



(51) International Patent Classification:
G06F 3/06 (2006.01)

(21) International Application Number:
PCT/US2017/046142

(22) International Filing Date:
09 August 2017 (09.08.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/403,340 03 October 2016 (03.10.2016) US
15/388,892 22 December 2016 (22.12.2016) US

(71) Applicant: CYPRESS SEMICONDUCTOR CORPORATION [US/US]; 198 Champion Court, San Jose, California 95134 (US).

(72) Inventors: MCCLAIN, Mark; 12598 Dormouse Road, San Diego, California 92129 (US). OBEREINER, Willy; 2416 Whitethorne Dr, San Jose, California 95128 (US). HOEHLER, Rainer; 1645 Oakmont Dr, Colorado Springs, Colorado 80921 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

— as to the identity of the inventor (Rule 4.17(i))

Published:

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

(54) Title: SYSTEMS, METHODS, AND DEVICES FOR USER CONFIGURABLE WEAR LEVELING OF NON-VOLATILE MEMORY




FIG. 1

(57) Abstract: Disclosed herein are systems, methods, and devices for user configurable wear leveling of non-volatile memory devices. Devices include a non-volatile memory including a plurality of physical memory portions, where each of the plurality of physical memory portions is configured to be mapped to a logical memory portion of a plurality of logical memory portions. Devices may also include wear leveling control circuitry configured to receive a plurality of wear leveling parameters, and further configured to determine a plurality of wear leveling characteristics based on the plurality of wear leveling parameters received, where the plurality of wear leveling characteristics identify an amount of wear leveling allowed for each of the plurality of logical memory portions.

SYSTEMS, METHODS, AND DEVICES FOR USER CONFIGURABLE WEAR LEVELING OF NON-VOLATILE MEMORY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is an International Application of US Non-Provisional Application No. 15/388,892, filed on December 22, 2016, which claims the benefit under 35 U.S.C. § 119(e) of US Provisional Patent Application No. 62/403,340, filed on October 03, 2016, all of which are incorporated by reference herein in their entirety for all purposes.

TECHNICAL FIELD

[0002] This disclosure generally relates to non-volatile memories and, more specifically, to wear leveling associated with non-volatile memories.

BACKGROUND

[0003] Non-volatile memory devices include memory cell arrays that may be configured as a memory component that stores data values for access by various other components of a system. Memory cells may include memory elements that may degrade over time, by retaining valid data for less time, as the cells experience an increased number of program and erase operations (PE cycles). Such memory cells may include physical storage locations grouped in physical sectors that are mapped to logical addresses and sectors that represent the memory space of the memory cell array. Wear leveling may be implemented to change a physical location of data and spread wear from PE cycles across several physical sectors. However, such wear leveling may also move less frequently cycled data to physical sectors that have already experienced numerous PE cycles. Accordingly, wear leveling techniques may inadvertently reduce the retention time of data that rarely experiences PE cycles.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] Figure 1 illustrates an example of a configurable wear leveling memory device including sectors of more than one size, configured in accordance with some embodiments.

5 [0005] Figure 2 illustrates an example of a configurable wear leveling memory device including uniform sector sizes, configured in accordance with some embodiments.

[0006] Figure 3 illustrates an example of a configurable wear leveling memory device where the parameter sectors are excluded from wear leveling, and other sectors are grouped in fixed partition sizes, configured in accordance with some embodiments.

10 [0007] Figure 4 illustrates a flow chart of an example of a method for configuring wear leveling of a memory device, implemented in accordance with some embodiments.

[0008] Figure 5 illustrates a flow chart of another example of a method for configuring wear leveling of a memory device, implemented in accordance with some embodiments.

[0009] Figure 6 illustrates an example of a configurable wear leveling memory system, configured in accordance with some embodiments.

15

DETAILED DESCRIPTION

[0010] In the following description, numerous specific details are set forth in order to provide a thorough understanding of the presented concepts. The presented concepts may be practiced without some or all of these specific details. In other instances, well known
5 process operations have not been described in detail so as to not unnecessarily obscure the described concepts. While some concepts will be described in conjunction with the specific examples, it will be understood that these examples are not intended to be limiting.

[0011] In various embodiments, memory devices may be limited as to the number of
10 program and erase (PE) cycles they may undergo because the memory cells degrade to a point they are unable to reliably retain data for a sufficient time period. As the number of PE cycles (also referred to herein as endurance) is increased, the length of time data is reliably retained (also referred to herein as retention) is reduced. In some embodiments, wear leveling techniques may be implemented to move frequently modified information
15 to different areas of the memory array to limit the number of PE cycles implemented in any one area. Such wear leveling techniques may be used to reduce the effective PE cycles on the area of memory storing a logical block of information. However, when applied to an entire memory array, wear leveling also increases the effective PE cycles on areas of memory storing logical blocks of information which are modified less frequently,
20 thereby reducing the retention period of the less frequently modified information.

Moreover, such reduction in retention period due to movement of the data may occur without the knowledge of the user of the memory. The user may believe that low cycled areas of the memory array have longer retention periods, but wear leveling techniques may have actually reduced the retention period by moving information from low cycled
25 physical sectors to higher cycled physical sectors. In various embodiments, a sector may be a unit of erase and/or program or overwrite. Accordingly, a size of a sector may be determined or defined by a size of a portion of memory included in an erase, program, or overwrite operation.

[0012] As will be discussed in greater detail below, wear leveling techniques may include various methods of mapping logical sectors to physical sectors and tracking the erase count of each physical sector so that when a PE cycle threshold value is reached for a physical sector, the data in the sector is moved to another physical sector with fewer erase
5 cycles and the logical to physical map is updated. This makes the logical sector appears to support many more PE cycles than is possible for any one physical sector thus spreading the PE cycle wear of a logical sector across many physical sectors such that over time all physical sectors have a roughly equal number of PE cycles and wear.

