



(51) International Patent Classification:  
*G11C 16/34* (2006.01)

(21) International Application Number:  
PCT/EP2010/053943

(22) International Filing Date:  
25 March 2010 (25.03.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
12/420,628 8 April 2009 (08.04.2009) US

(71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).

(71) Applicant (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; P.O. Box 41, North Harbour, Portsmouth, Hampshire PO6 3AU (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **KABELAC, William, John** [US/US]; IBM Corporation, Almaden Research Centre, Intellectual Property Law C4TA/J2B, 650 Harry Road, San Jose, California 95120-6099 (US). **HETZLER, Steven, Robert** [US/US]; IBM Corporation, MD K01/E2, 650 Harry Road, Almaden Research Center, San Jose, California 95120-6099 (US).

(74) Agent: **LITHERLAND, David, Peter**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester, Hampshire SO21 2JN (GB).

(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, CL, 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, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD,

[Continued on next page]

(54) Title: ANALYZING MONITOR DATA INFORMATION FROM MEMORY DEVICES HAVING FINITE ENDURANCE AND/OR RETENTION

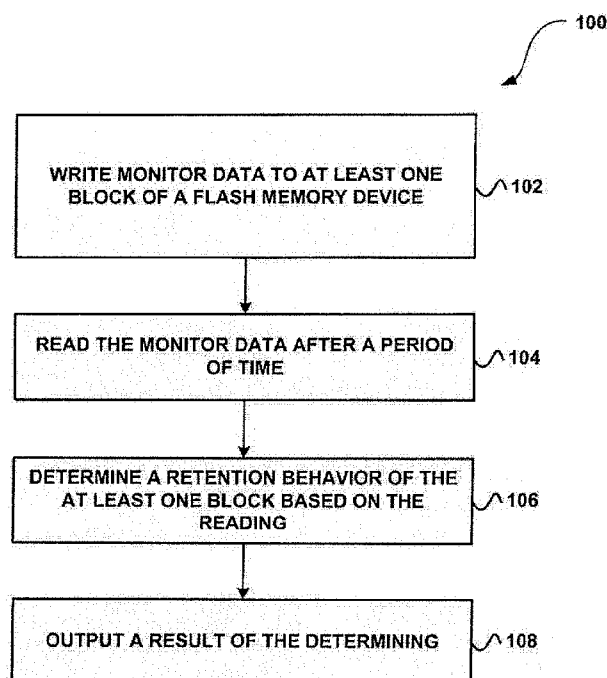


FIG. 1

(57) Abstract: A method according to one embodiment includes gathering monitor data information from a plurality of memory devices having finite endurance and/or retention, the monitor data being data of known content stored in dedicated memory cells of known write cycle count; analyzing the monitor data information; and taking an action relating to at least one of the devices based on the analyzing. Additional systems, methods, and computer program products are also disclosed.



SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**(84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, 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,

**Published:**

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

## **ANALYZING MONITOR DATA INFORMATION FROM MEMORY DEVICES HAVING FINITE ENDURANCE AND/OR RETENTION**

### **FIELD OF THE INVENTION**

5 The present invention relates generally to memory devices, and more particularly to analyzing monitor data information from memory devices having finite endurance and/or retention.

### **BACKGROUND OF THE INVENTION**

10 Some memory devices such as NAND flash memory devices possess finite write/erase endurance and/or limited data retention. Tracking the write/erase endurance and data retention is often desired to recognize a current usefulness and reliability of such memory.

15 To date, however, proposed solutions have failed to provide acceptable solutions for these and/or other issues associated with such memory devices.

### **SUMMARY OF THE INVENTION**

20 A method according to one embodiment includes gathering monitor data information from a plurality of memory devices having finite endurance and/or retention, the monitor data being data of known content stored in dedicated memory cells of known write cycle count; analyzing the monitor data information; and taking an action relating to at least one of the devices based on the analyzing.

25 A computer program product according to one embodiment includes a computer usable medium having computer usable program code embodied therewith. The computer usable program code comprises computer usable program code configured to gather monitor data information from a plurality of memory devices having finite endurance and/or retention, the

30 monitor data being data of known content stored in dedicated memory cells of known write cycle count; computer usable program code configured to analyze the monitor data

information; and computer usable program code configured to take an action relating to at least one of the devices based on the analyzing.

5 A method according to another embodiment includes gathering monitor data information from a plurality of memory devices having finite endurance and/or retention, the monitor data being data of known content stored in dedicated memory cells of known write cycle count; analyzing the monitor data information; and setting a behavioral target for at least one of the devices based on the analyzing.

10 A system according to yet another embodiment includes a plurality of memory devices having finite endurance and/or retention, each of the memory devices having: a plurality of memory blocks, at least one of the blocks having monitor data written therein, wherein the at least one block has been written to a plurality of times prior to writing the monitor data; and circuitry for addressing the blocks. The system also includes a processor in communication  
15 with the memory devices; and a clock, the processor being in communication with the clock.

Other aspects and advantages of the present invention will become apparent from the following detailed description, which, when taken in conjunction with the drawings, illustrate by way of example the principles of the invention.

#### 20 BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention will next be described by way of example with reference to the accompanying drawings in which:

25 FIG. 1 illustrates a method for determining a retention behavior for at least one block of a flash memory device, in accordance with one embodiment.

FIG. 2 illustrates a representative plot showing a shape of a flash error rate surface, in  
30 accordance with one embodiment.

FIG. 3 shows an exemplary cumulative distribution function (CDF) illustrating the probability that a given block will exceed the correction power of a sector-level error correcting code (ECC) due to wearout at some data age, in accordance with one embodiment.

FIG. 4 illustrates a representation of the overlap between the access distribution and the CDF at some data age, in accordance with one embodiment.

FIG. 5 shows representative cycle count histograms for a user and monitor data at some point in time, in accordance with one embodiment.

FIG. 6 shows one possible set of measurements from a monitor data system.

FIG. 7 illustrates a method for analyzing monitor data information from a plurality of flash memory devices, in accordance with one embodiment.

FIG. 8 illustrates a representative storage system, in accordance with one embodiment.

FIG. 9 shows a graph of results, in accordance with one example.

FIG. 10 illustrates a network architecture, in accordance with one embodiment.

FIG. 11 shows a representative hardware environment that may be associated with the servers and/or clients of FIG. 10, in accordance with one embodiment.

### DETAILED DESCRIPTION

The methodology and systems described herein may be implemented in conjunction with, or include, any type of memory device having finite endurance and/or retention. Illustrative types of memory devices having finite write/erase endurance or limited data retention include devices comprising flash memory, phase change memory, ferroelectric memory, etc.

For ease of understanding and to place the teachings presented herein in a context, much of the present description is set forth in terms of a flash device, it being understood that this is done by way of example only and is not to be construed as limiting.

5 NAND flash devices possess both finite write/erase endurance and limited data retention. Current techniques for detecting wearout and retention have escape rates that are insufficient for information technology (IT) and some consumer applications. The consequences of an escape may be severe. For example, data may be corrupted due to an escape.

10 The error rate in flash is a 3-D surface, as it depends both on the write cycle count and the age of the data. See, e.g., FIG. 2 and discussion thereof, below. Since cells can be driven to complete failure, the error rate can exceed the detection power of any sector-level error correcting code (ECC). This can result in corrupt data being passed as valid on a subsequent read operation.

