contains an amount of data and corresponding ECC check bytes that protect that data. In some embodiments, a flash page (such as flash page 100), or in other embodiments, a group of flash pages treated as a unit for 5 purposes of writing, is divided into an integer number of read units, as illustrated in FIG. 3.

In various embodiments, the number of read units per flash page is allowed to vary. For example, some portions of the nonvolatile memory use stronger error correction code than others 10 (using more bytes in the flash page for error correction coding), and have fewer read units and/or less usable data per read unit. In another example, the number of read units per flash page varies as the nonvolatile memory is used, since the program/erase cycles tend to weaken the nonvolatile memory, resulting in stronger error 15 correction codes as the nonvolatile memory is used (worn) more.

According to various embodiments, the error correction code used is one or more of: a Reed-Solomon (e.g., RS) code; a Bose Chaudhuri Hocquenghem (e.g., BCH) code; a turbo code; a low-density 20 parity-check (e.g., LDPC) code; a polar code; a non-binary code; a redundant array of inexpensive/independent disks (e.g., RAID) code; an erasure code; any other error correction code; any combination of the foregoing including compositions, concatenations, and

interleaving. Typical codeword sizes range from 512 bytes (plus ECC bytes) to 2176 bytes (plus ECC bytes). Typical numbers of ECC bytes range from only a few bytes to several hundred bytes.

Referring to FIG. 4, an illustration of selected details 5 of an embodiment of mapping a logical block address to a variable-sized region spanning one or more read units is shown. In some embodiments, VFTL mapping maps the address of a variable-sized (e.g., compressed) mapping unit (e.g., LBA[M:U] 110) to a number of read units, represented in each entry of the map as a read unit 10 address 404 and a span (a number of read units) 406. The read units referenced by one of the map entries are in one or more (logically and/or physically) sequential flash pages, for example, the number of read units optionally and/or selectively cross a flash page boundary. An entry of the map alone is generally not 15 sufficient to locate the associated data (as the entry only references the read units, and not a location of the data within the read units), and further information (such as headers) within the referenced read units is used to precisely locate the associated data.

20 In some embodiments, data is written into flash pages in a manner that is striped across multiple dies of the nonvolatile memory. Striping write data across multiple dies advantageously

enables greater write bandwidth by only writing a flash page into a given die once per stripe. A stripe of blocks across multiple dies is termed a redundancy block, because in further embodiments and/or usage scenarios, RAID-like redundancy is added on a 5 redundancy block basis using, for example, one redundant die. In various embodiments, some blocks of the nonvolatile memory are defective and are skipped on writing, so that the striping occasionally has "holes" where one of the die is skipped (rather than writing into flash pages of a bad block). In such 10 embodiments, "sequential" flash pages are sequential in a logical order determined by an order in which the flash pages are written.

Referring to FIG. 5, an illustration of selected details of an embodiment of a read unit comprising headers and data is shown. In various embodiments, the mapping illustrated in FIG. 4 15 engenders a criterion to locate the variable-sized data within the read units. As illustrated in FIG. 5, each read unit (e.g., read units 500 and 510) has a set of headers 501, and the headers are written, typically by hardware, as the variable-sized data is "tiled" (e.g., densely packed without wasted space) into one or 20 more read units. The headers are interpreted, typically by other hardware, to extract the variable-sized data when the nonvolatile memory is read. Variable-sized data is located by a respective

offset and length in one of the headers having a matching logical block address, and data optionally and/or selectively spans read units (such as illustrated by the variable-sized data illustrated by "DATA, START" and "DATA, CONTINUE").

5                   In various embodiments, the headers are also used as part of recycling (e.g., garbage collection) - including the logical block address (or equivalently, the mapping unit address) in the headers both enables finding the variable-sized data within a read unit, and provides a way to determine when a particular one of the  
10 read units is read, if the variable-sized data within is still valid or has been overwritten (by looking up the logical block address in the map and determining if the map still references a physical address of the particular read unit, or has been updated to reference another one of the read units).

15                   In some embodiments, dedicated hardware to extract data from the read units based on the logical block addresses is implemented to operate with high efficiency for random reads. The dedicated hardware parses headers within one or more read units to find the one of the headers with a given logical block address, and  
20 then uses the respective length and offset to extract the associated variable-sized data. However, a hardware-based solution is costly (in silicon area and power). For a low-end and/or mobile

environment where sequential performance is more important than random, changes are implemented to the variable-size flash translation layer to reduce silicon area, save power, and achieve high sequential throughput rates.

5           In some embodiments, a sequential-read-optimized variable-size flash translation layer (e.g., SRO-VFTL) tiles data into flash pages (or, in some embodiments, a group of flash pages treated as a unit for purposes of writing) without any gaps for headers within the data - all the headers are grouped in one 10 portion of the flash page. In further embodiments, the headers are not used dynamically to access data (as in some variable-size flash translation layers), but are only used for recycling and recovery. Instead, entries of the map comprise complete information used to find variable-sized (e.g., compressed) data within the flash pages. 15 Separating headers and data into different portions of the flash page leads to read units that comprise only headers, read units that comprise a mixture of headers and data (but only one such read unit per flash page), and read units that comprise only data.

20           While being configured for sequential read throughput at low cost, a sequential-read-optimized-variable-size flash translation layer is able to perform comparatively well on other metrics such as random read input/output operations per second

(e.g., IOPs), random write input/output operations per second, and sequential write throughput. However, removal of hardware assists for functions such as VFTL-style data tiling with headers in each read unit places a larger burden on a control processor.

5 Referring to FIG. 6, an illustration of a first embodiment of an SRO-VFTL flash page is shown. Referring to FIG. 7, an illustration of a second embodiment of an SRO-VFTL flash page is shown in accordance with an embodiment of the invention. A difference between the embodiments of FIG. 6 and FIG. 7 is whether 10 the continuation data from a previous flash page 640 is before or after the headers. Various embodiments and arrangements of the data within the flash page are contemplated.