[0013] Some wear leveling techniques remain limited because they do not provide any
10 user control of which portions of the memory array are wear leveled to increase cycling capacity, and which portions of the memory array are not wear leveled to maintain longer data retention. For example, some memory devices with an internal wear leveling mechanism use most or all physical sectors in the device in the wear leveling operation. In such embodiments, a memory device is unable to control which logical sectors may be
15 relocated into physical sectors that have received a higher number of PE cycles than the physical sector currently in use by a logical sector. This can cause a logical sector containing data that needs long retention time to be moved to a physical sector that has undergone a higher number of PE cycles and therefore reduced retention time, without the knowledge of the flash device user. Accordingly, areas of memory where high
20 retention is desirable, such as rarely updated code, may be moved into highly PE cycled sectors by a wear leveling algorithm, thus reducing the retention time for that memory area without the user's knowledge.

[0014] Various embodiments are disclosed herein that are configured to provide memory devices that implement user configurable wear leveling. Accordingly, some areas of the
25 memory array may undergo wear leveling for higher apparent endurance of some logical sectors, and other areas of the memory array do not undergo wear leveling to maintain retention. Such portions of the memory array that undergo wear leveling and that do not undergo wear leveling may be explicitly defined in a number of ways disclosed herein. In one example, address pointers may be used to define boundaries between partitions
30 (address ranges) of the memory array and to indicate for each partition whether that

partition is within a pool of sectors used in wear leveling or not within the wear leveling pool. Such explicit definition of logical address ranges gives the memory user control over which memory areas receive lower PE cycles and therefore have longer retention, and which areas are effectively able to have higher PE cycles than is possible for a single
5 area of memory to support. Additional details of such configurable wear leveling are provided below.

[0015] Figure 1 illustrates an example of a configurable wear leveling memory device including sectors of more than one size, configured in accordance with some embodiments. As will be discussed in greater detail below, a memory device, such as memory device 100,
10 may be configured to enable a user to explicitly select some logical portions of memory for inclusion in a set or pool of physical memory portions used in wear leveling, and further enable a user to explicitly exclude other logical memory portions from use in such wear leveling. A contiguous group of logical portions is referred to as a partition, and each partition may be designated for use in wear leveling or excluded from wear leveling.

15 Accordingly, various embodiments disclosed herein provide flexibility to a user to choose which logical portions may have low PE cycles and therefore longer data retention versus which logical portions may have high PE cycle endurance but shorter retention. As will be discussed in greater detail below, logical portions may describe logical sectors, and physical portions may describe physical sectors. However, in various embodiments, such portions
20 might not be aligned with sectors, and a portion might not directly correspond with a sector in memory space.

[0016] In various embodiments, memory device 100 may be a storage device that includes several non-volatile memory elements in a memory array, such as memory array 102.

Accordingly, memory array 102 may include numerous columns and rows of memory
25 elements that may be grouped or divided into sectors. In various embodiments, such sectors may be a group of data units, such as words, that may be erased as a unit. In various embodiments, the sectors shown in Figure 1 are logical sectors that may be mapped to physical sectors. In some embodiments, logical sectors may be sectors of data that are referenced by a memory array address that is independent of the physical location of the data
30 within the memory array. Accordingly, a host system may view a logical sector as residing at

particular address (e.g. starting at address zero of the memory array). In this example, the data may actually be located in a different part of the memory array address space, at a different physical address. In various embodiments, memory device 100 maintains a mapping table from the logical address of the data to corresponding actual physical locations. In this way, memory device 100 may include various physical memory sectors that may be mapped to logical memory sectors. In various embodiments, memory array 102 may include memory sectors of varying sizes. For example, memory array 102 may include supervisory memory sector 130 which may have a first size, such as 4 kilobytes (KB), and may further include first memory sector 104 which may be a large memory sector and have a second size, such as 256 KB. In some embodiments, memory array 102 may also include second memory sector 106, which may be a last memory sector in memory array 102, and may have a first size of 4 KB as well.

[0017] Furthermore, the sectors may be grouped into partitions. Accordingly, memory array 102 may include various memory partitions that are contiguous group of logical sectors which may be defined by boundaries. For example, memory array 102 may include first boundary 110, second boundary 112, third boundary 114, and fourth boundary 116 which may define first partition 118, second partition 120, third partition 122, fourth partition 124, and fifth partition 126. As will be discussed in greater detail below, such boundaries may be configurable based on wear leveling parameters which may be user specified. Accordingly, a user may specify and designate sizes and boundaries associated with partitions. Moreover, the wear leveling parameters may also indicate whether or not a particular partition is to be included in physical sectors used in wear leveling. As will be discussed in greater detail below, the sizes of partitions may be flexible or configurable based on user input.

[0018] As previously discussed, wear leveling may be a sector erase management technique used to improve the logical sector endurance in memory devices, such as memory device 100. Wear leveling spreads the PE cycles applied to one logical sector, across multiple physical sectors which are in a set or pool of physical sectors in the memory device. The pool of physical sectors is available for the relocation of logical sector data. For example, when implementing wear leveling, logical sectors may be mapped to physical sectors, and an erase count of each physical sector may be tracked so that when a PE cycle threshold value is

reached or exceeded for a particular physical sector, the data in the physical sector is moved to another physical sector with fewer erase cycles, and the logical to physical map is updated. Such tracking of PE cycles of physical sectors, and movement of logical sectors across physical sectors makes the logical sector appear to support many more PE cycles than is possible for any one physical sector. Accordingly, the PE cycle wear of a logical sector is spread across many physical sectors such that over time all physical sectors have roughly equal number of PE cycle wear.

[0019] In a specific example, various sectors may utilize a form of wear leveling that places one smaller logical sector in a larger physical sector in which sub-sections of the physical sector may be erased individually. As the portion of the physical sector used by the smaller logical sector receives a certain number of erase cycles, the logical sector data is moved to the next sub-section of the larger physical sector. All sub-sections not currently in use by the smaller logical sector are kept erased to minimize column leakage. This movement of the logical sector through the physical sector spreads the PE cycle wear across the entire larger physical sector. In various embodiments, the smaller logical sector data remains within the same larger physical sector.