15 In some cases, NAND storage devices may implement some form of wear leveling to increase the time before the first wearout event. A side effect of such an approach is to further enhance wearout clustering, as many blocks may approach wearout at the same time. Using the sector ECC to monitor the bit error rate may suffer from an aliasing effect and the  
20 ability to detect the onset of wearout may be limited.

Some other techniques for detecting wearout are similarly lacking in accuracy. These include testing for successful erase and successful programming. These techniques cannot detect the situation where the data retention time is longer than an operation verify time of  
25  $\sim 100\text{ns}$ , but significantly shorter than the stated endurance of  $3 \times 10^{17}\text{ns}$  ( $\sim 10$  years), which is a spread of 15 orders of magnitude. At least some embodiments provided herein address these and other limitations by providing techniques for determining a retention behavior of a flash memory device.

30 FIG. 1 illustrates a method 100 for determining a retention behavior for at least one block of a flash memory device, in accordance with one embodiment. As shown, monitor data is written to at least one block of a flash memory device. See operation 102.

The flash memory device may include any flash memory device including blocks of memory. For example, in one embodiment, the flash memory device may be a NAND device.

5 In the context of the present description, the data patterns used by the monitor data may refer to any data or data pattern that is known. For example, in one embodiment, the monitor data may include data patterns whose entire value is known to a memory system, a system including the memory, or another host system. In this case, errors may be characterized according to number, location, and pattern.

10 Additionally, in one embodiment, the monitor data may be generated. In this case, the monitor data may not be passed on as user data, but may be generated internally by a host system. The monitor data may be written to physical blocks, bypassing any wear leveling. In this way, the monitor data may not be subject to wear leveling.

15 Furthermore, multiple blocks of the flash memory device may be selected for storing the monitor data, where at least two of the blocks are not physically adjacent each other, e.g., to provide sampling of behaviors of different regions of the device. In any case, the monitor data may be written to multiple groups of blocks, where a number of write cycles performed for each group is different.

20 It should be noted that the at least one block may have been written to a plurality of times prior to writing the monitor data. For example, the same data may be rewritten consecutively, different data may be written on each rewrite, a combination of the two, etc. Thus, the monitor data that is ultimately stored on the device may not necessarily be the monitor data that was used during each cycle of writing.

25 After a period of time, the monitor data is read. See operation 104. The period of time may include any length of time after a point where the monitor data is written (e.g. seconds, minutes, hours, days, etc.). In one embodiment, this period of time may include a user or system defined period of time. As an option, the period of time may be configurable. In any approach, an output of a run time clock, real time clock, calendar, etc. may be used to

measure the data age, periods of time, etc. Use of a real time clock is advantageous in that the data age may be readily determined in spite of periods of being powered off, power loss, etc.

5        Once the monitor data is read, a retention behavior of the at least one block is determined based on the reading. See operation 106. In the case that the monitor data is known to the host system, the determining of the retention behavior may include comparing the read monitor data to a stored version thereof.

10       Furthermore, the reading of the monitor data and the determining of the retention behavior of the at least one block may occur at periodic intervals. In this case, the periodic intervals may include intervals defined by a user or defined automatically by a system.

15       Once the retention behavior is determined, a result of the determining is output. See operation 108. In one embodiment, an operational limit of the flash memory device, or portion thereof, may be determined based on the result of the determining of the retention behavior of the at least one block. As an option, writes may no longer be allowed to the flash memory device or a portion thereof when the operation limit is met.

20       Preferably, the monitor data is write protected by any manner known in the art, e.g., by omitting it from a mapping of available data space. It should be noted that, in one embodiment, at least a portion of the at least one block having the monitor data may be write protected. For example, at least a portion of the at least one block having the monitor data may be write protected to prevent being overwritten with user data.

25       In one approach, the write protection may be accomplished by marking the blocks as unwritable in a memory controller. In another approach, an indication of a protected block, or portion thereof, may be stored on the flash memory device. Furthermore, in one embodiment, the monitor data may include less than about 1% of a memory capacity of the  
30       flash memory device, and more preferably less than about 0.5%. In other embodiments, less than 0.1% of the memory capacity may be used for storing the monitor data.



Using the method 100, flash wearout and retention may be measured by the addition of monitor data to the flash memory device. The monitor data may include data patterns where all values are known to a host system. Thus, errors may be characterized according to number, location, and pattern.

5

A set of blocks may be positioned at high cycle counts, prior to any user data achieving these cycle counts. A sample of the error rate surface may be made using this data. Once positioned at high cycle count, the monitor data may be checked at intervals to measure the retention behavior. As an option, the measured wearout and data retention limits may be reported via self-monitoring, analysis, and reporting technology (S.M.A.R.T.).

10

As noted above, NAND flash devices may suffer from finite endurance for erase/write cycles and finite data retention. This is a consequence of the poor reversibility of the data storage process, which drives charge through an insulator. As a result, the bit error rate in flash may be characterized as a 3-D surface.

15

FIG. 2 illustrates a representative plot 200 showing a shape of a flash error rate surface, in accordance with one embodiment. As shown, the X-axis shows the data age in seconds, from 1 millisecond to 10 years. The depth axis shows the write cycle count, from 1 to  $1 \times 10^6$  cycles. The vertical axis is the error rate multiplier. The error rate surface of most flash devices has the general characteristics as shown in FIG. 2.

20

If cells are written to a sufficiently high cycle count, they fail completely (e.g. an error rate of  $\frac{1}{2}$ ). Thus, the error rate at the left rear corner of the surface illustrated in FIG. 2 increases to failure at some cycle count. Further, the failure occurs at lower cycle counts when the age of the data increases. Thus, the error rate at the right rear corner goes to failure even faster. Designing an ECC, or ECC plus a cyclic redundancy check (CRC) depends on knowing the error rate or the shape of the surface.

25

Many NAND devices have a number of issues that may impact data integrity. These may include charge detrapping, stress-induced leakage current (SILC), read disturb and write disturb, and/or other issues.

30

The manufacturing tests for flash (e.g. JEDEC, JESD47F, etc.) are inadequate for IT data integrity. The SILC test uses only 38 devices per lot for 500 hours. Zero errors in such a test indicate an escape rate as high as 6%, to 90% confidence. Further, a 500 hour test allows testing a maximum of 10% of the device to the endurance limit (assuming the rest of the device is not tested at all), but at a data age of only 100nS.

The data integrity requirements for IT systems may be very strict. Returning corrupt data is generally not acceptable. A loose target of one event per 10 M unit-years may be set. At a 10 kIOPS rate, this gives an uncorrected error rate of  $1 \times 10^{-23}$  per bit transferred for a typical 4 kB IO. Thus, designing IT flash systems using the stated endurance specifications may entail significant risk.

Wear leveling may be utilized to increase the time until there are no free blocks below the endurance cycle count. Free blocks may be achieved through spares, or even by moving the static data, for example. One goal of a static wear leveling algorithm is to maximize the cycle count until there are no free or freeable blocks below the endurance limit.