According to various embodiments, a flash page comprises one or more of:

15 - Headers, including a master header 610, optionally and/or selectively a redundancy block header 620 (e.g., a header added in the first page of each block in a redundancy block), and zero or more additional packed headers 630. Every flash page has at least a count of the number of followings headers and a pointer to where 20 data (associated with the headers) start in the flash page. In some embodiments, the headers may be byte-aligned, but are only 6 bytes (e.g., B) each. The headers may include, but are not limited

to, data headers, epoch headers and padding. Data headers utilize a mapping unit address and a length. The offset is implied because all data is consecutively packed.

5 - Optionally and/or selectively continuation data from a previous flash page (a portion of variable-sized data of a mapping unit) 640.

10 - Packed (e.g., optionally and/or selectively compressed) data of one or more mapping units 650 to fill the flash page, the last of which optionally and/or selectively continues in a subsequent flash page.

15 - Optional padding at the end of the flash page (included in 650). In various embodiments, the data is byte-packed (e.g., no holes), though possibly padded at the end of the flash page if highly compressed (e.g., too many headers). Padding is, for example, used if: (i) the last variable-sized piece of data added to the flash page left fewer unused bytes than a size of a header (so a new header could not be added to begin another variable-sized piece of data) and (ii) optionally and/or selectively, a specified number of headers per flash page is exceeded (so the number of 20 mapping units stored in the flash pages is limited by the specified number of headers and not by a size of data of the mapping units).

In some embodiments, recovery and/or recycling (e.g., garbage collection) with a sequential-read-optimized-variable-size flash translation layer is advantageously enabled to read and/or error correct and/or examine only a header portion of each of the 5 flash pages, and not every read unit as in a non-sequential-read-optimized-variable-size flash translation layer. If recycling determines that data of a flash page may be rewritten, that data may also be read and may also be error corrected. In some embodiments, an entire flash page is read for recycling, but only 10 the header portion is error corrected until a determination is made that some data in the flash page should be recycled.

In various embodiments, a number of headers per flash page is limited to bound a number of read units per flash page that may be read to ensure all the headers have been read from the 15 nonvolatile memory. In the embodiment of FIG. 6, only a number of read units sufficient to contain the maximum number of headers are read. In the embodiment of FIG. 7, an additional number of read units are read to account for a largest size of data completing from a previous flash page (e.g., continuation data 640). The 20 embodiment of FIG. 7, however, enables a number of reads units to access the completion of data from a previous flash page to be determined from the associated map entry, as the number of bytes in

the completion of data is determinable based on the respective offset and length of the associated map entry, and the number of bytes of user (non-error correction code) data in the previous flash page. Further, the only headers prior to the completion of 5 data are the optional redundancy block header (only present in known flash pages, such as the first page in each block) and the master header (always present in each flash page). In the embodiment of FIG. 6, to read the completion of data without having to access nonvolatile memory twice, the maximum number of headers 10 are assumed to be present (or the entire flash page is read).

In some embodiments, the sequential-read-optimized-variable-size flash translation layer uses a single-level map having a plurality of map entries. In other embodiments, the sequential-read-optimized-variable-size flash translation layer 15 uses a multi-level map, such as a two-level map having a first-level map (e.g., FLM) pointing to second-level map (e.g., SLM) pages, where each of the second-level map pages comprises a plurality of leaf-level map entries. In further embodiments, the multi-level map has more than two levels, such as three levels. In 20 some embodiments and/or usage scenarios, use of a multi-level map enables only a relevant (e.g., in use) portion of the map to be stored (e.g., cached) in local (e.g., on-chip) memory, reducing a

cost of maintaining the map. For example, if typical usage patterns have 1 gigabyte (e.g., GB) of the logical block address space active at any point in time, then only a portion of the map sufficient to access the active 1GB portion of the logical block address space is locally stored for fast access versus being stored in the nonvolatile memory. References outside of the active portion of the logical block address space fetch requested portions of one or more levels of the multi-level map from the nonvolatile memory, optionally and/or selectively replacing other locally-stored portions of the map.

Each of the leaf-level map entries is associated with (corresponds to) an address of one of a plurality of mapping units. A logical block address is converted to a mapping unit address, such as by removing zero or more least-significant bits (e.g., 15 LSBs) of the logical block address and/or adding a constant to the logical block address for alignment purposes, and the mapping unit address is looked up in the map to determine a corresponding entry of the map.

Referring to FIG. 8, an illustration of details of an embodiment of various types of headers is shown. In the example of FIG. 8, the headers have been formatted to fit in six bytes each. According to various embodiments, the various types of headers are

one or more of: all of a same size; optionally and/or selectively of different sizes; each comprises a respective field specifying a size of the header; vary in size in different flash pages; and any combination of the foregoing.

5                   According to various embodiments, the headers in the flash page comprise one or more of:

                  - Data headers (810) indicating information associated with a variable-sized data portion. In some embodiments, data associated with a data header starts in a same flash page as the data header appears. In further embodiments and/or usage scenarios, if a flash page only has remaining space for a data header, all of the associated data starts in a subsequent flash page.

10                   - Map headers, such as second-level map (e.g., SLM) headers (820). The second-level map headers comprise a first-level map index (e.g., FLMI) to indicate (such as for second-level map recycling and/or recovery) which second-level map page is being stored.

15                   Log/Checkpoint headers (820). Log/Checkpoint headers indicate data used for recycling, recovery, error handling, debugging, or other special conditions.

- Epoch headers (830) are used as part of recovery to associate data with corresponding map/checkpoint information. Typically, there is at least one Epoch header per flash page.

5 - Master headers (870) are used once per flash page to provide information as to a number of headers in the flash page and where non-header data starts within the flash page. Various techniques determine a start of non-header data, such as illustrated in the embodiments of FIG. 6 and FIG. 7.