[0020] Accordingly, wear leveling is based on a mapping of logical sectors to physical sectors. During the lifetime of a memory device, such as memory device 100, this mapping is changed as logical sector content is moved from high PE cycle physical sectors to lower PE cycle physical sectors. This provides a more uniform distribution of PE cycles over all physical sectors. The logical to physical mapping information is stored in a dedicated memory array which is updated when logical sector data is moved from one physical sector to another. Sector data movement may occur when a program and/or an erase command is given by the host system or a component, such as a processor.

[0021] As will be discussed in greater detail below, control circuitry, such as control circuitry 128 and a dedicated storage location or register, such as supervisory memory sector 130, may be utilized to configure the mapping of logical sectors to physical sectors that undergo wear leveling. As disclosed herein, control circuitry may also be referred to herein as wear leveling control circuitry. More specifically, wear leveling parameters may be

received and used to explicitly identify which logical sectors should be mapped to physical sectors included in a wear leveling pool and do undergo wear leveling. The wear leveling parameters may also be used to explicitly identify which logical sectors should be mapped to physical sectors which are not included in the wear leveling pool and do not undergo wear leveling. As discussed above, a memory device, such as memory device 100, may include various memory sectors that are included in a wear leveling pool and undergo wear leveling, and may further include memory sectors that are not included in a wear leveling pool and do not undergo wear leveling. In this way, wear leveling parameters may be configured to identify which partitions and corresponding logical memory sectors should be included in a pool of physical sectors undergoing wear leveling, and which partitions and corresponding logical memory sectors should be excluded.

[0022] More specifically, supervisory memory sector 130 may be configured to store one or more data values that identify a mapping that identifies logical memory sectors, further identifies partitions that such memory sectors are included in, and further identifies a wear leveling setting associated with such a partition. An example of a representation of such a mapping is shown below in Table 1:

Sectors	Partitions	Wear Leveling Setting (0=excluded, 1= included)
Sector 1	First Partition	0
Sector 2		
Sector 3	Second Partition	1
Sector 4		
Sector 5	Third Partition	1
Sector 6		
Sector 7	Fourth Partition	0
Sector 8		

Table 1

[0023] Accordingly, as shown in Table 1, a mapping may identify logical address sectors as well as their associated partitions. Moreover, the mapping may further identify a wear leveling setting for each partition. As discussed above, such a wear leveling setting or characteristic may be user specified and determined based on a received user input. In this way, a user may explicitly assign logical memory sectors and partitions to be included or

excluded from a wear leveling pool of physical memory sectors. Furthermore, if a partition, such as fifth partition 126, is not identified in the mapping, it may be assigned a default value, such as “1” indicating it should be included. Similarly, such a default value may be “0”.

5 **[0024]** Various embodiments are contemplated and disclosed herein regarding the assigning of logical memory sectors to wear leveling pools. For example, while explicit partition labeling was previously described, other suitable techniques may be implemented as well. For example, a storage location, such as supervisory memory sector 130, may store several pre-defined partition maps, and a configuration parameter or value may be used to select a
 10 map that is used. In some embodiments, such pre-defined maps may be generated based on different views of the address space of memory array 102. For example, the pre-defined maps may include a first map that excludes the lowest 1/8th addresses of the memory address space from wear leveling. The pre-defined maps may further include a second map that excludes the lowest 1/4 of the address space from wear leveling. The pre-defined maps may
 15 also include a third map that excludes the lower 1/2 of the address space from wear leveling. In this example, a user may provide an input that includes the configuration parameter and selects a particular pre-defined map to apply. Examples of such pre-defined partition maps are illustrated in Tables 2, 3, and 4 below:

Sectors	Partitions	Wear Leveling Setting (0=excluded, 1= included)
Sector 0	Partition 0	0
Sector 1	Partition 1	1
Sector 2		
Sector 3		
Sector 4		
Sector 5		
Sector 6		
Sector 7		

Table 2

20 **[0025]** As shown in Table 2, the lower 1/8th of the address space is excluded from wear leveling.

Sectors	Partitions	Wear Leveling Setting (0=excluded, 1= included)
Sector 0	Partition 0	0
Sector 1		
Sector 2	Partition 1	1
Sector 3		
Sector 4		
Sector 5		
Sector 6		
Sector 7		

Table 3

[0026] As shown in Table 3, the lower 1/4th of the address space is excluded from wear leveling.

Sectors	Partitions	Wear Leveling Setting (0=excluded, 1= included)
Sector 0	Partition 0	0
Sector 1		
Sector 2		
Sector 3		
Sector 4	Partition 1	1
Sector 5		
Sector 6		
Sector 7		

Table 4

5 **[0027]** As shown in Table 4, the lower 1/2 of the address space is excluded from wear leveling. While the above examples illustrate some embodiments, various other embodiments are contemplated and disclosed herein. For example, if more memory sectors are implemented and available, memory space allocation may be implemented at granularities of 1/16 or 1/32 of memory space. Furthermore, memory portions may be assigned from either

10 side of memory, such as from an address of “0” or from the end or top of memory. Accordingly the above examples are not intended to limit the scope of the embodiments disclosed herein.

[0028] Furthermore, the assigning of logical memory sectors to wear leveling pools may be implemented by utilizing address pointers to define address ranges for partitions to be

included or excluded from wear leveling. For example, address pointer pairs may identify beginning and ending points of address ranges that define partitions, and further identify a wear leveling setting associated with such partitions. As illustrated below in Table 5, a first address pointer pair including a first and second address pointer may identify a beginning and ending point of partition 0, and further identify a wear leveling setting indicating that partition 0 should be excluded from the wear leveling pool. Moreover, a second address pointer pair including a third and fourth address pointer may identify a beginning and ending point of partition 1, and further identify a wear leveling setting indicating that partition 1 should be included in the wear leveling pool. Furthermore a third address pointer pair including a fifth and sixth address pointer may identify a beginning and ending point of partition 2, and further identify a wear leveling setting indicating that partition 2 should be excluded from the wear leveling pool. In some embodiments, address pointer pairs may only be provided for partitions included in the wear leveling pool, and other partitions may be set to a default value of “0” and excluded from the wear leveling pool. In various embodiments, the opposite may be true, and address pointer pairs may only be provided for partitions excluded from the wear leveling pool.