An ideal static wear level algorithm would produce the narrowest distribution at any point in time. This would give the maximum total cycle count for the device. In this case, however, wear leveling assumes the existence of a robust technique for determining that a block has worn out. That is, wear leveling assumes that a cumulative distribution function (CDF) for wearout has an onset that is sudden, and at a known cycle count. Additionally, it is assumed that all blocks have the same wearout characteristics. Further, the effects of retention are ignored.

FIG. 3 shows an exemplary CDF 300 illustrating the probability that a given block will exceed the correction power of a sector-level ECC due to wearout at some data age, in accordance with one embodiment. Sector level ECC may be utilized in flash to protect against soft errors and other effects. However, there is a finite limit to the power of the ECC, and this may be exceeded at some point due to block wearout. When this happens, corrupt data may be returned. Additional sector checks, such as a CRC, move the point at which this happens, but may not eliminate this issue.

The curve in FIG. 3 is representative of what the actual field behavior may be at a short data age, in accordance with one embodiment. Assuming a device with an endurance specification of  $10^5$  cycles, this may be assigned to a 6% probability, given the JEDEC test results. In this example, it was assumed that 94% of the blocks fail at  $10^8$  cycles. At longer data ages, the curve would shift to the left.

The reliability impact may be seen in FIG. 4. FIG. 4 illustrates a representation of the overlap between the access distribution and the CDF at some data age, showing the region of concern, in accordance with one embodiment.

Using FIG. 4, the overlap between the access distribution and the CDF may be seen. The narrower the access distribution, the more likely a large percentage of blocks are in the wearout zone before the wearout is noticed. Thus, it is more likely that a power of the sector ECC will be exceeded and corrupt data will be returned.

To address these and other issues, the error rate surface may be sampled to aid in the determination of the endurance and retention limits for a flash device. In one embodiment, this may be achieved by reserving a number of physical blocks that are not subject to wear leveling. These blocks may include data that is known to a host system such that the values may be fully verified, avoiding any aliasing problems inherent in using error rates determined from ECC.

Further, such blocks may be used at an error rate beyond the ECC capabilities without compromising data integrity. The monitor data may also be verified at intervals to measure the error rate versus retention.

In one embodiment, a set of full erase blocks that are positioned near the expected rise of the CDF may be utilized to aid in determining wearout before it leads to data corruption. For example, a flash device may have 16,000 erase blocks. Using 15 blocks as monitor data would take less than 0.1% of the device capacity.

As an option, 15 blocks arranged in 5 groups of 3 may be utilized. Each group may have the same cycle count. In this case, having multiple blocks in a group may improve the quality of the measurement. The locations (in terms of cycle count) for each group may be chosen to be near the region where the CDF is expected to begin to have a magnitude that may impact data integrity.

As an example, Table 1 shows characteristics representative of flash device, in accordance with one embodiment.

Table 1

Capacity	8GB	Page Write	800 $\mu$ s
Page	4kB	Block Erase	1500 $\mu$ s
Erase block	128 pages	Page Read	60 $\mu$ s
Blocks	16,384	Block Read	21ms
Endurance	10 <sup>5</sup>	Block Write	116ms

In this example, the time to erase/write/read a block is 138ms. Thus, a monitor data target layout of 5 monitor groups may be chosen. Table 2 shows a representation of a monitor data target layout of 5 monitor groups, in accordance with one embodiment.

Table 2

Target Cycles (K)	20	36	63	112	200	Total
Time to Deploy (H)	2.3	4.1	7.3	12.9	23.0	50

In this case, the cycle counts are chosen to be power-law distributed assuming that wearout has a log-normal behavior. In various embodiments, other distributions may be chosen to maximize the effectiveness of the monitor data. In this example, it would take a total of 50 hours to cycle all the monitor data to the final counts.

FIG. 5 shows representative cycle count histograms 500 for a user and monitor data at some point in time, in accordance with one embodiment. In this case, the histograms of user data and monitor data are shown as function of write cycle count. Furthermore, the data includes wear leveling.

As shown, the monitor data is positioned near the anticipated wearout cycle count. In some cases, the time required to cycle all the monitor data blocks may be longer than desirable to perform prior to shipment. In such a case, the monitor data may be partially cycled prior to shipment, with the remainder of the cycling occurring in the field.

5

Once the monitor data is at the final cycle count, the blocks may be read at intervals and the read data may be compared with the known data. This allows complete identification of errors, including the ability to maintain a rate of change.

10

FIG. 6 shows one possible set of measurements 600 from a monitor data system. Each monitor group is shown as a column group and each in the group represents a different data age. The height of each column is the number of errors per sector and the lower region labeled "ECC Limit" represents a possible ECC correction limit. As shown, some measurements may exceed the ECC capability and thus pose a risk for data integrity.

15

In one embodiment, operational limits may be derived from this data. For example, Table 3 shows operational limits that may be generated for the exemplary data in FIG. 6, in accordance with one embodiment.

20

Table 3

Minimum Cycle Count	Maximum Cycle Count	Maximum Data Age
<40k		4
40k	60k	3
60k	80k	2
80k	100k	2
100k	120k	1
>120k		Not allowed

A monitoring data pattern may also be chosen. In various embodiments, this may be a static pattern or be computed for a particular block. As an option, a pattern that stresses the endurance and/or retention may be utilized. In either case, a system may be configured such that the expected contents of each monitor block may be known without reading anything

25

from the flash device. When one of the limits is reached, the system may be informed by a S.M.A.R.T. flag, or another technique such as no longer accepting writes (e.g. returning error on write, etc.).

5 While much of the above description has described the monitor data units as including entire flash blocks, the techniques discussed are not limited to using entire flash blocks. For example, in one embodiment, the monitor data units may be as small as a sector size.

10 Although, using sector size monitor data units will retain the monitoring characteristics, in some cases, this may leave exposure to hard errors. For example, if a page size monitor is utilized with user data in the rest of the block, the user data may be marked as lost (i.e. a hard error) if the monitor page exhibited an error rate close to the ECC limits. This is because the sector ECC may not be reliable in such situations. Using this technique, however, the probability of corruption events would still be reduced at a lower overhead.

15 FIG. 7 illustrates a method 700 for analyzing monitor data information from a plurality of flash memory devices, in accordance with one embodiment. As shown, monitor data information is gathered from a plurality of flash memory devices, the monitor data being data of known content stored in dedicated memory cells of known write cycle count. See  
20 operation 702.

The monitor data may include any data or data pattern that is known. In one embodiment, the monitor data information may include error rate information relating to at least one of the data age of the monitor data and the write cycle count. Furthermore, monitor data and/or the  
25 monitor data information may be gathered using a variety of techniques. In one embodiment, the monitor data may be generated.

Once the monitor data information is gathered, the monitor data information is analyzed. See operation 704. Additionally, an action is taken relating to at least one of the devices  
30 based on the analyzing. See operation 706.

The action may include a variety of actions. In one embodiment, taking the action may include setting a behavioral target for at least one of the devices. In this case, the behavioral target may be a limit on data write cycle count. As another option, the behavioral target may be a limit on data age.

5

In another embodiment, taking the action may include setting a behavioral target for at least one of the devices. In this case, the behavioral target may be a limit on data write cycle count. As another option, the behavioral target may be a limit on data age.