10 - Redundancy block headers (880) are used in certain flash pages, such as the first flash page in each block in a redundancy block.

- Other types of headers (840), such as padding headers, checkpoint headers supporting larger lengths, etc.

15 In some embodiments, some headers comprise a TYPE field to provide multiple subtypes of the header. In various embodiments, some headers comprise a LEN (length) field containing a length of data associated with the header. In other embodiments, rather than a LEN field, some headers comprise an OFFSET (offset) field (not shown) containing an offset (within the flash page) to 20 the end of data associated with the header. (If the last one of the variable-sized pieces of data spans a flash page, the OFFSET is an offset within a subsequent flash page or a number of bytes

within the subsequent flash page.) Only one of a LEN field or an OFFSET field is generally implemented since with the variable-sized pieces of data packed with no wasted space, the starting location and ending location of each of the variable-sized pieces of data in 5 a flash page is implied by the starting location of the first variable-sized pieces of data in the flash page (e.g., immediately after the headers as in FIG. 7), and the list of LEN or OFFSET fields.

Referring to FIG. 9, an illustration of selected details 10 of an embodiment of a map entry 900 is shown. According to various embodiments, the entries of the map comprise one or more of:

- A physical flash page address,
- An offset within the flash page to a variable-sized data item,
- 15 - A length of the variable-sized data item, and
- Other control information.

In some embodiments, the length is encoded, for example by being offset such that a value of zero corresponds to a specified minimum length. In further embodiments, data that is 20 compressed to less than the specified minimum length is padded to be at least the specified minimum length in size.

In various embodiments, the SRO-VFTL map entries are larger than VFTL map entries since the SRO-VFTL map entries store a full offset and byte length of corresponding data. Accordingly, reducing a size of the map entries when stored in the nonvolatile 5 memory may be advantages. In typical use, data is often sequentially read and written, at least with some granularity and/or an average number of sequential mapping units greater than one, and a map entry compression format taking advantage of the sequential nature of writing is relatively inexpensive to implement 10 and produces a high map compression rate. Compression of map entries is further aided by sequentially-written data going into same flash pages, until a flash page boundary is crossed.

Referring to FIG. 10, an illustration of selected details of an embodiment of various compressed map entries is shown. The 15 various map entries include uncompressed (1010), having a same flash page address as a previous map entry (1020), having a same flash page address as the previous map entry and starting at an offset where the previous data ended (1030), and having a same flash page address as the previous map entry, starting at an offset 20 where the previous data ended, and having a same length as the previous map entry (1040).

In some embodiments having a multi-level map, a cache is maintained of lower-level (such as leaf-level) map pages. The cached map pages are in an uncompressed form, providing quick access by the processor. When map pages are moved (such as from 5 nonvolatile memory or dynamic random access memory (e.g., DRAM)) into the cache, the map pages are uncompressed. When the map pages are flushed from the cache (such as due to being modified), the map pages are compressed for storage (such as in the nonvolatile memory). According to various embodiments in which DRAM is used to 10 reduce latency by storing some or all of the map pages in the dynamic random access memory, the map pages in the dynamic random access memory are stored in one or more of: compressed form; uncompressed form; a selectively compressed or uncompressed form; and with an indirection table used to access the compressed 15 versions of the map pages in the dynamic random access memory

In some embodiments, host write data of a host write command is optionally and/or selectively compressed as the host write data arrives, and stored in a first-in-first-out (e.g., FIFO)-like fashion in a local (such as an on-chip) memory. For 20 example, in some embodiments the host write data is stored in a unified buffer (e.g., UBUF in FIG. 12), along with firmware data structures, flash statistics, portions of the map such as a cache

holding one or more pages of the map, read data from the nonvolatile memory including recycling read data, headers of data to be written to the nonvolatile memory, firmware code, and other uses. In other embodiments, one or more dedicated memories are 5 used for various local storage criteria of the solid-state drive.

With each mapping unit of data that arrives from the host, a control processor of the solid-state drive (e.g., central processing unit, CPU, in FIG. 12) is notified of one or more of: a respective mapping unit address, a respective local memory address, 10 and/or a respective length of each mapping-unit of variable-sized (e.g., compressed) host data. The control processor is enabled to determine an order of writing of flash pages, and a total number of non-ECC bytes available in each of the flash pages. According to the total number of non-ECC bytes available in a given one of the 15 flash pages, the control processor is enabled to determine an amount of headers and an amount of data placed in the given flash page. For example, the control processor accumulates headers for the given flash page (and tracks a number of bytes of headers used so far), and adds variable-sized data of mapping units and headers 20 to the given flash page one at a time until the given flash page is full. When the given flash page is full, a last portion of data of a final one of the mapping units added to the given flash page

possibly did not fit in the given flash page, and is used as a data completion portion of a subsequent one of the flash pages, reducing the total number of non-ECC bytes available in the subsequent flash page for new headers and data.

5           At a particular point in time, one or more flash pages are enabled to be filled with host write data, and one or more flash pages are enabled to be filled with recycled data. For example, at least two bands (e.g., FIFO-like series of redundancy blocks) may be filled, one band with "hot" data (e.g., fresh from the host) and the other band with "cold" data (e.g., recycled).  
10           Continuing the example, in various embodiments, the host write data is enabled to be directed into either the hot or the cold band, and recycled data is enabled to be directed into either the hot or the cold band.

15           The control processor is enabled to convert the series of respective mapping unit addresses, local memory addresses and lengths into one or more of: a series of headers to be written to a flash page as a header portion of the flash page; a first starting address and a first length of a sequential portion of the  
20           local memory to be written to the flash page as a user data portion of the flash page, the user data portion of the flash page comprising at least a portion of data of at least one mapping unit;

a second starting address and a second length of a sequential portion of the local memory to be written to a subsequent flash page as a user data completion portion of the subsequent flash page, the user data completion portion comprising a portion of data 5 of one mapping unit or being empty; a number of zero or more padding bytes to be written to the flash page, where the padding bytes are, for example, used if the user data completion portion is empty and the flash page is not full. Advantageously, the control processor is enabled to simply convert the series of respective 10 mapping unit addresses, local memory addresses and lengths into the series of headers by reformatting, and to generate a small number of direct memory access (e.g., DMA) commands to transfer the portions comprising the flash page (the series of headers, a completion portion of a previous flash page, the user data portion, 15 and any padding bytes) to the nonvolatile memory.

