

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
2 April 2009 (02.04.2009)


PCT

(10) International Publication Number
WO 2009/042554 A1

- (51) **International Patent Classification:**
G11C 29/00 (2006.01) *G06F 11/20* (2006.01)
G06F 11/10 (2006.01) *G06F 11/16* (2006.01)
- (21) **International Application Number:**
PCT/US2008/077267
- (22) **International Filing Date:**
22 September 2008 (22.09.2008)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
11/861,146 25 September 2007 (25.09.2007) US
- (71) **Applicant (for all designated States except US):** SAN-DISK CORPORATION [US/US]; 601 McCarthy Blvd., Milpitas, California 95035-7932 (US).
- (72) **Inventor; and**
- (75) **Inventor/Applicant (for US only):** CHENG, Steven, S. [US/US]; 1121 South Bernardo Avenue, Sunnyvale, California 94087 (US).
- (74) **Agents:** GALLAGHER, Peter, A. et al.; WEAVER AUSTIN VILLENEUVE & SAMPSON LLP, P.O. Box 70250, Oakland, California 94612-0250 (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, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
 — with international search report
 — before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

(54) **Title:** NONVOLATILE MEMORY WITH SELF RECOVERY

FIG. 8B



(57) **Abstract:** A nonvolatile memory array includes two or more devices, each device containing data that is scrambled using a different scrambling scheme. When the same data is provided and stored in both devices, different data patterns occur in each device, so that if one of the patterns causes data pattern induced errors, the original data can be recreated from another copy that does not share the same data pattern.

WO 2009/042554 A1

NONVOLATILE MEMORY WITH SELF RECOVERY

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to and benefit of U.S. Patent Application No:
5 11/861,146, filed September 25, 2007, which is incorporated herein by reference in its
entirety for all purposes.

BACKGROUND OF THE INVENTION

10 [0002] This invention relates to nonvolatile memories and methods of forming
nonvolatile memories. In particular, this application relates to nonvolatile memory
arrays in which floating gate memory cells individually hold one or more bits of data.

[0003] Nonvolatile memory systems are used in various applications. Some
nonvolatile memory systems are embedded in a larger system such as a personal
15 computer. Other nonvolatile memory systems are removably connected to a host
system and may be interchanged between different host systems. Examples of such
removable memory systems include memory cards and USB flash drives. Electronic
circuit cards, including non-volatile memory cards, have been commercially
implemented according to a number of well-known standards. Memory cards are
20 used with personal computers, cellular telephones, personal digital assistants (PDAs),
digital still cameras, digital movie cameras, portable audio players and other host
electronic devices for the storage of large amounts of data. Such cards usually
contain a re-programmable non-volatile semiconductor memory cell array along with
a controller that controls and supports operation of the memory cell array and
25 interfaces with a host to which the card is connected. Several of the same type of card
may be interchanged in a host card slot designed to accept that type of card.
However, the development of the many electronic card standards has created different
types of cards that are incompatible with each other in various degrees. A card made
according to one standard is usually not useable with a host designed to operate with a
30 card of another standard. Memory card standards include PC Card, CompactFlash™
card (CF™ card), SmartMedia™ card, MultiMediaCard (MMC™), Secure Digital
(SD) card, a miniSD™ card, Subscriber Identity Module (SIM), Memory Stick™,

Memory Stick Duo card and microSD/TransFlash™ memory module standards. There are several USB flash drive products commercially available from SanDisk Corporation under its trademark “Cruzer®.” USB flash drives are typically larger and shaped differently than the memory cards described above. Solid State Drives (SSDs) use nonvolatile memory systems in applications where Hard Disk Drives (HDDs) have traditionally been used, such as in laptop computers.

[0004] Different types of memory array architecture are used in nonvolatile memory systems. Flash memory arrays that use floating gates to store charge are one common type of memory array. In one type of architecture, a NAND array, a series of strings of more than two memory cells, such as 16 or 32, are connected along with one or more select transistors between individual bit lines and a reference potential to form columns of cells. Word lines extend across cells within a large number of these columns. An individual memory cell may hold one bit of data in what is known as a Single Level Cell (SLC) design. In some examples, a memory cell may hold two or more bits of data in what is known as a Multi Level Cell (MLC) design.

[0005] Reliability is generally an important feature in memory products, and is especially important for certain industrial applications. Thus, there is a need for nonvolatile memory systems with high reliability.