Sectors	Partitions	Address Range Pointers
Sector 0	Partition 0 included	Start 0
Sector 1		End 0
Sector 2	Partition 1 included	Start 1
Sector 3		
Sector 4		
Sector 5		End 1
Sector 6	Partition 2 excluded	Start 0
Sector 7		End 0

Table 5

[0029] Moreover, the assigning of logical memory sectors to wear leveling pools may be implemented by utilizing a single address pointer per partition. In some embodiments, partitions may be defined relative to a reference address or other pointers. For example, a first partition (Partition 0) may be defined by an address range identified by a first pointer (Pointer A) to a base address, which may be a starting address of the memory. A second

partition (Partition 1) may be defined by an address range identified by a second pointer (Pointer B) to the first pointer (Pointer A), or more specifically, to an increment of the first pointer (Pointer A + 1). Moreover, a third partition (Partition 2) may be defined by an address range identified by another reference address, such as the end of the address space of the memory array, to the second pointer (Pointer B) , or more specifically, to an increment of the first pointer (Pointer B + 1). As shown in Table 6 below, the mapping may also store wear leveling settings for the defined partitions, and a toggle or flag bit for validity of the address pointers.

Sectors	Partitions	Address range pointers	Address pointer valid (0=invalid, 1= valid)	Wear Leveling Setting (0=excluded, 1= included)
Sector 0	Partition 0			
Sector 1				
Sector 2	Partition 1	Pointer A+1	1	1
Sector 3				
Sector 4				
Sector 5				
Sector 6	Partition 2	Pointer B+1	1	0
Sector 7				

Table 6

[0030] While the above description has described inclusion in and exclusion from wear leveling, as will be discussed in greater detail below, various amounts of wear leveling applied to sectors may also be specified. For example, the wear leveling parameters may include wear leveling characteristics or settings that specify or identify an amount of wear leveling to be applied to each logical sector and each partition associated with such logical sectors. Thus, in addition to an amount of wear leveling characterizing inclusion in wear leveling, the amount may also characterize a quantity or degree of wear leveling. For example, the wear leveling characteristics may include a first wear leveling setting indicating a first limit of wear leveling that should be applied to a first partition, a second wear leveling

setting indicating a second limit of wear leveling that should be applied to a second partition, a third wear leveling setting indicating a third limit of wear leveling should be applied to a third partition, and a fourth wear leveling setting indicating no wear leveling should be applied to a fourth partition.

- 5 **[0031]** In various embodiments, the wear leveling characteristics may be determined based on one or more additional types of information or parameter included in the wear leveling parameters. For example, the wear leveling parameters may specify a number of PE cycles associated with memory sectors. In various embodiments, such a number may be user specified, and may characterize an expected number of PE cycles. In some embodiments, a component, such as control circuitry 128, may be configured to identify an amount of wear leveling applied to the memory sectors based on the specified number of PE cycles. For example, if the number is above a threshold value of PE cycles, a memory sector may be included in a wear leveling pool. Moreover, if the number is below a threshold value of PE cycles, a memory sector may be excluded from a wear leveling pool.
- 10
- 15 **[0032]** As discussed above, the wear leveling parameters may specify logical sizes of partitions. In some embodiments, a component, such as control circuitry 128, may be configured to identify an amount of wear leveling applied to the memory sectors based on the size of the partition. For example, if the identified size is larger than a designated size, a memory sector may be included in a wear leveling pool. Moreover, if the number is smaller than a designated size, a memory sector may be excluded from a wear leveling pool. In this example, the designated size may be a default value or may be user specified. Additional details of such configuration of wear leveling is discussed in greater detail below with reference Figure 3, Figure 4, Figure 5, and Figure 6.
- 20

- 25 **[0033]** Figure 2 illustrates an example of a configurable wear leveling memory device including uniform sector sizes, configured in accordance with some embodiments. As discussed above, a memory device, such as memory device 200, may be configured to enable a user to explicitly select some logical sectors of memory for inclusion in a set or pool of physical memory sectors used in wear leveling, and is further configured to enable a user to explicitly exclude other logical memory sectors from use in such wear leveling. As also

discussed above, a memory array, such as memory array 102 may include various partitions defined by various boundaries and may be communicatively coupled with control circuitry 128. For example, memory array 102 may include first boundary 110, second boundary 112, third boundary 114, and fourth boundary 116 which may define first partition 118, second partition 120, third partition 122, fourth partition 124, and fifth partition 126. As discussed above, such boundaries may be configurable based on wear leveling parameters which may be user specified. Accordingly, a user may specify and designate sizes and boundaries associated with partitions. Moreover, the wear leveling parameters may also indicate inclusion or exclusion of partitions from sets or pools of physical sectors that undergo wear leveling, as may be managed by control circuitry 128 and supervisory memory sector 130.

[0034] As shown in Figure 2, memory array 102 may include uniformly sized memory sectors. For example, all logical memory sectors included in memory array 102 may be of the same size and may be included in fixed size partitions. In this example, first memory sector 104, which may be a first memory sector in memory array 102, may be the same size as another memory sector, such as a last memory sector in memory array 102. More specifically, such memory sectors may have a size of 256 KB. In this way, embodiments disclosed herein may be implemented utilizing memory devices that have uniform sector sizes.

[0035] Figure 3 illustrates an example of a configurable wear leveling memory device where the parameter sectors are excluded from wear leveling, and other sectors are grouped in fixed partition sizes, configured in accordance with some embodiments. Figure 3 also illustrates an example of a configuration of memory array 102 in which there are small memory sectors at the top and bottom of memory array 102, with large memory sectors in the middle of memory array 102. As discussed above, a memory device, such as memory device 300, may include a memory array, such as memory array 102 may include various partitions defined by various boundaries and may be communicatively coupled with control circuitry 128. For example, memory array 102 may include first boundary 110, second boundary 112, third boundary 114, and fourth boundary 116, which may define first partition 118, second partition 120, third partition 122, fourth partition 124, and fifth partition 126. In various embodiments, the end of fifth partition 126 may be defined by the end of memory array 102

which may provide a fifth boundary. As discussed above, such partitions may include memory sectors such as first memory sector 104, which may be a first memory sector 104 in memory array 102, and second memory sector 106, which may be a last memory sector in memory array 102. Moreover, memory array 102 may also include supervisory memory sector 130.