In various embodiments, compression of the host write data is optionally and/or selectively enabled. In a first example, information of the host write command selectively enables compression. In a second example, compression is selectively enabled as a function of a logical block address of the host write 20 command. In a third example, compression is selectively disabled if compression of the host write data did not reduce a size of the

host write data. If compression is not enabled, the host write data is stored uncompressed. According to various embodiments, entries of the map indicate if the corresponding data is compressed or uncompressed by one or more of: a respective bit in each entry of the map; and/or a value of the length stored in each map entry. For example, if mapping units are 4KB, a length of 4KB in a map entry indicates that associated data of the map entry is uncompressed, whereas a length less than 4KB indicates that the associated data is compressed.

10 In some embodiments, data is recycled by selecting a redundancy block to be recycled, reading flash pages of the redundancy block in an order in which the flash pages were written, processing solely read units that contain headers of the flash pages, looking up a logical block address (or equivalently a mapping unit address) of each header that is a data header in the map to see if the data is still valid, and if the data is still valid constructing appropriate new headers and DMA commands to assemble the data to be recycled as part of a new flash page. The new flash page is then written to the nonvolatile memory.

20 Referring to FIG. 11, an illustration of a flow diagram 1100 of reading a nonvolatile memory is shown. In contrast with the non-sequential-read-optimized variable-size flash translation

layer, headers within the read units (or within the flash pages) are not used to extract read data. Both the non-sequential-read-optimized variable-size flash translation layer and the sequential-read-optimized-variable-size flash translation layer advantageously are enabled to access variable-sized data, and to only access the read units that contain the desired read data.

In some embodiments, in response to receiving a read command from the host comprising a logical block address (step 1110), the control processor and/or various hardware units are enabled to perform one or more of:

- converting the logical block address to a mapping unit address (step 1114);

- looking up the mapping unit address in a map comprising a plurality of map entries to determine an associated one of the map entries (step 1118);

- extracting a respective flash page address of the associated map entry (step 1122) and determining if the associated flash page is in a flash page cache, or is read from the nonvolatile memory (step 1130);

- extracting a respective offset and length from the associated map entry, and according to the respective offset and length, determining (step 1126);

- a. number of a plurality of read units to access in the associated flash page,

                  - b. a read unit offset and total read unit length within the flash page of the accessed read units, and

5                  - c. a DMA command to extract and process (such as by uncompressed) data associated with the mapping unit address from decoded versions of the accessed read units;

10                 - in response to determining that the associated flash page is not in the flash page cache, reading the accessed read units of the associated flash page from nonvolatile memory (step 1134), and performing error correction decoding on the accessed read units (step 1138) to produce corrected data;

15                 - extracting the associated data within the corrected data according to the respective offset and length of the associated map entry, and uncompressing the extracted data (step 1142); and

                  - providing the uncompressed data to the host in response to the read command (step 1146).

Typically for random read, the number of read units to access in the associated flash page to read the associated data is less than all of the read units in the associated flash page. Further, as the associated data is variable-sized, a number of read units to access in the associated flash page for a first read

command referencing a first logical block address is different from a number of read units to access in the associated flash page for a second read command referencing a second logical block address, the second logical block address different from the first logical 5 block address. In some embodiments, solely the number of read units to access in the associated flash page are read from the associated flash page. That is, only the ones of the read units that contain a portion of the associated data are read in order to access the associated data.

10                 In some embodiments and/or usage scenarios, a particular one of the read units comprises at least a portion of data associated with a first logical block address, and at least a portion of data associated with a second different logical block address.

15                 Referring to FIG. 12, an illustration selected details of an embodiment of a solid-state drive controller 1200 is shown. In some embodiments, the solid-state drive controller 1200 is enabled to implement a sequential-read-optimized-variable-size flash translation layer. In various embodiments, the controller 1200 may 20 be implemented as one or more integrated circuits.

               As illustrated in FIG. 12, a host interface (e.g., HIF), via an input/output receiver such as a SerDes (e.g., serialization-

deserialization), receives commands, such as read and write commands, receives write data, and sends read data. Commands are sent to a CPU via a shared memory (e.g., OpRAM). The CPU interprets the commands and controls other portions of the SSD controller via the shared memory. For example, the CPU communicates DMA commands to and receives responses from various datapath transmit and receive units, such as the host receive datapath (e.g., HDRx) or the flash transmit datapath (e.g., FDTx), via the shared memory.

10 Write data from the host interface is transferred via a host receive datapath (e.g., HDRx) to a unified buffer (e.g., UBUF). In various embodiments, the host receive datapath includes logic to optionally and/or selectively compress and/or encrypt the host write data. The optionally and/or selectively compressed and/or encrypted host write data is then sent from the unified buffer to the nonvolatile memory via a flash transmit datapath (e.g., FDTx) and a generic flash interface (e.g., GAFI). In various embodiments, the flash transmit datapath includes logic to perform encryption and/or scrambling and/or error correction encoding. In response to host read commands, data is read from the nonvolatile memory via the generic flash interface (e.g., GAFI) and sent to the unified buffer via a flash receive datapath (e.g.,

FDRx). In various embodiments, the flash receive datapath incorporates error correction decoding and/or decryption and/or de-scrambling. In other embodiments, a separate error correction decoder (e.g., LDPC-D to implement LDPC codes) is enabled to 5 operate on "raw" data stored in the unified buffer by the flash receive datapath. Decoded read data in the unified buffer is then sent to the host Interface via a host transmit datapath (e.g., HDTx). In various embodiments, the host transmit datapath includes logic to optionally and/or selectively decrypt and/or decompress 10 decoded read data. In some embodiments, a RAID-like and soft-decision processing unit (e.g., RASP) is enabled to generate RAID-like redundancy to additionally protect host write data and/or system data stored in the nonvolatile memory, and/or to perform soft-decision processing operations for use with the LDPC-D.