SUMMARY OF THE INVENTION

[0006] A combination of Redundant Array of Independent Disk (RAID) techniques and data scrambling can provide high reliability in nonvolatile memory arrays. In particular, where different data scrambling is used in different devices that provide redundancy, each device contains a different data pattern, even though the original (unscrambled) data is the same for each device. Thus, any errors induced by a particular data pattern in a device is not repeated in multiple devices because of the different scrambling used. A RAID system using such scrambling incorporates protection against data pattern induced errors as well as device failure.

[0007] A method of securely storing data in a nonvolatile memory array according to an embodiment of the present invention comprises: transforming a portion of data

according to a first transformation to obtain first transformed data; storing the first transformed data in a first portion of the memory array; transforming the portion of data according to a second transformation to obtain second transformed data; and storing the second transformed data in a second portion of the memory array while the first transformed data is stored in the first portion of the memory array, to provide redundant storage of the portion of data.

[0008] A nonvolatile memory system according to an embodiment of the present invention comprises: a first portion of a memory array; a first data scrambler unit connected to the first portion of the memory array to scramble data according to a first transformation prior to storage in the first portion of the memory array; a second portion of the memory array; and a second data scrambler unit connected to the second portion of the memory array to scramble data according to a second transformation prior to storage in the second portion of the memory array, the second data scrambler unit connected in parallel with the first data scrambler unit.

15 BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 shows a memory system using mirroring, where all data stored in a first device is also stored in a second device.

[0010] FIG. 2 shows the memory system of FIG. 1 where data A2, because of its pattern when stored, is uncorrectable by ECC when it is read. Both copies of data A2, in device 1 and device 2, are uncorrectable.

[0011] FIG. 3 shows data being stored using striping to distribute the data across four devices.

[0012] FIG. 4 shows an alternative striping arrangement with larger stripes.

[0013] FIG. 5A shows a memory system using mirroring and striping of data.

25 [0014] FIG. 5B shows the memory system of FIG. 5A with data A2 and A7 being uncorrectable by ECC when read because of data pattern induced errors.

[0015] FIG. 6A shows an alternative memory system using mirroring and striping.

- [0016] FIG. 6B shows the memory system of FIG. 6A, with data A2 being uncorrectable by ECC when read because of data pattern induced errors.
- [0017] FIG. 7A shows a memory system using mirroring and striping that has a data scrambler unit to transform all data before the data is stored in the memory system.
- 5 [0018] FIG. 7B shows the memory system of FIG. 7A with transformed data A2' and A3' being uncorrectable by ECC when read because of data pattern induced errors.
- [0019] FIG. 8A shows an alternative memory system having a first data scrambler unit for device 1 and a second data scrambler unit for device 2 so that data in each device is subject to a different transformation.
- 10 [0020] FIG. 8B shows transformed data A1' being uncorrectable by ECC when read because of data pattern induced errors, while transformed data A1* is still correctable.
- [0021] FIG. 9A shows an alternative arrangement with data striping between devices, where data is striped across devices that share a data scrambler.
- [0022] FIG. 9B shows the memory system of FIG. 9A with transformed data A2' and
15 A7* being uncorrectable by ECC because of data pattern induced errors, while transformed data A2* and A7' remain correctable by ECC.
- [0023] FIG. 10A shows another memory system where some devices have dedicated data scramblers and some devices share a data scrambler.
- [0024] FIG. 10B shows the memory system of FIG. 10A with transformed data A2'
20 and A3" being uncorrectable by ECC when read because of data pattern induced errors, while transformed data A2*, A2", A3' and A3* remain correctable by ECC when read.
- [0025] FIG. 10C shows replacement of uncorrectable transformed data A2' using data A2*, and replacement of uncorrectable transformed data A3" using data A3*.
- 25 [0026] FIG. 11A shows a set of 32 keys used sequentially to transform 32 consecutive pages before the keys are repeated.

[0027] FIG. 11B shows 32 keys used to transform data that is striped across four devices, with keys repeating within each device after 8 pages.

[0028] FIG. 12A shows data stored using a single data scrambler unit that transforms data using a set of keys that change for each page of data.

5 [0029] FIG. 12B shows an alternative system with separate data scrambler units for each device, with each data scrambler using keys sequentially so that keys are not frequently repeated within a device.

[0030] FIG. 13 shows a memory system that uses a dedicated device to store parity data, and which has dedicated data scrambler units for each device.

10 [0031] FIG. 14 shows a memory system in which parity data is distributed across four devices, with each device having a dedicated data scrambler unit.