[0036] In various embodiments, memory array 102 may be configured to have partitions that have fixed size. Accordingly, sizes of the partitions may be specified by a hardware manufacturer, or may be determined based on hardware properties of memory device 100, such as a size of memory array 102. Moreover, in some embodiments, such partitions may be uniformly sized. As discussed above, the inclusion and exclusion of logical memory sectors included in partitions may be determined based on wear leveling parameters stored in supervisory memory sector 130. As will be discussed in greater detail below with reference to Figure 4 and Figure 5, such wear leveling parameters may be user specified, and may be received from a user via a communications interface.

[0037] Figure 4 illustrates a flow chart of an example of a method for configuring wear leveling of a memory device, implemented in accordance with some embodiments. As similarly discussed above, a memory device may be utilized to enable a user to explicitly select some logical sectors of memory for inclusion in a set or pool of physical memory sectors used in wear leveling, and further enable a user to explicitly exclude other logical memory sectors from use in such wear leveling. As will be discussed in greater detail below, such configurability may provide increased flexibility to a user when implementing wear leveling.

[0038] Method 400 may commence with operation 402 during which a partition identifier associated with a memory array may be set to a designated value. As discussed above with reference to Figure 1, Figure 2, and Figure 3, logical sectors of a memory array may be represented as partitions. Accordingly, during method 400, such partitions may be configured and defined. In various embodiments, method 400 may commence with configuring an initial partition, and may start by setting the designated value identifying a partition to a default value that may be a partition number, such as 0.

[0039] Method 400 may proceed to operation 404 during which partition boundary identifiers and wear leveling identifiers may be generated. In various embodiments, the partition boundary identifiers may be one or more data values that identify the boundaries of the partition that is being configured and defined. Accordingly, the partition boundary
5 identifiers may include data values characterizing address pointers, or any other suitable identifier discussed above with reference to Figure 1. In this way, during operation 404, boundaries may be defined for a particular partition within a logical memory space of a memory array. As discussed above, such boundary identifiers may be user specified, and based on a received user input. More specifically, partition boundary identifiers and wear
10 leveling identifiers may be generated based on the received wear leveling parameters.

[0040] Accordingly, a component, such as control circuitry, may generate the boundary identifiers based on the received input that includes the wear leveling parameters.

Furthermore, wear leveling identifiers may also be generated based on a received input. Such wear leveling identifiers may also be referred to herein as wear leveling characteristics or
15 settings. Accordingly, during operation 404, wear leveling identifiers may be generated that identify whether or not the partition should be included in a wear leveling pool. In some embodiments, the wear leveling identifiers may further identify an amount or limit of wear leveling applied to the partition if the partition is included in the wear leveling pool.

[0041] Method 400 may proceed to operation 406 during which it may be determined if
20 additional partitions should be generated. As discussed above, the memory space of a memory array may be divided into numerous partitions, and various iterations of operation 404 may be implemented to generate partition boundary identifiers and wear leveling identifiers for each partition. If it is determined that no additional partitions should be generated, method 400 may terminate. However, if it is determined that additional partitions
25 should be generated, method 400 may proceed to operation 408. In various embodiments, such a determination may be made based on a user input. For example, a user may specify that additional partitions should be generated, and method 400 may proceed to operation 408. In some embodiments, such a determination may also be made based on a comparison of the partition boundary defined during operation 404 and a particular address in the memory
30 space of the memory array. For example, if the defined boundary, which may be an address

pointer, of the partition defined during operation 404 is not the last address in the memory array, and some memory space remains undefined, method 400 may proceed to operation 408.

[0042] Accordingly, during operation 408 a partition identifier may be incremented. In this way, the partition identifier may be modified or adjusted to identify a new and different partition, and method 400 may return to operation 404 during which partition boundary identifiers and wear leveling identifiers may be generated for the new partition. For example, if the previously defined partition identifier was “0”, the partition identifier may be incremented to “1”, and partition boundary identifiers and wear leveling identifiers may be generated for partition 1 during another iteration of operation 404. In this way, numerous partitions may be defined, and wear leveling identifiers associated with those partitions may also be defined such that an amount or limit of wear leveling, or no wear leveling at all, is specified for each defined partition.

[0043] Figure 5 illustrates a flow chart of another example of a method for configuring wear leveling of a memory device, implemented in accordance with some embodiments. As similarly discussed above, a memory device may be utilized to enable a user to explicitly select some logical sectors of memory for inclusion in a set or pool of physical memory sectors used in wear leveling, and further enable a user to explicitly exclude other logical memory sectors from use in such wear leveling. Moreover, such selections may form the basis of logical sector/physical sector mappings that underlie subsequent wear leveling operations. In this way, a user may be provided with the ability to configure which logical sectors undergo how much wear leveling.

[0044] Method 500 may commence with operation 502 during which a plurality of partition boundary identifiers and a plurality of wear leveling identifiers may be generated. As discussed above with reference to Figure 4, partition boundary identifiers and wear leveling identifiers may be generated for various different partitions of a logical memory space of a memory array. As also stated above, such identifiers may be generated iteratively as several partitions are defined by partition boundary identifiers, which may be address pointers, and wear leveling characteristics are defined for each partition based on the designation of wear

leveling identifiers. Moreover, such partition boundary identifiers and wear leveling identifiers may be generated based on user inputs received as wear leveling parameters. Accordingly, the identifiers may be user configurable, thus enabling user configuration of the sizes and number of partitions, as well as the amount of wear leveling allowed for each
5 partition.