15 According to various embodiments, any operations of the control processor are performed by any of one or more CPUs, by one or more hardware units, and/or by any combination of the foregoing. For example, for writing, conversion of the series of respective mapping unit addresses, local memory addresses and lengths into the 20 series of headers is aided by hardware supplying the series of respective mapping unit addresses, local memory addresses and

lengths in a format that is same as and/or similar to a format of the series of headers.

According to various embodiments, a solid-state drive controller coupled to a nonvolatile memory is enabled to use one or 5 more of: a traditional flash transition layer; a variable-sized flash transition layer; a sequential read optimized variable-sized flash transition layer; any combination of the foregoing in different physical portions of the nonvolatile memory; any combination of the foregoing in different logical portions of a 10 logical address space of the SSD controller; raw physical access to the nonvolatile memory; and any combination of the foregoing under control of a host coupled to the SSD controller.

According to various embodiments, host write data is optionally encrypted prior to being written to the nonvolatile 15 memory, and decrypted after being read from the nonvolatile memory. In further embodiments, encryption happens subsequent to compression of the host write data, and decryption happens prior to decompression of data being read to return to the host.

While example embodiments have used solid-state drives, 20 the techniques described herein are generally applicable to other input/output devices and/or data storage devices such as hard disk drives.

The following is a collection of example embodiments, including at least some explicitly enumerated example combinations (e.g., ECs), providing additional description of a variety of embodiment types in accordance with the concepts described herein; 5 the examples are not meant to be mutually exclusive, exhaustive, or restrictive; and the invention is not limited to these example embodiments but rather encompasses all possible modifications and variations within the scope of the issued claims and their equivalents.

10 A method EC1 comprising: receiving, at an input/output device and via a host to input/output device interface, a read request to read data corresponding to a logical block address of the read request from nonvolatile memory of the input/output device; and in response to receiving the read request, reading a 15 particular one of a plurality of entries of a map, the particular map entry associated with the logical block address of the read request, to obtain a physical address of a particular one of a plurality of pages of the nonvolatile memory, an offset in the particular page to compressed data previously stored in response to 20 writing data corresponding to the logical block address, and a length in bytes of the compressed data, converting the offset in the particular page to the compressed data and the length in bytes

of the compressed data to an address of a first one of a plurality of read units in the particular page and a number of the read units to be read from the particular page, reading from the particular page solely the number of read units, performing error correction 5 decoding on each of the read units read from the particular page to obtain corrected data, extracting the compressed data from the corrected data according to the offset in the particular page to the compressed data and the length in bytes of the compressed data, decompressing the compressed data to produce return data, and 10 returning the return data to the host.

A method EC2 comprising: receiving, at an input/output device and via a host to input/output device interface, a read request to read data corresponding to a logical block address of the read request from nonvolatile memory of the input/output device; and in response to receiving the read request, reading a particular one of a plurality of entries of a map, the particular map entry associated with the logical block address of the read request, to obtain a physical address of a particular one of a plurality of pages of the nonvolatile memory, an offset in the 15 particular page to compressed data previously stored in response to writing data corresponding to the logical block address, and a length in bytes of the compressed data, converting the offset in 20

the particular page to the compressed data and the length in bytes of the compressed data to an address of a first one of a plurality of read units in the particular page and a number of the read units to be read from the particular page, reading from the particular 5 page at least the number of read units and less than all of the read units in the particular page, performing error correction decoding on each of the read units read from the particular page to obtain corrected data, extracting the compressed data from the corrected data according to the offset in the particular page to 10 the compressed data and the length in bytes of the compressed data, decompressing the compressed data to produce return data, and returning the return data to the host.

A method EC3 according to EC1 or EC2, wherein the number of the read units to be read is less than all of the read units in 15 the particular page.

A method EC4 according to either EC1 or EC2, further comprising: determining according to the offset in the particular page to the compressed data and the length in bytes of the compressed data combined with an amount of user data in the 20 particular page that at least a portion of the compressed data is in one or more read units of a subsequent one of the pages of the nonvolatile memory.

A method EC5 according to EC4, wherein, in response to  
the update of the global redundant data on the second processing  
node, the respective local redundancy computation unit of the  
second processing node is enabled to compute second redundant data  
5 according to data of the update of the global redundant data for  
storing on at least some of the respective disks of the second  
processing node.

A method EC6 according to EC1 or EC2, wherein a first one  
of the pages of the nonvolatile memory comprises a first number of  
10 read units, a second one of the pages of the nonvolatile memory  
comprises a second number of read units, and the first number of  
read units is different from the second number of read units.

A method EC7 according to EC1 or EC2, wherein a first one  
of the pages of the nonvolatile memory comprises a first amount of  
15 user data, a second one of the pages of the nonvolatile memory  
comprises a second amount of user data, and the first number of  
user data is different from the second amount of user data.

A method EC8 according to EC1 or EC2, further comprising:  
receiving, at the input/output device and via the host to  
20 input/output device interface, a write request to write the data  
corresponding to the logical block address; in response to  
receiving the write request, compressing the data corresponding to

the logical block address to form compressed write data that is smaller than the data corresponding to the logical block address, writing in the particular page at least a first portion of the compressed write data, and storing in the particular entry, the 5 physical address of the particular page, the offset in the particular page to the compressed write data, and the length in bytes of the compressed write data.

10 A method EC9 according to EC8, further comprising: in response to receiving the request to write data, writing in the particular page a header, the header comprising at least a portion of the logical block address of the request and a length in bytes of the compressed data.