DETAILED DESCRIPTION OF ILLUSTRATED EMBODIMENTS

[0032] Various methods are used to achieve high reliability data storage in nonvolatile memory systems. Certain methods that improve reliability through
15 redundancy use a Redundant Array of Independent Disks (RAID). There are different types of RAID systems that use mirroring, striping, and parity data, or some combination of these techniques. Although the term RAID is generally associated with Hard Disk Drives (HDDs), RAID techniques may be used with other types of memory including nonvolatile memory arrays such as flash memory.

20 [0033] FIG. 1 shows two devices, device 1 and device 2, which are connected in a mirroring arrangement (also referred to as RAID 1). Mirroring means that the same data that is stored in device 1, A0, A1, A2, A3, is also stored in device 2 so that there is a redundant copy of each unit of data in device 2. In other examples, more than two devices may be connected in this way to provide more than two copies of the same
25 data. Generally, maintaining a separate copy of the same data in two or more devices provides protection against failure of any individual device. Thus, if device 1 fails in FIG. 1, all the data stored in device 1 can be recovered from device 2 instead.

[0034] While a redundant copy of data in a separate device protects against device failure, it does not protect against all possible forms of data loss. In particular, in certain nonvolatile memory arrays, particular patterns of data storage may increase the chances of data becoming corrupted. This may occur because of the interaction
5 between adjacent cells, with floating gates of cells affecting not only their underlying channels, but also affecting neighboring channels and thus affecting threshold voltages of neighboring cells. This is especially a problem for MLC memory systems in which small changes in a cell's threshold voltage may result in one or more bits stored in the cell being incorrectly read. The problem tends to become more severe as
10 dimensions in memory arrays become smaller and interactions between cells become stronger. For example, an MLC cell that is initially programmed to a low threshold voltage (little or no charge on its floating gate) may later have its threshold voltage raised because neighboring cells are programmed to high threshold voltages (large charges on their floating gates). This may be considered a form of program disturb.
15 Which specific data patterns cause program disturb, and the degree of disturb caused (the number of bits that are later erroneously read) depends on the memory architecture. Thus, there may be particular patterns of data that will cause data stored in particular memory arrays to be read erroneously even though the memory array is working within its design limits. Avoiding such data patterns is generally desirable.

[0035] Many memory systems include some form of Error Correction Coding (ECC) in order to allow data containing some erroneous bits to be detected and corrected. In order to do this, the data is stored with some redundancy. When the data is read, the redundancy allows the data to be checked for errors, and up to some limited number of errors may be corrected. However, in order to correct a large number of errors, a
25 high degree of redundancy is needed (a large number of extra bits) which is not an efficient use of memory space. Typically, the degree of redundancy used is based on expected error rates (e.g. based on a number of expected defective cells) and may not be sufficient to correct a high number of errors, such as caused by a particular data pattern. In particular, data pattern induced errors may occur together in the same
30 portion of data in cells that are exposed to similar conditions, and ECC is generally limited to correcting only a limited number of bits in such a portion of data. Thus, ECC alone may not be sufficient to overcome data pattern induced errors in certain

memory systems.

[0036] FIG. 2 shows an example where device 1 and device 2 store identical copies of the same data. In this case, a portion of data A2 has a data pattern such that when A2 is stored and later read, the data that is read has a high number of errors because of
5 interaction between cells. This may include interaction of cells within A2 and interaction with cells outside A2. In particular, cells in a row tend to be affected by cells in adjacent rows of the same block. In this case, the number of errors in data A2 is beyond the correcting capability of the ECC scheme being used and the data is considered to be uncorrectable by ECC. Data portion A2 may be any unit of data,
10 such as a sector of 512 bytes (plus overhead), a page (unit of writing) or a block (minimum unit of erase). Data A2 is stored identically in both device 1 and device 2, so the same pattern occurs in each device. Therefore, copies of data A2 in both devices are affected and neither copy is recoverable. Thus, in this case, mirroring alone does not provide a sufficient level of reliability to enable all data to be
15 recovered.

[0037] FIG. 3 shows a first example of data striping between different devices (RAID 0). Data is received as three files, File A, File B and File C as shown. These files are then spread out across devices 1-4 as shown. In particular, File A is stored as A0 in device 1 and A1 in device 2. File B is stored as B0-B6, which are spread out across
20 devices 1-4. And File C is stored as C0-C9, which are spread out across devices 1-4. A stripe of data in this example consists of one unit of data in each device, e.g. A0, A1, B0 and B1 together make up a stripe. Each stripe contains 2KBytes of data in this example. FIG. 4 shows another example in which the same data is striped across the same devices using a larger stripe size where each stripe contains 4KBytes of data.
25 Although striping does not provide any redundancy, it can improve reliability because data is broken up and data patterns that might otherwise induce errors may be avoided. Striping data may improve performance because it facilitates programming with a high degree of parallelism across multiple devices.