[0045] In some embodiments, a component, such as control circuitry, may generate, store, and maintain the partition boundary identifiers and wear leveling identifiers in one or more data structures. For example, the control circuitry may generate a first data structure that identifies the mapping of logical sector addresses to partitions and their associated partition
10 identifiers. Moreover, the control circuitry may generate and store a second data structure that identifies logical sectors and partitions, and maps them to the wear leveling pool, as may be indicated by a flag or bit, as well as their associated wear leveling characteristics or settings. Furthermore, as will be discussed in greater detail below, the control circuitry may also generate a third data structure that stores an address mapping of logical to physical
15 addresses, and such an address mapping may be generated based, at least in part, on the plurality of partition boundary identifiers and the plurality of wear leveling identifiers.

[0046] Method 500 may proceed to operation 504 during which an erase command may be received. In various embodiments, the erase command may be associated with a logical address included in a logical sector. Accordingly, such an erase command may be received
20 from a component, such as a processor of the memory, and may be received during operation of the memory device. Thus, operation 502 may be implemented during a configuration phase of the memory device, and operation 504 may occur subsequently and during operation of the memory device itself.

[0047] Method 500 may proceed to operation 506 during which it may be determine if the
25 logical address associated with the erase command is included in a wear leveling pool. If it is determined that the logical address associated with the erase command is not included in a wear leveling pool, method 500 may proceed to operation 512. However, if it is determined the logical address associated with the erase command is included in a wear leveling pool, method 500 may proceed to operation 508. In various embodiments, such a determination

may be made based on a comparison of the logical address with the generated partition boundary identifiers and wear leveling identifiers discussed above with reference to operation 502. Accordingly, the logical address associated with the erase command may be compared with the address ranges specified by the partition boundary identifiers, and a partition that includes the logical address may be identified. Furthermore, wear leveling identifiers may be looked up based on the identified partition, and wear leveling characteristics or settings may be identified. For example, the wear leveling identifiers may indicate that the identified partition that includes the logical address associated with the erase command is included in the wear leveling pool and should potentially be wear leveled.

5 [0048] If it is determined that the logical address is included in a wear leveling pool, method 500 may proceed to operation 508 during which it may be determined if a physical sector associated with the logical address should be wear leveled. If it is determined that the physical sector associated with the logical address should not be wear leveled, method 500 may proceed to operation 512. However, if it is determined that the physical sector associated with the logical address should be wear leveled, method 500 may proceed to operation 510.

[0049] In various embodiments, such a determination may be made based on a comparison of a number of PE cycles associated with the physical sector and a designated threshold value. In various embodiments, a component, such as a processor or control circuitry, may keep track of and count a number of PE cycles applied to each physical sector during operation of the memory device. In some embodiments, the PE cycle count associated with the physical sector identified during operations 504 and 506 may be looked up, and the PE cycle count may be compared with a threshold value. In various embodiments, the threshold value may be a designated value that has been specified by a user. As previously discussed, a user may specify a threshold value for a partition to identify when wear leveling is applied to sectors included in that partition. Moreover, as also discussed above, different threshold values may be utilized to enable different amounts of wear leveling being applied to different partitions. In some embodiments, the threshold value may be a default value set by a manufacturer. Thus, if it is determined that the PE cycle count does not exceed the threshold value, it may be determined that wear leveling should not be applied, and method 500 may proceed to operation 512. Moreover, if it is determined that the PE cycle count does exceed

the threshold value, it may be determined that wear leveling should be applied, and method 500 may proceed to operation 510.

[0050] Accordingly, if it is determined that the logical address is included in a wear leveling pool and that wear leveling should be applied, method 500 may proceed to operation 510

5 during which one or more wear leveling operations may be implemented. In various embodiments, such wear leveling operations may include swapping the physical sector with another physical sector in the wear leveling pool that currently has the lowest number of PE cycles. Such swapping may include swapping of data between physical sectors, and also updating of their associated logical address mapping, as discussed in greater detail below. In
10 this way, during operation of the memory device, PE cycles may be tracked for the physical sectors included in the memory device, including all those undergoing wear leveling, and data may be moved when PE cycle thresholds are met. As previously discussed, during such a wear leveling operation, the physical address of data is changed, typically to a location that has undergone fewer PE cycles, and the logical address remains the same.

15 **[0051]** Accordingly, during operation 510, an address mapping of logical addresses to physical addresses may be updated. In various embodiments, the control circuitry may have previously generated a third data structure that includes a mapping of the logical sectors and/or partitions to physical sectors. Such a mapping may have been generated during operation 502, and may have also been generated based on the identified wear leveling
20 characteristics specifying whether or not a logical sector is included in a wear leveling pool. For example, if a logical address and associated partition have a wear leveling identifier or characteristic that indicates it should undergo wear leveling, that partition, as well as the logical sectors included in the partition, are mapped to physical sectors identified as being included in the wear leveling pool. Such mapping may randomly assign the logical sectors to
25 physical sectors, or such assignment may occur sequentially. Similar mapping may also be implemented for partitions and logical sectors not undergoing wear leveling. In this way, a component, such as the control circuitry, may store and maintain a mapping of logical addresses to physical addresses, and such a mapping complies with the wear leveling characteristics initially specified by the user.

[0052] Accordingly, as stated above, this mapping may be updated based on the one or more wear leveling operations. More specifically, the mapping of logical sectors to physical sectors that was generated during operation 502 may be updated every time data is moved as part of the wear leveling process when the memory device is in use. In this way the memory device may accurately track the physical storage location of data that is associated with a particular logical address.

[0053] Method 500 may proceed to operation 512 during which an erase operation may be implemented. Accordingly, a component, such as a processor, may implement an erase operation on the logical sector identified during operation 504 as well as its associated physical sector. In an example in which operation 512 is being implemented after wear leveling, the corresponding physical sector may be the new physical sector selected during wear leveling. Furthermore, a counter associated with the physical sector may be updated to accurately identify a number of PE cycles that have been applied to that physical sector.