15 A method EC10 according to EC1 or EC2, wherein the logical block address is a first one of a plurality of logical block addresses, and at least one of the read units of the number of read units comprises at least some data corresponding to a different one of the logical block addresses.

20 A method EC11 according to EC1 or EC2, wherein at least one of the read units of the number of read units comprises one or more headers in addition to a portion of the compressed data.

A method EC12 comprising: receiving, at an input/output device and via a host to input/output device interface, a read

request to read data corresponding to a logical block address of the read request from nonvolatile memory of the input/output device; and in response to receiving the read request, reading a particular one of a plurality of entries of a map, the particular 5 map entry associated with the logical block address of the read request, to obtain a physical address of a particular one of a plurality of pages of the nonvolatile memory, an offset in the particular page to variable-sized data previously stored in response to writing data corresponding to the logical block 10 address, and a length in bytes of the variable-sized data, converting the offset in the particular page to the variable-sized data and the length in bytes of the variable-sized data to an address of a first one of a plurality of read units in the particular page and a number of the read units to be read from the 15 particular page, reading from the particular page solely the number of read units, performing error correction decoding on each of the read units read from the particular page to obtain corrected data, extracting the variable-sized data from the corrected data according to the offset in the particular page to the variable- 20 sized data and the length in bytes of the variable-sized data, and returning the extracted data to the host.

A method EC13 according to EC1 or EC22, further

comprising: receiving, at the input/output device and via the host to input/output device interface, a write request to write the variable-sized data corresponding to the logical block address and a size of the variable-sized data; in response to receiving the 5 write request, writing in the particular page at least a first portion of the variable-sized data, and storing in the particular entry, the physical address of the particular page, the offset in the particular page to the variable-sized data, and the length in bytes of the variable-sized data according to the size of the 10 variable-sized data.

In some embodiments, various combinations of all or portions of operations performed by a multi-node storage device or portion(s) thereof, for instance a hard disk drive or a solid-state disk controller of an input/output device enabled for 15 interoperation with a processor (such as a CPU), an input/output controller (such as a RAID-on-chip die), and portions of a processor, microprocessor, system-on-a-chip, application-specific-integrated-circuit, hardware accelerator, or other circuitry providing all or portions of the aforementioned operations, are 20 specified by a specification compatible with processing by a computer system. The specification is in accordance with various descriptions, such as hardware description languages, circuit

descriptions, netlist descriptions, mask descriptions, or layout descriptions. Example descriptions include, but are not limited to: Verilog, VHDL, SPICE, SPICE variants such as PSpice, IBIS, LEF, DEF, GDS-II, OASIS, or other descriptions. In various 5 embodiments, the processing includes any combination of interpretation, compilation, simulation, and synthesis to produce, to verify, or to specify logic and/or circuitry suitable for inclusion on one or more integrated circuits. Each integrated circuit, according to various embodiments, is designable and/or 10 manufacturable according to a variety of techniques. The techniques include a programmable technique (such as a field or mask programmable gate array integrated circuit), a semi-custom technique (such as a wholly or partially cell-based integrated circuit), and a full-custom technique (such as an integrated 15 circuit that is substantially specialized), any combination thereof, or any other technique compatible with design and/or manufacturing of integrated circuits.

In some embodiments, various combinations of all or portions of operations as described by a computer readable medium 20 having a set of instructions stored therein, are performed by execution and/or interpretation of one or more program instructions, by interpretation and/or compiling of one or more

source and/or script language statements, or by execution of binary instructions produced by compiling, translating, and/or interpreting information expressed in programming and/or scripting language statements. The statements are compatible with any 5 standard programming or scripting language (such as C, C++, Fortran, Pascal, Ada, Java, VBscript, and Shell). One or more of the program instructions, the language statements, or the binary instructions, are optionally stored on one or more computer readable storage medium elements. In various embodiments, some, 10 all, or various portions of the program instructions are realized as one or more functions, routines, subroutines, in-line routines, procedures, macros, or portions thereof.

Certain choices have been made in the description merely for convenience in preparing the text and drawings, and unless 15 there is an indication to the contrary, the choices should not be construed per se as conveying additional information regarding structure or operation of the embodiments described. Examples of the choices include, but are not limited to: the particular organization or assignment of the designations used for the figure 20 numbering and the particular organization or assignment of the element identifiers (e.g., the callouts or numerical designators) used to identify and reference the features and elements of the

embodiments.

The words "includes" or "including" are specifically intended to be construed as abstractions describing logical sets of open-ended scope and are not meant to convey physical containment 5 unless explicitly followed by the word "within".

Although the foregoing embodiments have been described in some detail for purposes of clarity of description and understanding, the invention is not limited to the details provided. There are many embodiments of the invention. The 10 disclosed embodiments are exemplary and not restrictive.

Many variations in construction, arrangement, and use are possible consistent with the description, and are within the scope of the claims of the issued patent. For example, interconnect and function-unit bit-widths, clock speeds, and the type of technology 15 used are variable according to various embodiments in each component block. The names given to interconnect and logic are merely exemplary, and should not be construed as limiting the concepts described. The order and arrangement of flowchart and flow diagram process, action, and function elements are variable 20 according to various embodiments. Also, unless specifically stated to the contrary, value ranges specified, maximum and minimum values used, or other particular specifications (such as input/output

device technology types, and the number of entries or stages in registers and buffers), are merely those of the described embodiments, are expected to track improvements and changes in implementation technology, and should not be construed as 5 limitations.