[0038] FIG. 5A shows an example of a memory system using a combination of
30 mirroring and striping. Data is striped between device 1 and device 2, and the data in

devices 1 and 2 is mirrored in devices 3 and 4. This hybrid arrangement applying aspects of RAID 0 and RAID 1 may be referred to as RAID 0+1. As shown in FIG. 5B, this arrangement still suffers problems when certain patterns occur in the stored data. In particular, if storage of data A2 in device 1 causes an uncorrectable number of errors due to the data pattern, then the copy of data A2 stored in device 3 also contains an uncorrectable number of errors. Similarly, the two copies of data A7, in device 2 and device 4, are both uncorrectable, so there is no recoverable copy of either A2 or A7.

[0039] FIG. 6A shows an alternative combination of RAID 0 and RAID 1, which may be referred to as RAID 10. In this arrangement data in device 1 is mirrored in device 2. The same data that is stored in devices 1 and 2 is striped across device 3 and device 4. However, as shown in FIG. 6B, where one copy of data A2 is uncorrectable because of its data pattern, then the other copies of data A2 in other devices are also uncorrectable.

[0040] One approach to data patterns that may cause data to be uncorrectable when it is read is to scramble data before storage. Scrambling subjects data to some transformation so that the transformed data has a different pattern to untransformed data. In particular, such scrambling may introduce some randomization into a repeated pattern in the data. A simple example of a transformation is to invert (or flip) all bits, i.e. change all zeros to ones and change all ones to zeros. Other transformations involve performing exclusive OR (XOR) operations on data to be stored and some series of bits. Unlike ECC, scrambling does not add redundant bits, so the scrambled data includes the same number of bits as the original data. In some cases, a transformation may compress data, reducing the number of bits to be stored. Some transformations use data security algorithms that may break up data patterns.

[0041] FIG. 7A shows an example where a data scrambler unit, or Data Scrambler (DS), is connected to devices 1-4 so that all data to be stored in devices 1-4 is scrambled before storage. Scrambling is indicated in FIG. 7A by an apostrophe so that A0 is transformed to become A0' etc. FIG. 7A shows a RAID 10 arrangement as in FIG. 6A, with data mirrored between device 1 and device 2, and the same data striped

across device 3 and device 4. While scrambling data may result in transformed data that has a pattern that does not cause disturbs, in some cases the transformed data will have a pattern that causes errors when reading. In some cases, data that without scrambling would not cause errors is transformed into a pattern that does cause errors.

5 FIG. 7B shows the result of scrambling where transformed data A2' and A3' have patterns that cause errors when read. Because these patterns cause errors, and the patterns are replicated in each copy of the data, no recoverable copy of the data is available. Thus, scrambling using a single data scrambler unit, in addition to mirroring and striping is not sufficient to ensure reliability in this case.

10 [0042] FIG. 8A shows a data storage system according to an embodiment of the present invention. The memory system receives a sequence of units of data A0, A1, A2, A3... which are transformed by the first data scrambler DS' according to a first transformation. The data transformed by DS' (A0', A1', A2', A3'...) is stored in device 1. Data A0, A1, A2, A3... are also transformed by a second data scrambler DS* (into
15 A0*, A1*, A2*, A3*...) and stored. Thus, data A0 is transformed into A0' for storage in device 1 and data A0 is transformed into A0* for storage in device 2. Data A0' has a different pattern to data A0*. Thus, the threshold voltages of cells of device 1 storing data A0' may not be the same as corresponding cells in device 2 that store data A0*. Similarly, for other portions of data, different copies undergo different
20 transformations before storage so that the same data pattern is not reproduced. As discussed above, sometimes scrambled data may have a pattern that causes the stored data to be uncorrectable. FIG. 8B shows transformed data A1' being uncorrectable due to its data pattern when stored. However, transformed data A1* has a different data pattern when stored, so transformed data A1* is not uncorrectable. Thus, even
25 though data A1' is not recoverable, a copy of original (untransformed) data A1 can be obtained from data A1*. Similarly, if any other unit of data was to be uncorrectable due to its data pattern when transformed and stored, the original data could still be recovered because the data was also subject to another transformation prior to storage in another device. Thus, the arrangement of FIG. 8B provides protection against
30 device failure (i.e. if either device 1 or device 2 fails) and data pattern induced errors. While this method is similar to mirroring in that the data is stored in two independent devices, the two devices hold different data pattern as a result of scrambling. Thus,

the data in device 2 does not mirror the data in device 1 and this is not true mirroring.