10

It should be noted that the gathering of the monitor data information and the analyzing of the information may occur any number of times. For example, the gathering and analyzing steps may be performed at least twice, and a statistical trend may be generated based thereon. Additionally, the gathering and analyzing steps may be performed at least twice, and a time-dependence of the monitor data may be tracked.

15

By implementing this technique, a system and method may be provided for increasing the reliability of flash-based storage systems by gathering monitor data information from a set of devices, analyzing the data, and setting behavioral targets for the devices. The monitor data may include error rate information in terms of cycle count and data age for each device.

20

Additionally, the behavioral targets may include setting device level limits on cycle count and data age. Aggregating the data from multiple devices allows the sample size to be increased and may help identify statistical trends. Further, the aggregation can also track the time-dependence of the data.

25

The behavioral target may be updated after a period of time. In one approach, monitor data may be gathered again and analyzed to generate an updated behavioral target after a period of time has elapsed from when a prior behavioral target was set. This will thus take into account the effect of the longer age of the monitor data, potentially improving the behavioral target. In another approach, the behavioral target may be updated based at least in part on other factors such as failure events of memory devices, additional information or parameters that became available since the target was set, etc.

30

In one embodiment, the monitor data approach may be utilized to provide a measurement of the error rate surface in a set of flash devices in a storage system. Furthermore, a measure of the error rate vs. cycle count and error rate vs. data age may be provided for a given flash device.

5

FIG. 8 illustrates a representative storage system 800, in accordance with one embodiment. As shown, the storage system 800 may include a storage controller 802 and a plurality of flash devices 804. The behavior of the system 800 may be adjusted in response to monitor data information. The controller 802 may be in communication with a remote data repository 806 that aggregates field data from such storage systems.

10

The controller 802 may determine likely initial cycle count targets for the set of local flash devices 804 based on this data, or as a factory preset. These values may be passed to the flash devices 804, which will cycle their monitor data to place it at the predetermined cycle counts. The storage controller 802 may then collect the error characteristics of the monitor data at different times.

15

In one embodiment, the storage controller 802 may include, or have access to, a real-time clock. This clock may be used to record the write times of the monitor data (either in the monitor data or in another reserved area) and to determine the age of the monitor data.

20

The error characteristics reported by the monitor data may be collected over an interface to the flash device, such as via S.M.A.R.T. The storage system 800 may aggregate the collected data from the flash devices 804 to improve the statistics, and identify outliers.

25

For example, the storage system 800 may comprise 16 flash units (e.g. SSDs), each of which may comprise 16 flash devices. This provides 256 flash devices, thus improving the quality of the samples by 16x (assuming Poisson statistics). If each flash device had 3 blocks in each data set, the system data set size would be 768 block samples for each set. Thus, it is possible to obtain good statistics and to perform outlier detection.

30



The monitor data may report the error configuration for each monitor data set. Each set may have a different cycle count. Thus, the storage controller 802 may use them to estimate the error rate curve for the population. The error rate curve may be used to select the endurance and retention limit criteria.

5

As an example of how the data may be used in the system 800, a Monte Carlo simulation may be performed to generate a set of error count measurements for the monitor data on each device. In this case, it is assumed that the monitor data reports at least the mean error count for each monitor data block. The error count values for each cycle count are merged

10

Assuming a Gaussian spread, an escape limit target may be computed, beyond which an outer ECC will be able to correct. It is also assumed that the target probability of escape is  $1 \times 10^{-6}$ , which should be sufficient to prevent device-to-device correlation from overwhelming an outer ECC. This works out to be 4.8 standard deviations for this example (assuming Gaussian statistics). Table 4 shows the results of the Monte Carlo simulation for this example.

15

Table 4

Cycle	20	36	63	112	200
Count (k)					
Mean	1.4	2.4	4.2	7.4	13.2
errors/sector					
Sigma	0.2	0.3	0.4	0.6	1.1
errors/sector					
Sigma Clip	2.2	3.7	6.1	10.4	18.4
Curve					
Number	1	1	1	1	1
outliers					

20

In this case, the monitor data is sampled at 5 separate cycle counts. The “Mean” row lists the mean number of errors per sector, and the “Sigma” row lists the corresponding standard deviation. The “Sigma Clip Curve” row is the 4.8 standard deviation value at each cycle

count. The last row is the number of devices which have error counts exceeding 4.8 standard deviations.

In this example, there is only a single device that exceeds the clip, thus it is assumed that this is an outlier. Further, it is assumed that an error limit of 8 errors per sector is chosen. This would be determined by the capabilities of the sector-level error correction.

Given such a value, the cycle count limit may be estimated for the non-outliers by interpolating the sigma clip curve. For this example, the limit value is 85k write cycles. Similarly, the outlier limit may be computed to be 61k cycles in this example. A more conservative approach may also be utilized with the outliers. For example, a lower cycle count may be chosen rather than simple interpolation. Similarly, other techniques of estimation, such as curve fitting may be used.

FIG. 9 shows a graph 900 of the results, in accordance with this example. The graph 900 illustrates that the outlier is well beyond the edge of the distribution, and therefore applying the normal clip limit would likely result in an error exposure. Once the cycle count limits have been computed, they may be fed back to the flash devices or a controller responsible for the endurance limits. This may be repeated at different times, thus reflecting different data ages. As such, a set of endurance and retention targets may both be generated for a storage system.

This approach is an improvement over allowing each device to set its own limits as it benefits from population statistics, and can more accurately distinguish outliers from sampling noise. Each monitor data set may include a plurality of erase blocks, each of which comprises many data sectors. For example, a typical NAND device may have 528kB erase blocks, including 1024 sectors, each with 528 bytes. Assuming that each monitor set used 3 data blocks, for the entire set of 256 devices there are  $3.3 \times 10^9$  bits per set. Thus, reliability is enhanced and costs are decreased since overkill is avoided.

It should be noted that the invention can take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software

elements. For example, in one embodiment, a system may include a flash memory device having a plurality of memory blocks, where at least one of the blocks has monitor data written therein, and the at least one block has been written to a plurality of times prior to writing the monitor data. The system may also include circuitry for addressing the blocks and a secondary memory storing a copy of the monitor data. Additionally, the system may include a storage controller and a plurality of flash memory devices.

As will be appreciated by one skilled in the art, the present invention may be embodied as a system, method or computer program product. Accordingly, the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, the present invention may take the form of a computer program product embodied in any tangible medium of expression having computer-usable program code embodied in the medium.

Any combination of one or more computer usable or computer readable medium(s) may be utilized. The computer-usable or computer-readable medium may be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or propagation medium. More specific examples (a non-exhaustive list) of the computer-readable medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a transmission media such as those supporting the Internet or an intranet, or a magnetic storage device. Note that the computer-usable or computer-readable medium could even be paper or another suitable medium upon which the program is printed, as the program can be electronically captured, via, for instance, optical scanning of the paper or other medium, then compiled, interpreted, or otherwise processed in a suitable manner, if necessary, and then stored in a computer memory. In the context of this document, a computer-usable or computer-readable medium may be any medium that can contain, store, communicate, propagate, or transport the

program for use by or in connection with the instruction execution system, apparatus, or device. The computer-usable medium may include a propagated data signal with the computer-usable program code embodied therewith, either in baseband or as part of a carrier wave. The computer usable program code may be transmitted using any appropriate  
5 medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc.