Functionally equivalent techniques known in the art are employable instead of those described to implement various components, sub-systems, operations, functions, routines, sub-routines, in-line routines, procedures, macros, or portions thereof. Many functional aspects of embodiments are realizable 10 selectively in either hardware (e.g., generally dedicated circuitry) or software (e.g., via some manner of programmed controller or processor), as a function of embodiment dependent design constraints and technology trends of faster processing 15 (facilitating migration of functions previously in hardware into software) and higher integration density (facilitating migration of functions previously in software into hardware). Specific variations in various embodiments include, but are not limited to: differences in partitioning; different form factors and 20 configurations; use of different operating systems and other system software; use of different interface standards, network protocols, or communication links; use of different coding types; and other

variations to be expected when implementing the concepts described herein in accordance with the unique engineering and business constraints of a particular application.

The embodiments have been described with detail and 5 environmental context well beyond a minimal implementation of many aspects of the embodiments described. Those of ordinary skill in the art will recognize that some embodiments omit disclosed components or features without altering the basic cooperation among the remaining elements. Much of the details disclosed are not 10 utilized to implement various aspects of the embodiments described. To the extent that the remaining elements are distinguishable from the prior art, components and features that are omitted are not limiting on the concepts described herein.

All such variations in design are insubstantial changes 15 over the teachings conveyed by the described embodiments. The embodiments described herein have broad applicability to other computing and networking applications, and are not limited to the particular application or industry of the described embodiments. The invention is thus to be construed as including all possible 20 modifications and variations encompassed within the scope of the claims of the issued patent.

The functions performed by the diagrams of FIGS. 1-12 may

be implemented using one or more of a conventional general purpose processor, digital computer, microprocessor, microcontroller, RISC (reduced instruction set computer) 5 processor, CISC (complex instruction set computer) processor, SIMD (single instruction multiple data) processor, signal processor, central processing unit (CPU), arithmetic logic unit (ALU), video digital signal processor (VDSP) and/or similar computational machines, programmed according to the teachings 10 of the specification as will be apparent to those skilled in the relevant art(s). Appropriate software, firmware, coding, routines, instructions, opcodes, microcode, and/or program modules may readily be prepared by skilled programmers based on the teachings of the disclosure, as will also be apparent to 15 those skilled in the relevant art(s). The software is generally executed from a medium or several media by one or more of the processors of the machine implementation.

Embodiments of the present invention may also be 20 implemented by the preparation of ASICs (application specific integrated circuits), Platform ASICs, FPGAs (field programmable gate arrays), PLDs (programmable logic devices), CPLDs (complex programmable logic devices), sea-of-gates, RFICs (radio frequency integrated 25 circuits), ASSPs (application specific standard products), one or more monolithic integrated circuits, one or more chips or die arranged as flip-chip

modules and/or multi-chip modules or by interconnecting an appropriate network of conventional component circuits, as is described herein, modifications of which will be readily apparent to those skilled in the art(s).

5

Embodiments of the present invention thus may also include a computer product which may be a storage medium or media and/or a transmission medium or media including instructions which may be used to program a machine to 10 perform one or more processes or methods in accordance with the invention. Execution of instructions contained in the computer product by the machine, along with operations of surrounding circuitry, may transform input data into one or more files on the storage medium and/or one or more output 15 signals representative of a physical object or substance, such as an audio and/or visual depiction. The storage medium may include, but is not limited to, any type of disk including floppy disk, hard drive, magnetic disk, optical disk, CD-ROM, DVD and magneto-optical disks and circuits such 20 as ROMs (read-only memories), RAMs (random access memories), EPROMs (erasable programmable ROMs) , EEPROMs (electrically erasable programmable ROMs), UV PROM (ultra-violet erasable programmable ROMs), Flash memory, magnetic cards, optical cards, and/or any type of media suitable for storing 25 electronic instructions.

The elements of embodiments of the present invention may form part or all of one or more devices, units, components, systems, machines and/ or apparatuses. The devices may include, but are not limited to, servers, workstations, 5 storage array controllers, storage systems, personal computers, laptop computers, notebook computers, palm computers, personal digital assistants, portable electronic devices, battery powered devices, set-top boxes, encoders, decoders, transcoders, compressors, decompressors, pre- 10 processors, post-processors, transmitters, receivers, transceivers, cipher circuits, cellular telephones, digital cameras, positioning and/or navigation systems, medical equipment, heads-up displays, wireless devices, audio recording, audio storage and/or audio playback devices, video 15 recording, video storage and/or video playback devices, game platforms, peripherals and/or multi-chip modules. Those skilled in the relevant art (s) would understand that the elements of the invention may be implemented in other types of devices to meet the criteria of a particular application.

20

The terms "may" and "generally" when used herein in conjunction with "is(are)" and verbs are meant to communicate the intention that the description is exemplary and believed to be broad enough to encompass both the specific examples 25 presented in the disclosure as well as alternative examples that could be

derived based on the disclosure. The terms "may" and "generally" as used herein should not be construed to necessarily imply the desirability or possibility of omitting a corresponding element.

5

While the invention has been particularly shown and described with reference to embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made without departing from the 10 scope of the invention as defined by the claims.

The disclosures in United States patent application no. 14/055336 and 61/866,672, from which this application claims priority, and in the abstract accompanying this application 15 are incorporated herein by reference.

CLAIMS

1. A method for using a variable-size flash transition layer, comprising the steps of:

5 (A) receiving, at an apparatus from a host, a read request to read particular data corresponding to a logical block address from a nonvolatile memory of said apparatus;

(B) reading a particular one of a plurality of entries in a map to obtain (i) a physical address of a particular one 10 of a plurality of pages of said nonvolatile memory, (ii) an offset in said particular page to compressed data previously stored in response to writing said compressed data corresponding to said logical block address and (iii) a length of said compressed data, wherein said particular entry 15 is associated with said logical block address;

(C) converting said offset and said length to (i) an address of a given one of a plurality of read units in said particular page and (ii) a number of said read units to be read from said particular page; and

20 (D) reading from said particular page at most said number of said read units starting from said given read unit, wherein a granularity of said offset and said length is finer than a size of one of said read units.