[0043] FIG. 9A shows another example using RAID 0+1 with two different scrambler units DS', DS* performing two different transformations of data. Data scrambler DS' serves device 1 and device 2, while data scrambler DS* serves device 3 and device 4. Thus, all data stored in device 1 and device 2 is subject to a first transformation, while the same data is subject to a second transformation before storage in device 3 and device 4. Data is striped between device 1 and device 2, and data is also striped between device 3 and device 4. In other examples, banks of more than two devices may share a data scrambler and have data striped in this manner. FIG. 9B shows data pattern induced errors occurring in transformed data A2' and in transformed data A7*. Original data A2 may be recovered by reading A2* and transforming it (reversing the transformation that created A2*). Because A2* has a different data pattern than A2' when stored, it is unlikely that both patterns induce errors. Similarly, even though A7* is unrecoverable, original data A7 can be recreated from A7'. The striping used in this example may allow a greater degree of parallelism during programming, which improves performance. While a different data scrambler unit may be provided for each device (i.e. separate data scrambler for devices 1, 2, 3, and 4), it is generally sufficient to provide one data scrambler for a group of devices that store striped data because such devices generally do not contain more than one copy of any particular portion of data.

[0044] FIG. 10A shows an alternative memory system that has a RAID 10 configuration. Data received by the memory system is subject to a first transformation by data scrambler DS' before it is stored in device 1. The received data is also subject to a second transformation by data scrambler DS* before being stored in device 2. The received data is also subject to a third transformation by data scrambler DS'' before being stored in device 3 and device 4. The transformed data is striped across device 3 and device 4. Thus, in this example, if any transformation results in data that induces errors when the data is read, there are two other locations from which the data can be recovered. For example, FIG. 10B shows data A2' and data A3'' being unrecoverable because of data pattern induced errors. However, for each of these units of data, there are two other versions (A2*, A2'', A3' and A3*) from which the original

data can be recovered.

[0045] FIG. 10C shows how portions of data A2' and A3" of FIG. 10B that are uncorrectable by ECC when read may be replaced by data generated from another device. For example, when A2' in device 1 is found to be uncorrectable by ECC, original data A2 is recovered from data A2* stored in device 2. The original data A2 may also be subject to yet another transformation to provide data A2**. This transformation may be performed by modifying the transformation performed by data scrambler DS' to be different to the earlier transformation. After this replacement, there are three devices from which original data A2 can be recovered. Similarly, data A3" in device 4 is uncorrectable by ECC, but original data A3 is recovered from data A3* in device 2. Original data A3 is then transformed to provide data A3**, which replaces uncorrectable data A3". The transformation of data A3 to A3** is performed by data scrambler DS", which is modified to perform a different transformation than was performed for previously stored data. In this example, data A2** and A3** are the result of a similar transformation. However, in other cases, different transformations may be used in different devices when performing replacement.

[0046] In some memory systems, data transformation is performed according to a key. For example transformed data may be the product of an XOR operation on the original data and a key. To recover the original data, the operation is reversed. The key may be changed periodically (e.g. changed for each sector, each page, each block or some other pattern). Generally a limited number of keys are used, so that there is a repetitive pattern of key usage. However, where such a system is used with data striping, only a fraction of the total number of keys may be used in any one device resulting in the key pattern repeating sooner than expected. FIG. 11A shows device 1 in which 32 keys (key 0 to key 31) are used to scramble data within a memory system having four devices (device 1 – device 4), the key changing from page to page. With 32 keys, the pattern of key use repeats after 32 pages. However, if this scheme is applied in a memory system using striping as shown in FIG. 11B, then the key repeats much sooner. In this case, after just 8 pages. It is generally undesirable for the pattern to repeat so soon because the risk of memory cells interacting is increased if the keys repeat sooner.

[0047] FIG. 12A shows another example where a single data scrambler DS uses a limited number of keys to scramble data before the data is striped across devices 1-4. As can be seen, the keys used in each device are limited to a particular subset of keys, and the keys repeat within a device sooner as a result.