Computer program code for carrying out operations of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural  
10 programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network,  
15 including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

The present invention is described herein with reference to flowchart illustrations and/or  
20 block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose  
25 computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

30 These computer program instructions may also be stored in a computer-readable medium that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable medium

produce an article of manufacture including instruction means which implement the function/act specified in the flowchart and/or block diagram block or blocks.

5 The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

10 The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which  
15 comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It  
20 will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

25 A data processing system suitable for storing and/or executing program code will include at least one processor coupled directly or indirectly to memory elements through a system bus. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk  
30 storage during execution.

Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, etc.) can be coupled to the system either directly or through intervening I/O controllers.

5 Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modem and Ethernet cards are just a few of the currently available types of network adapters.

10 FIG. 10 illustrates a network architecture 1000, in accordance with one embodiment. As an option, the functionality and architecture of FIGS. 1-9 may be implemented in the context of the network architecture 1000. However, the functionality and architecture of FIGS. 1-9 may be carried out in any desired environment. It should be noted that the aforementioned definitions may apply during the present description.

15 As shown in FIG. 10, a plurality of remote networks 1002 are provided including a first remote network 1004 and a second remote network 1006. A gateway 1007 may be coupled between the remote networks 1002 and a proximate network 1008. In the context of the present network architecture 1000, the networks 1004, 1006 may each take any form including, but not limited to a LAN, a WAN such as the Internet, PSTN, internal telephone network, etc.

20 In use, the gateway 1007 serves as an entrance point from the remote networks 1002 to the proximate network 1008. As such, the gateway 1007 may function as a router, which is capable of directing a given packet of data that arrives at the gateway 1007, and a switch, which furnishes the actual path in and out of the gateway 1007 for a given packet.

25 Further included is at least one data server 1014 coupled to the proximate network 1008, and which is accessible from the remote networks 1002 via the gateway 1007. It should be noted that the data server(s) 1014 may include any type of computing device/groupware. Coupled to each data server 1014 is a plurality of user devices 1016. Such user devices 1016 may include a desktop computer, lap-top computer, hand-held computer, printer or any other type

of logic. It should be noted that a user device 1017 may also be directly coupled to any of the networks, in one embodiment.

5 A facsimile machine 1020 or series of facsimile machines 1020 may be coupled to one or more of the networks 1004, 1006, 1008. It should be noted that databases and/or additional components may be utilized with, or integrated into, any type of network element coupled to the networks 1004, 1006, 1008. In the context of the present description, a network element may refer to any component of a network.

10 FIG. 11 shows a representative hardware environment associated with a user device 1016 of FIG. 10, in accordance with one embodiment. Such figure illustrates a typical hardware configuration of a workstation having a central processing unit 1110, such as a microprocessor, and a number of other units interconnected via a system bus 1112.

15 The workstation shown in FIG. 11 includes a Random Access Memory (RAM) 1114, Read Only Memory (ROM) 1116, an I/O adapter 1118 for connecting peripheral devices such as disk storage units 1120 to the bus 1112, a user interface adapter 1122 for connecting a keyboard 1124, a mouse 1126, a speaker 1128, a microphone 1132, and/or other user interface devices such as a touch screen and a digital camera (not shown) to the bus 1112,  
20 communication adapter 1134 for connecting the workstation to a communication network 1135 (e.g., a data processing network) and a display adapter 1136 for connecting the bus 1112 to a display device 1138.

The workstation may have resident thereon an operating system such as the Microsoft  
25 Windows® Operating System (OS), a MAC OS, or UNIX operating system. It will be appreciated that a preferred embodiment may also be implemented on platforms and operating systems other than those mentioned. A preferred embodiment may be written using JAVA, XML, C, and/or C++ language, or other programming languages, along with an object oriented programming methodology. Object oriented programming (OOP), which  
30 has become increasingly used to develop complex applications, may be used.

While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.



**CLAIMS**

1. A method, comprising:

gathering monitor data information from a plurality of memory devices having finite  
endurance and/or retention, the monitor data being data of known content stored in dedicated  
5 memory cells of known write cycle count;

analyzing the monitor data information; and

taking an action relating to at least one of the devices based on the analyzing.

2. The method of claim 1, further comprising generating the monitor data.

3. The method of claim 1, wherein the monitor data information includes error rate  
information relating to at least one of the data age of the monitor data and the write cycle  
count.

4. The method of claim 1, further comprising performing the gathering and analyzing  
steps at least twice, and generating a statistical trend based thereon.

5. The method of claim 1, further comprising performing the gathering and analyzing  
steps at least twice, and tracking a time-dependence of the monitor data.

6. The method of claim 1, wherein taking the action includes setting a behavioral target  
for at least one of the devices.

7. The method of claim 6, wherein the behavioral target is a limit on data write cycle  
count.

8. The method of claim 6, wherein the behavioral target is a limit on data age.

9. The method of claim 1, wherein the monitor data is not subject to wear leveling.

10. The method of claim 6, further comprising updating the behavioral target for the at least one of the devices after a period of time has elapsed,.

11. The method of claim 1, wherein the memory devices are NAND devices.

12. The method of claim 1, wherein the monitor data comprises less than about 1% of a memory capacity of each of the memory devices.

13. The method of claim 1, wherein the monitor data in each of the memory devices is written to multiple groups of blocks, wherein a number of write cycles performed for each group is different.

14. A computer program product, comprising a computer usable medium having computer usable program code embodied therewith, the computer usable program code when executed by a computer causing the computer to perform the steps of claims 1 to 13.

15. A system, comprising:

a plurality of memory devices having finite endurance and/or retention, each of the memory devices having:

a plurality of memory blocks, at least one of the blocks having monitor data written therein, wherein the at least one block has been written to a plurality of times prior to writing the monitor data; and

circuitry for addressing the blocks;

a processor in communication with the memory devices; and

a clock, the processor being in communication with the clock.