25 2. The method according to claim 1, further comprising the steps of:

performing error correction decoding on each of said read units as read from said particular page to generate corrected data; and

30 extracting said compressed data from said corrected data according to both (i) said offset in said particular page to said compressed data and (ii) said length of said compressed data.

3. The method according to claim 2, further comprising the steps of:

decompressing said compressed data to generate return  
5 data; and  
transferring said return data to said host.

4. The method according to claim 1, 2 or 3, wherein said number of said read units to be read is less than all of said  
10 read units in said particular page.

5. The method according to any preceding claim, further comprising the step of:

determining that at least a portion of said compressed  
15 data is in one or more subsequent read units of a subsequent  
one of said pages of said nonvolatile memory based on both  
(i) said offset in said particular page to said compressed  
data and (ii) said length of said compressed data combined  
with an amount of user data in said particular page.

20

6. The method according to claim 5, further comprising the step of:

reading from said subsequent page at most said one or  
more subsequent read units.

25

7. The method according to any claim preceding, wherein  
(i) a first one of said pages of said nonvolatile memory  
includes a first number of said read units, (ii) a second one  
of said pages of said nonvolatile memory includes a second  
30 number of said read units and (iii) said first number is  
different from said second number.

8. The method according to any of claims 1 to 6, wherein  
(i) a first one of said pages of said nonvolatile memory  
includes a first amount of user data, (ii) a second one of  
said pages of said nonvolatile memory includes a second  
5 amount of said user data and (iii) said first amount is  
different from said second amount.

9. The method according to any preceding claim, further  
comprising the steps of:

10 receiving a write request to write said particular data  
in said nonvolatile memory;

compressing said particular data to generate said  
compressed data that is smaller than said particular data;

15 writing in said particular page at least a portion of  
said compressed data; and

storing in said particular entry (i) said physical  
address of said particular page, (ii) said offset in said  
particular page to said compressed data and (iii) said length  
of said compressed data.

20

10. The method according to claim 9, further comprising  
the step of:

writing in said particular page a header that includes  
at least a portion of said logical block address of said  
25 write request and said length.

11. The method according to any preceding claim, wherein  
(i) said logical block address is one of a plurality of  
logical block addresses and (ii) at least one of said read  
30 units of said number of said read units includes at least  
some different data corresponding to a different one of said  
logical block addresses.

12. The method according to any preceding claim, wherein  
at least one of said read units of said number of said read  
units includes (i) one or more headers and (ii) a portion of  
5 said compressed data.

13. An apparatus comprising:  
a nonvolatile memory; and  
a circuit configured to (i) receive a read request from  
10 a host to read particular data corresponding to a logical  
block address from said nonvolatile memory, (ii) read a  
particular one of a plurality of entries in a map to obtain  
(a) a physical address of a particular one of a plurality of  
pages of said nonvolatile memory, (b) an offset in said  
15 particular page to compressed data previously stored in  
response to writing said compressed data corresponding to  
said logical block address and (c) a length of said  
compressed data, wherein said particular entry is associated  
with said logical block address, (iii) convert said offset  
20 and said length to (a) an address of a given one of a plurality  
of read units in said particular page and (b) a number of said  
read units to be read from said particular page and (iv) read  
from said particular page at most said number of said read  
units starting from said given read unit, wherein a granularity  
25 of said offset and said length is finer than a size of one of  
said read units.

14. The apparatus according to claim 13, wherein said  
circuit is further configured to (i) perform error correction  
30 decoding on each of said read units as read from said  
particular page to generate corrected data and (ii) extract  
said compressed data from said corrected data according to both

(a) said offset in said particular page to said compressed data and (b) said length of said compressed data.

15. The apparatus according to claim 14, wherein said 5 circuit is further configured to (i) decompress said compressed data to generate return data and (ii) transfer said return data to said host.

16. The apparatus according to claim 13, 14 or 15, wherein 10 said number of said read units to be read is less than all of said read units in said particular page.

17. The apparatus according to any of claims 13 to 16, wherein said circuit is further configured to determine that at 15 least a portion of said compressed data is in one or more subsequent read units of a subsequent one of said pages of said nonvolatile memory based on both (i) said offset in said particular page to said compressed data and (ii) said length of said compressed data combined with an amount of user data in 20 said particular page.

18. The apparatus according to claim 17, wherein said circuit is further configured to read from said subsequent page at most said one or more subsequent read units.

25  
19. The apparatus according to any of claims 13 to 18, wherein (i) a first one of said pages of said nonvolatile memory includes a first number of said read units, (ii) a second one of said pages of said nonvolatile memory includes a 30 second number of said read units and (iii) said first number is different from said second number.

20. The apparatus according to any of claims 13 to 19,  
wherein said apparatus is implemented as one or more integrated  
circuits.

5 21. A method as herein described and as illustrated in the  
accompanying drawings.

22. An apparatus as herein described and as illustrated in  
the accompanying drawings.



**Application No:** GB1414545.2

**Examiner:** Mr David Maskery

**Claims searched:** 1 - 20

**Date of search:** 18 February 2015

## Patents Act 1977: Search Report under Section 17

### Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
A,E	-	US 8843711 B (YADAV et AL) See columns 10 and 11.
A	-	US 2014/0223089 A (KANG et AL) See paragraphs 71 - 110, and WO 2013/042880.
A	-	US 2010/0017578 A (MANSSON et AL) See paragraphs 100 - 110.
A	-	US 2007/0143569 A (SANDERS et AL) See paragraphs 38 - 56.
A	-	US 6374341 B (NIJHAWAN et AL) See columns 4 and 8.

### Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

### Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC<sup>X</sup> :

Worldwide search of patent documents classified in the following areas of the IPC

G06F; G11C

The following online and other databases have been used in the preparation of this search report

EPODOC, WPI, TXTE.



**International Classification:**

Subclass	Subgroup	Valid From
G06F	0012/02	01/01/2006
G11C	0016/08	01/01/2006
G11C	0029/04	01/01/2006