5 [0048] FIG. 12B shows an alternative configuration where a different data scrambler is used for each device, with each data scrambler using a full set of keys so that all keys are used in each device. This increases the interval between reuse of the same key and so lowers the chances of memory cell states aligning in a manner that would result in data corruption. Thus, even where no redundancy is provided, using different
10 scrambling for each device may be advantageous. FIG. 12B shows the same pattern of key use by each data scrambler. In other examples, different data scramblers may use keys in different orders. For example, each data scrambler may start with a different key (e.g. data scrambler DS1 may start with key 0, data scrambler DS2 may start with key 8, data scrambler DS3 may start with key 16, and DS4 may start with key 24).
15 Each scrambler may have an initial key, or seed, from which subsequent keys are derived. One way to ensure different transformations for data in different devices is to initialize their respective data scramblers using different seeds. In other examples, data scramblers may be controlled so that different keys are used by each data scrambler when scrambling the same data.

20 [0049] FIG. 13 shows an example of maintaining parity data in a dedicated device (RAID 4). In this case, parity data is stored exclusively in device 4. Parity data in device 4 is derived from data stored in devices 1-3. Parity data may be derived before or after data is scrambled. Parity information adds an extra level of protection against failure by allowing data in one device to be reconstructed from data in other devices
25 including parity data. Generally, using parity data increases redundancy but also incurs a performance penalty because of the time to perform parity operations whenever data is written or read. Each device has a dedicated data scrambler, so that data scrambling keys are not frequently repeated within a device.

[0050] FIG. 14 shows an example where parity data is distributed across all devices
30 (RAID 5). In this example, each stripe includes parity data derived from the other data

in the stripe. For example, Parity0 is derived from A0, A1 and B0. Parity1 is derived from B1, B2 and B3. This example also shows dedicated data scramblers for each device.

5 [0051] Certain examples show a dedicated data scrambler unit attached to an individual device. This may be achieved through hardware by having a dedicated circuit (e.g. a peripheral circuit on a memory die). In other examples, data scrambling may be achieved by firmware in a controller so that the same hardware is used for different circuits, but the firmware ensures that different transformations are performed according to the destination device. Other hardware may also provide
10 different scrambling to more than one device at a time. Thus, different scrambling schemes may be achieved without necessarily requiring different hardware for each device.

[0052] Various schemes described above may be combined with, or used to replace, other schemes that provide high reliability. For example, the schemes are generally
15 combined with an ECC scheme that provides protection against failure due to a small number of errors in a portion of data. In general, the schemes described above are performed within a memory system (e.g. a removable memory card or USB flash drive) and may not be visible to a host. In some cases, the degree of protection, or a particular protection scheme, may be selected by a user, through the host, or
20 otherwise. In some cases, the host system may have its own protection scheme that is in addition to any protection scheme in the memory system.

[0053] All patents, patent applications, articles, books, specifications, other publications, documents and things referenced herein are hereby incorporated herein by this reference in their entirety for all purposes. To the extent of any inconsistency
25 or conflict in the definition or use of a term between any of the incorporated publications, documents or things and the text of the present document, the definition or use of the term in the present document shall prevail.

[0054] Although the various aspects of the present invention have been described with respect to certain preferred embodiments, it is understood that the invention is
30 entitled to protection within the full scope of the appended claims.

WHAT IS CLAIMED IS:

1. A method of securely storing data in a nonvolatile memory array comprising:
transforming a portion of data according to a first transformation to obtain first
5 transformed data;
storing the first transformed data in a first portion of the memory array;
transforming the portion of data according to a second transformation to
obtain second transformed data; and
storing the second transformed data in a second portion of the memory array
10 while the first transformed data is stored in the first portion of the memory array, to
provide redundant storage of the portion of data.
2. The method of claim 1 wherein the first transformation scrambles data based
on a first seed and the second transformation scrambles data based on a second seed.
15
3. The method of claim 1 wherein transforming the portion of data according to
the first transformation is performed in parallel with transforming the portion of data
according to the second transformation.
- 20 4. The method of claim 1 further comprising, subsequently, in response to
attempting to read the first transformed data from the first portion of the memory
array and determining that the first transformed data is uncorrectable by ECC, reading
the second transformed data from the second portion of the memory array.
- 25 5. The method of claim 4 further comprising, using the second transformed data
to recover the portion of data and subsequently transforming the portion of data
according to a third transformation to obtain third transformed data, storing the third
transformed data in a third portion of the memory array and erasing the stored first
transformed data.
30
6. The method of claim 1 wherein the first portion of the memory array is a first
block and the second portion of the memory array is a second block.