[0054] Method 500 may proceed to operation 514 during which it may be determined whether or not the memory device should wait for additional erase commands. In some embodiments, such additional erase commands may occur during operation of the memory device. Accordingly, if it is determined that the memory device should wait for additional erase commands, as may be the case if use of the memory device continues, method 500 may return to operation 504. If it is determined that the memory device should not wait for additional erase commands, as may be the case if use of the memory device has ceased, method 500 may terminate.

[0055] Figure 6 illustrates an example of a configurable wear leveling memory system, configured in accordance with some embodiments. According to particular example embodiments, system 600 may be suitable for implementing various components described above, such as memory device 100. In various embodiments, system 600 may include processor 601 which may be configured to implement one or more processing operations. For example, processor 601 may be configured to implement wear leveling operations associated with memory array 102, and may be further configured to implement read and write

operations associated with memory array 102. System 600 may also include bus 615 which may be configured to enable communication between various components of system 600.

[0056] In various embodiments, system 600 may further include memory array 102 and control circuitry 128 which may be configured to implement the configurable wear leveling

5 discussed above. In various embodiments, control circuitry 128 may include one or more processors and dedicated memory that may be configured to implement the previously described generations of data structures and mappings. In this way, system 600 may have a dedicated processing unit, such as control circuitry 128, which may be configured to implement the previously described configurable wear leveling. Moreover, in some
10 embodiments, control circuitry 128 may be implemented in an application specific integrated circuit (ASIC), or may be implemented in reprogrammable logic of a field programmable gate array. In some embodiments, control circuitry 128 may be implemented in a programmable system or controller that further includes a non-volatile memory, such as a Programmable System On a Chip or PSoC™ controller, commercially available from
15 Cypress Semiconductor of San Jose, California. In various embodiments, one or more components of system 600 may be implemented on the same circuit die and in the same package. For example, control circuitry 128 and memory 102 may be implemented on the same circuit die. In some embodiments, they may be implemented on different dies and in different packages.

20 **[0057]** In various embodiments, communications interface 611 may be configured to send and receive data to other system components, or may be configured to send and receive

packets or data segments over a network. For example, communications interface 611 may be communicatively coupled to a user interface device via a bus, such as bus 615, or via a communications network. As discussed above, communications interface 611 may be
25 configured to receive data from such a user interface device which may be included in a data processing system or computer system. In various embodiments, such data may include an input identifying wear leveling parameters from a user. In various embodiments, communications interface 611 may be a device that supports Ethernet interfaces, frame relay interfaces, cable interfaces, and DSL interfaces. In addition, various very high-speed
30 interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, HSSI

interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM.

5 **[0058]** Although the foregoing concepts have been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. It should be noted that there are many alternative ways of implementing the processes, systems, and devices. Accordingly, the present examples are to be considered as illustrative and not restrictive.

CLAIMS

What is claimed is:

1. A device comprising:

5 a non-volatile memory comprising a plurality of physical memory portions, wherein each of the plurality of physical memory portions is configured to be mapped to a logical memory portion of a plurality of logical memory portions; and

wear leveling control circuitry configured to receive a plurality of wear leveling parameters, and further configured to determine a plurality of wear leveling

10 characteristics based on the plurality of wear leveling parameters received, wherein the plurality of wear leveling characteristics identify an amount of wear leveling allowed for each of the plurality of logical memory portions.

2. The device of claim 1, wherein the plurality of wear leveling parameters is user
15 specified.

3. The device of claim 1, wherein the plurality of wear leveling parameters is configured to further identify portions of the non-volatile memory allocated to each of the
20 plurality of wear leveling characteristics.

4. The device of claim 1, wherein the plurality of wear leveling characteristics comprises a first wear leveling setting indicating wear leveling should be applied, and further comprises a second wear leveling setting indicating wear leveling should not be
25 applied.

5. The device of claim 4, wherein the first wear leveling setting indicates a first amount of wear leveling should be applied, and wherein the plurality of wear leveling characteristics further comprises:

a third wear leveling setting indicating a second amount of wear leveling should be applied; and

a fourth wear leveling setting indicating a third amount of wear leveling should be applied.

5

6. The device of claim 1, wherein the plurality of wear leveling parameters identifies a number of program and erase (PE) cycles associated with each of the plurality of logical memory portions, and wherein the plurality of wear leveling characteristics identify an amount of wear leveling applied to each of the plurality of logical memory portions based on the number of PE cycles.

10

7. The device of claim 1, wherein the plurality of wear leveling parameters identifies a logical size of each partition of a plurality of partitions of the non-volatile memory, and wherein the plurality of wear leveling characteristics identify an amount of wear leveling allowed for each of the plurality of logical memory portions based on the size of each logical portion.

15

8. The device of claim 1, wherein the plurality of logical memory portions is partitioned into a plurality of fixed size partitions.

20

9. The device of claim 1, wherein the plurality of logical memory portions is partitioned into a plurality of flexible size partitions configured based on the plurality of wear leveling parameters.

25

10. A system comprising:

a non-volatile memory comprising a plurality of physical memory portions, wherein each of the plurality of physical memory portions is configured to be mapped to a logical memory portion of a plurality of logical memory portions;

wear leveling control circuitry configured to receive a plurality of wear leveling parameters, and further configured to determine a plurality of wear leveling characteristics based on the plurality of wear leveling parameters received, wherein the

30

plurality of wear leveling characteristics identify an amount of wear leveling allowed for each of the plurality of logical memory portions;

a communications interface configured to communicatively couple the wear leveling control circuitry with a data processing system; and

one or more processors configured to execute program and erase operations associated with the non-volatile memory.

11. The system of claim 10, wherein the plurality of wear leveling parameters is user specified.

12. The system of claim 10, wherein the plurality of wear leveling parameters is configured to further identify portions of the non-volatile memory allocated to each of the plurality of wear leveling characteristics.

13. The system of claim 10, wherein the plurality of wear leveling characteristics comprises a first wear leveling setting indicating wear leveling should be applied, and further comprises a second wear leveling setting indicating wear leveling should not be applied.