16. The system of claim 15, further comprising a secondary memory storing a copy of the monitor data.

17. The system of claim 15, further comprising a computer usable medium as claimed in claim 14.

18. The system of any of claims 15 to 17, wherein the memory devices are NAND devices.

1/11

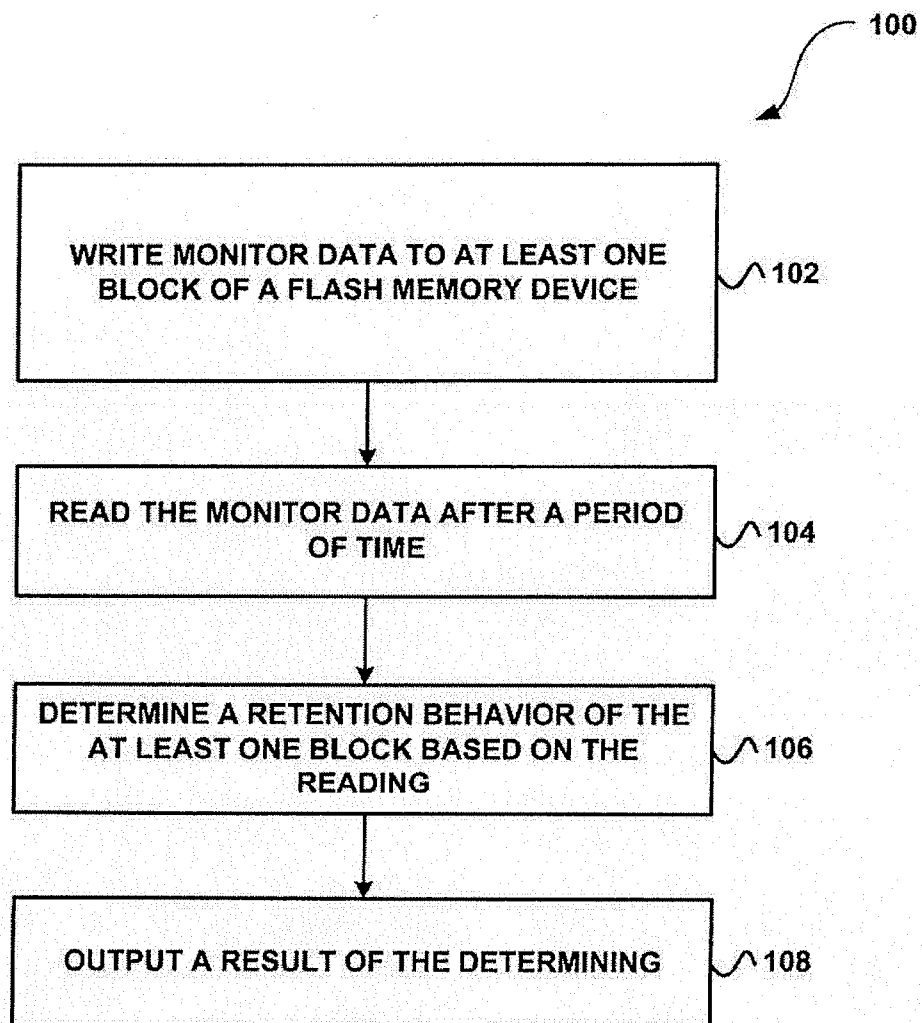


FIG. 1

2/11

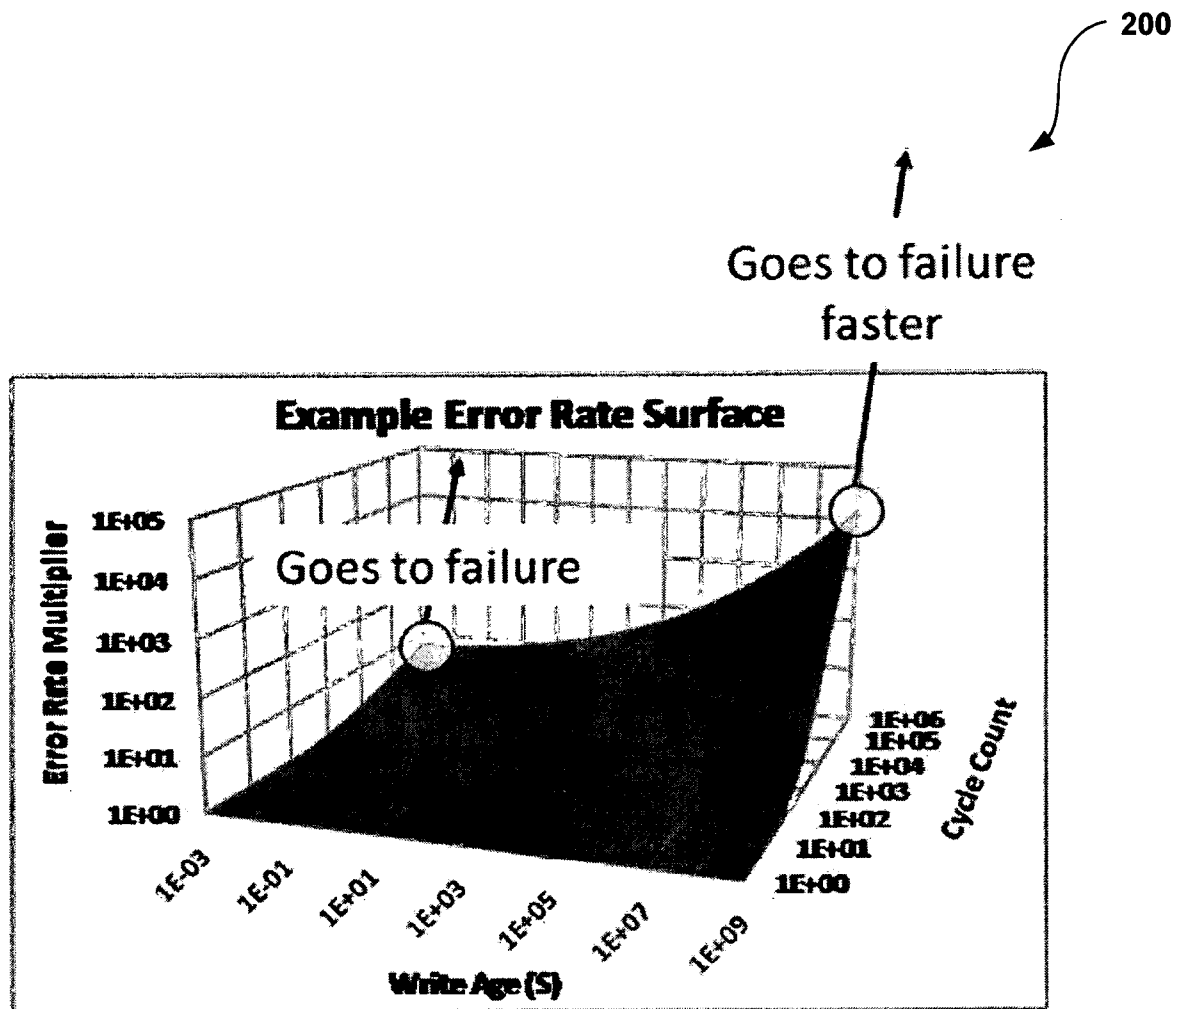
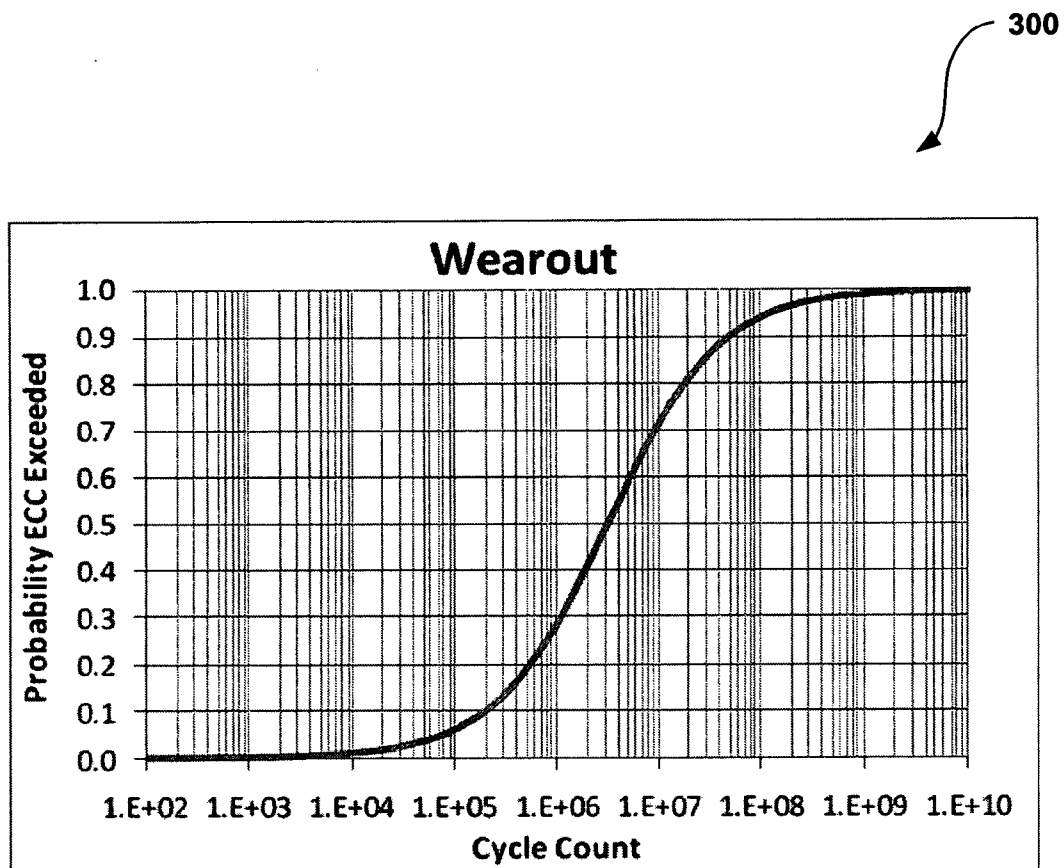
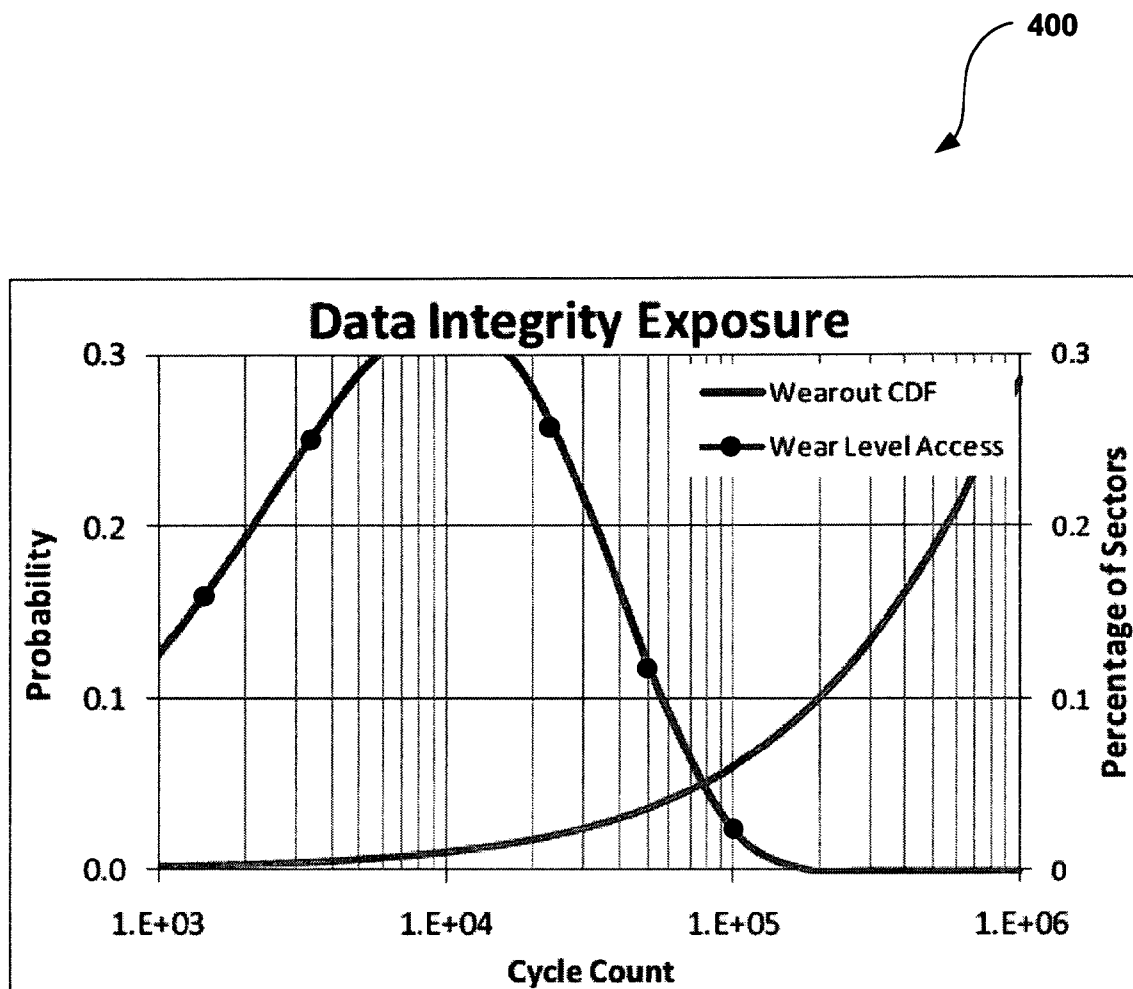