7. The method of claim 1 wherein the first portion of the memory array is a first plane and the second portion of the memory array is a second plane.
8. The method of claim 1 wherein the first portion of the memory array is a first
5 die and the second portion of the memory array is a second die.
9. A nonvolatile memory system comprising:
a first portion of a memory array;
a first data scrambler unit connected to the first portion of the memory array to
10 scramble data according to a first transformation prior to storage in the first portion of
the memory array;
a second portion of the memory array; and
a second data scrambler unit connected to the second portion of the memory
array to scramble data according to a second transformation prior to storage in the
15 second portion of the memory array, the second data scrambler unit connected in
parallel with the first data scrambler unit.
10. The nonvolatile memory system of claim 9 wherein the first scrambler unit
contains a first exclusive OR (XOR) circuit to perform XOR operations using a first
20 number and the second scrambler contains a second XOR circuit to perform XOR
operations using a second number.
11. The nonvolatile memory system of claim 9 wherein the first memory array
portion is a first block and the second memory array portion is a second block.
25
12. The nonvolatile memory system of claim 9 wherein the first memory array
portion is a first plane and the second memory array portion is a second plane.
13. The nonvolatile memory system of claim 9 wherein the first memory array
30 portion is a first die and the second memory array portion is a second die.

14. The nonvolatile memory system of claim 9 wherein the memory system is in a removable memory card.

15. The nonvolatile memory system of claim 9 wherein the first scrambler unit
5 and the second scrambler unit scramble the same data in parallel.




FIG. 1




FIG. 2




FIG. 4




FIG. 5A




FIG. 5B




FIG. 6A




FIG. 6B




FIG. 7A




FIG. 7B




FIG. 8A





FIG. 8B



7 / 14

FIG. 9A





FIG. 9B



8 / 14

FIG. 10A

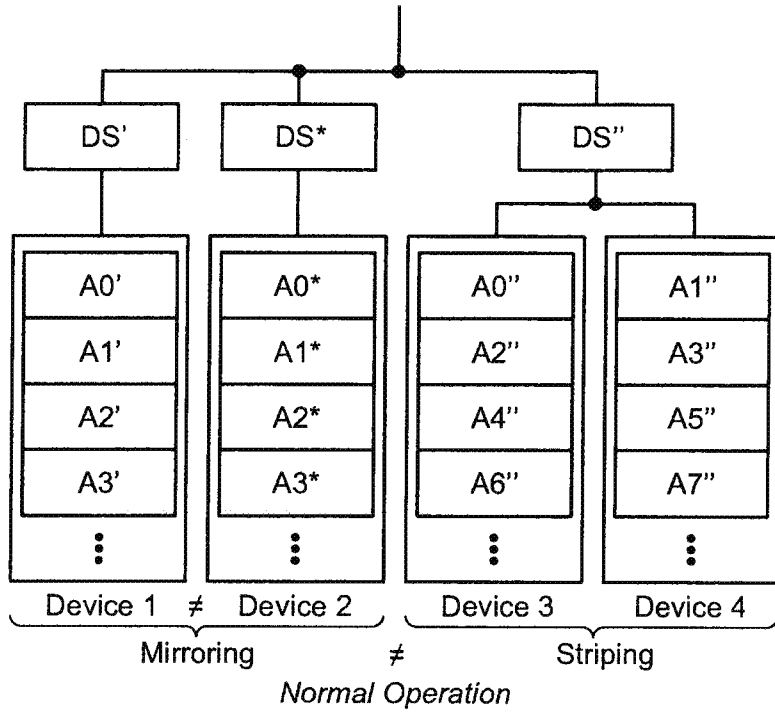
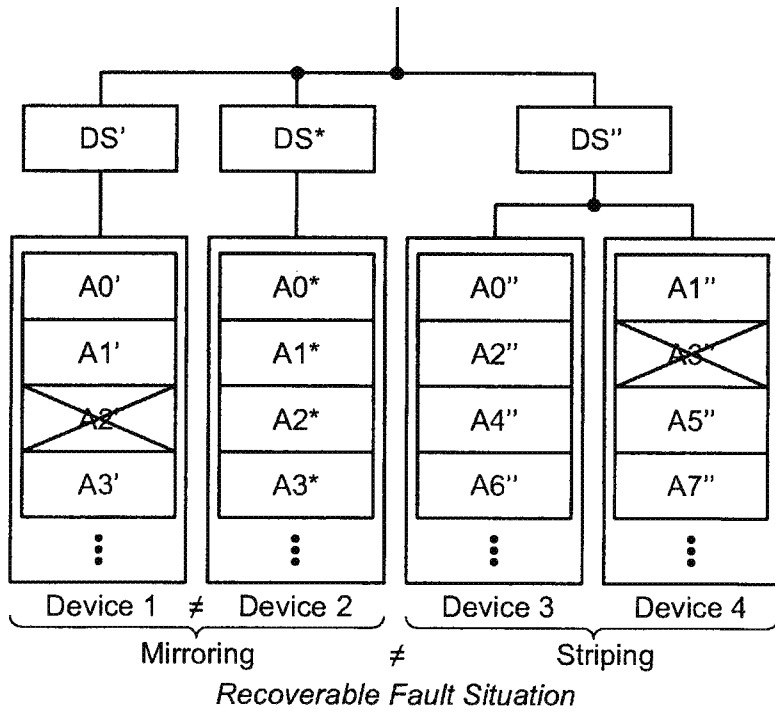