14. The system of claim 10, wherein the plurality of wear leveling parameters identifies a number of program and erase (PE) cycles associated with each of the plurality of logical memory portions, and wherein the plurality of wear leveling characteristics identify an amount of wear leveling applied to each of the plurality of logical memory portions based on the number of PE cycles.

15. The system of claim 10, wherein the plurality of logical memory portions is partitioned into a plurality of flexible size partitions configured based on the plurality of wear leveling parameters, and wherein the non-volatile memory and the wear leveling control circuitry are implemented on a same circuit die.

16. A method comprising:

generating, using wear leveling control circuitry, a plurality of partition boundary identifiers characterizing a plurality of logical address ranges associated with a plurality of partitions of a non-volatile memory comprising a plurality of physical memory portions, each of the plurality of physical memory portions being associated with a logical memory portion of a plurality of logical memory portions;

generating, using the wear leveling control circuitry, a plurality of wear leveling identifiers associated with the plurality of partitions, the wear leveling identifiers comprising wear leveling characteristics associated with the plurality of logical memory portions, the plurality of wear leveling characteristics identifying an amount of wear leveling allowed for each of the plurality of logical memory portions; and


generating, using the wear leveling control circuitry, an address map based, at least in part, on the plurality of wear leveling identifiers, the address map characterizing a first mapping of a first plurality of logical memory portions to a first plurality of physical memory portions included in wear leveling, and the address mapping further characterizing a second map of a second plurality of logical memory portions to a second plurality of physical memory portions excluded from wear leveling.


17. The method of claim 16, wherein the plurality of wear leveling identifiers is user specified.


18. The method of claim 16, wherein the plurality of wear leveling identifiers further identifies portions of the non-volatile memory allocated to each of the plurality of wear leveling characteristics.


19. The method of claim 16, wherein the plurality of wear leveling characteristics comprises a first wear leveling setting indicating wear leveling should be applied, and further comprises a second wear leveling setting indicating wear leveling should not be applied.

20. The method of claim 16, wherein the plurality of wear leveling identifiers identifies a logical size of each partition of a plurality of partitions of the non-volatile memory, and wherein the plurality of wear leveling characteristics identify an amount of wear leveling allowed for each of the plurality of logical memory portions based on the
- 5 size of each logical portion.

**FIG. 1**

**FIG. 2**

**FIG. 3**

**FIG. 4**




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT		International application No. PCT/US17/46142																		
A. CLASSIFICATION OF SUBJECT MATTER IPC: G06F 3/06 (2006.01) USPC: G06F 12/0246, 3/0679, 2212/2022, 2212/7211; G11C 16/102 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) U.S. : G06F 12/0246, 3/0679, 2212/2022, 2212/7211; G11C 16/102 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)																				
C. DOCUMENTS CONSIDERED TO BE RELEVANT <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 10%;">Category *</th> <th style="width: 70%;">Citation of document, with indication, where appropriate, of the relevant passages</th> <th style="width: 20%;">Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">X --- Y</td> <td>US 2006/0282610 A1 (Dariel et al) 14 December 2006, see entire document.</td> <td style="text-align: center;">1-3, 7, 9-12, 16-18, 20 ----- 5-6, 8, 14-15</td> </tr> <tr> <td style="text-align: center;">X --- Y</td> <td>US 2014/0075100 A1 (KANEKO et al.) 13 March 2014, see entire document.</td> <td style="text-align: center;">1, 4, 10, 13, 16, 19 ----- 5</td> </tr> <tr> <td style="text-align: center;">Y</td> <td>US 7,710,777 B1 (Montierth et al.) 04 May 2010, see entire document.</td> <td style="text-align: center;">6, 14-15</td> </tr> <tr> <td style="text-align: center;">Y</td> <td>US 2007/0174549 A1 (Gyl et al.) 26 July 2007, see entire document,</td> <td style="text-align: center;">8</td> </tr> <tr> <td style="text-align: center;">A</td> <td>US 2013/0036253 A1 (Baltar) 07 February 2013, see entire document.</td> <td style="text-align: center;">1-20</td> </tr> </tbody> </table>			Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	X --- Y	US 2006/0282610 A1 (Dariel et al) 14 December 2006, see entire document.	1-3, 7, 9-12, 16-18, 20 ----- 5-6, 8, 14-15	X --- Y	US 2014/0075100 A1 (KANEKO et al.) 13 March 2014, see entire document.	1, 4, 10, 13, 16, 19 ----- 5	Y	US 7,710,777 B1 (Montierth et al.) 04 May 2010, see entire document.	6, 14-15	Y	US 2007/0174549 A1 (Gyl et al.) 26 July 2007, see entire document,	8	A	US 2013/0036253 A1 (Baltar) 07 February 2013, see entire document.	1-20
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																		
X --- Y	US 2006/0282610 A1 (Dariel et al) 14 December 2006, see entire document.	1-3, 7, 9-12, 16-18, 20 ----- 5-6, 8, 14-15																		
X --- Y	US 2014/0075100 A1 (KANEKO et al.) 13 March 2014, see entire document.	1, 4, 10, 13, 16, 19 ----- 5																		
Y	US 7,710,777 B1 (Montierth et al.) 04 May 2010, see entire document.	6, 14-15																		
Y	US 2007/0174549 A1 (Gyl et al.) 26 July 2007, see entire document,	8																		
A	US 2013/0036253 A1 (Baltar) 07 February 2013, see entire document.	1-20																		
☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.																				
<table style="width: 100%;"> <tr> <td style="width: 50%; vertical-align: top;"> * 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 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 </td> <td style="width: 50%; vertical-align: top;"> "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 </td> </tr> </table>			* 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 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																
* 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 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 28 August 2017 (28.08.2017)		Date of mailing of the international search report 01 SEP 2017																		
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450 Alexandria, Virginia 22313-1450 Facsimile No. (571) 273-8300		Authorized officer Gail Hayes Telephone No. 571-272-4300																		

INTERNATIONAL SEARCH REPORT		International application No. PCT/US17/46142
C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2011/0066792 A1 (Shaeffer et al.) 17 March 2011, see entire document.	1-20