FIG. 10B



9 / 14

FIG. 10C





FIG. 11A



32 Keys Applied to 8 Pages in Each Device

FIG. 11B

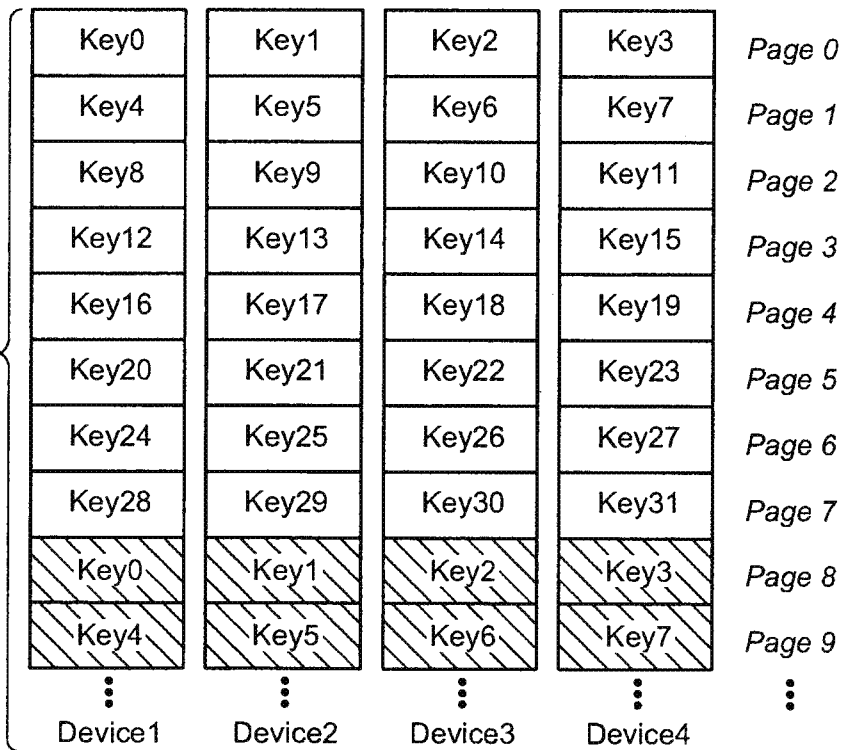


FIG. 12A




FIG. 12B





FIG. 13






FIG. 14

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2008/077267

A. CLASSIFICATION OF SUBJECT MATTER

INV. G11C29/00 G06F11/10 G06F11/20 G06F11/16

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)
G11C G06F

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

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

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 246 066 A (CONTI TEMIC MICROELECTRONIC [DE]) 2 October 2002 (2002-10-02) paragraph [0013]	1-3,6-15
Y		4
A		5
Y	EP 1 703 400 A (FUJITSU LTD [JP]) 20 September 2006 (2006-09-20) page 3, line 19 - line 24	4
A		1-3,5-15
A	EP 0 709 782 A (HEWLETT PACKARD CO [US]) HEWLETT PACKARD CO 1 May 1996 (1996-05-01) abstract	1-15
A	DE 36 19 174 A1 (FREIBERG BRENNSTOFFINST [DD]) 25 June 1987 (1987-06-25) abstract	1-15

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

4 February 2009

Date of mailing of the international search report

12/02/2009

Name and mailing address of the ISA/

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

Authorized officer

Ríos Báez, Abel

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2008/077267

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1246066	A	02-10-2002 DE 10113319 A1	26-09-2002
EP 1703400	A	20-09-2006 CN 1834930 A JP 2006260160 A KR 20060101151 A US 2006212656 A1	20-09-2006 28-09-2006 22-09-2006 21-09-2006
EP 0709782	A	01-05-1996 DE 69508034 D1 DE 69508034 T2 JP 3748117 B2 JP 8234922 A	08-04-1999 19-08-1999 22-02-2006 13-09-1996
DE 3619174	A1	25-06-1987 DD 245068 A1	22-04-1987