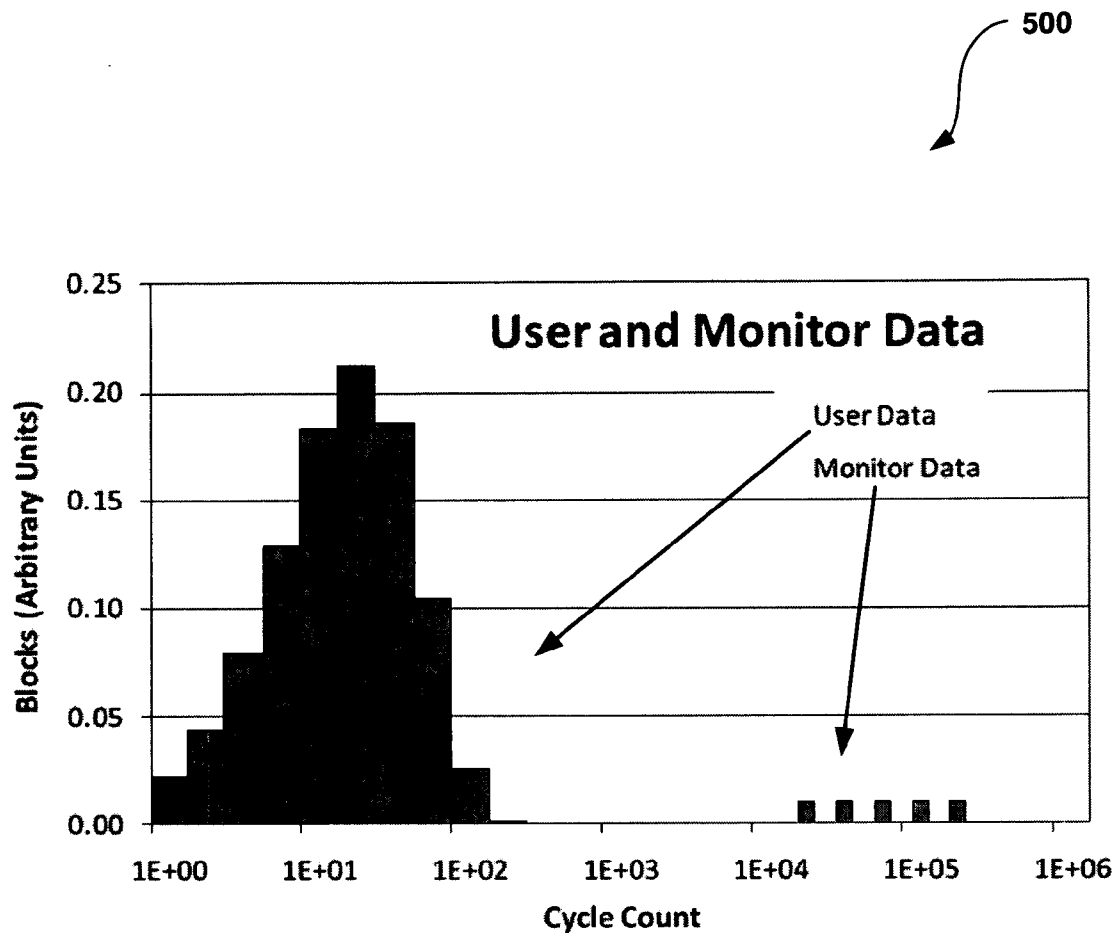
FIG. 2

3/11

**FIG. 3**

**FIG. 4**

5/11

**FIG. 5**



6/11

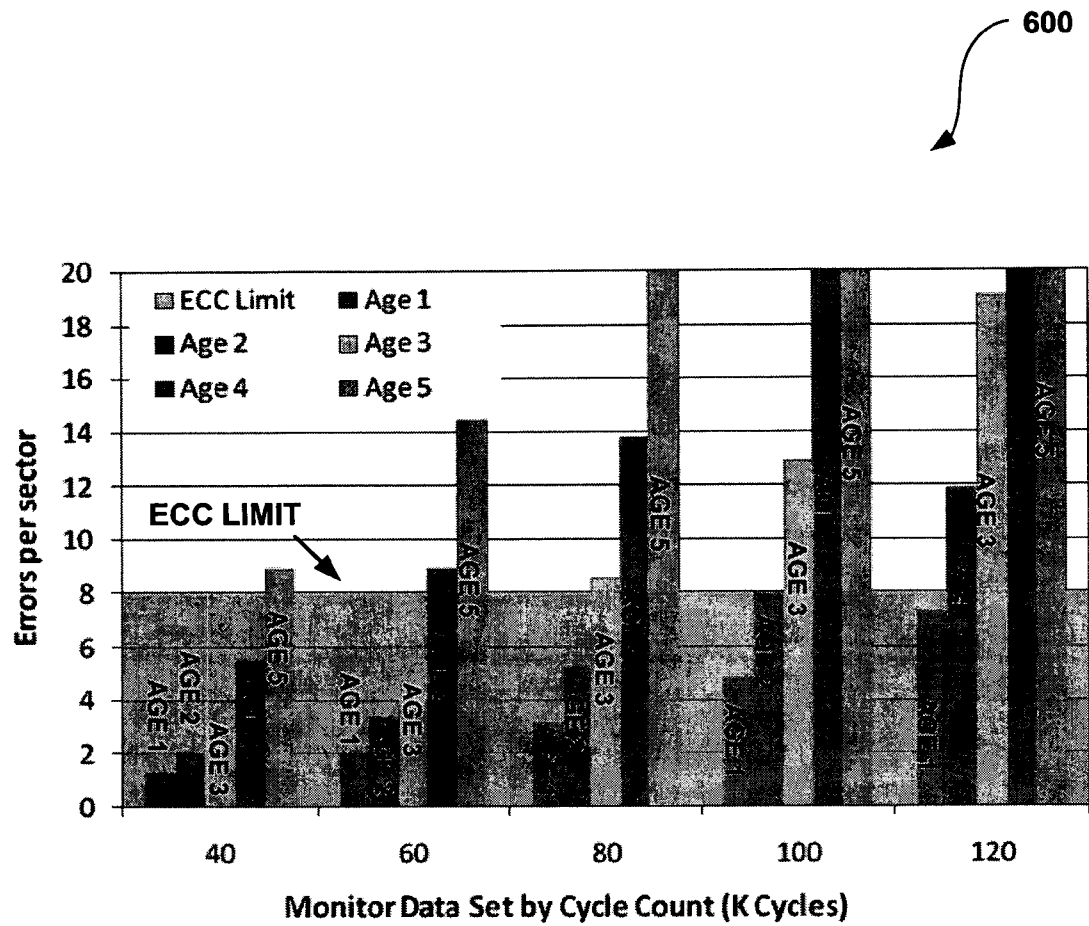


FIG. 6

7/11

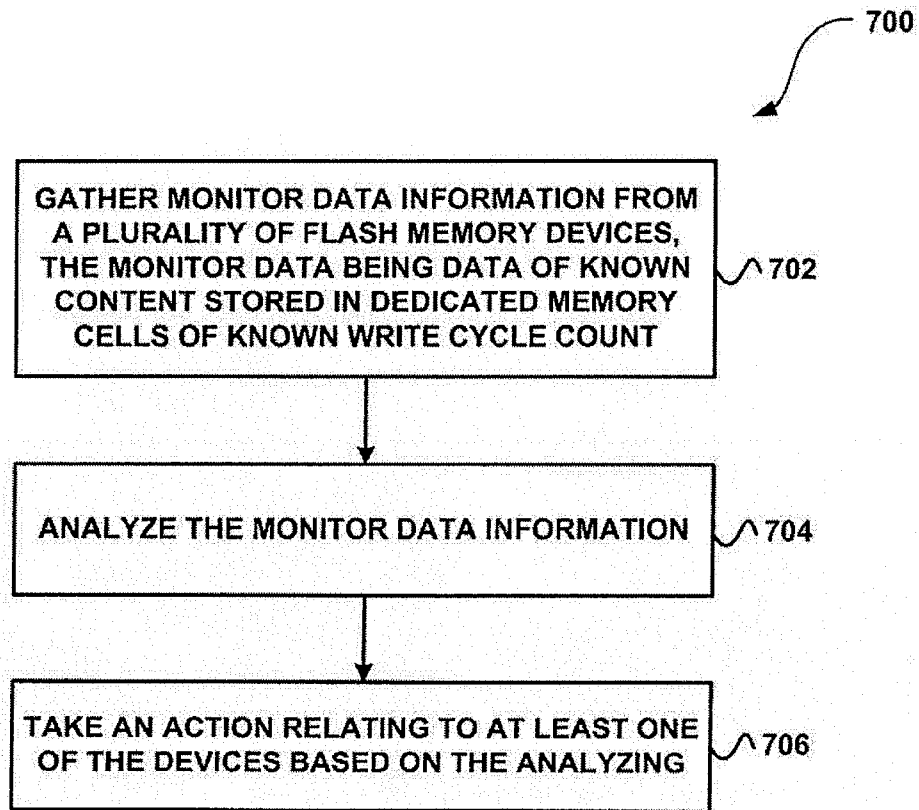


FIG. 7

8/11

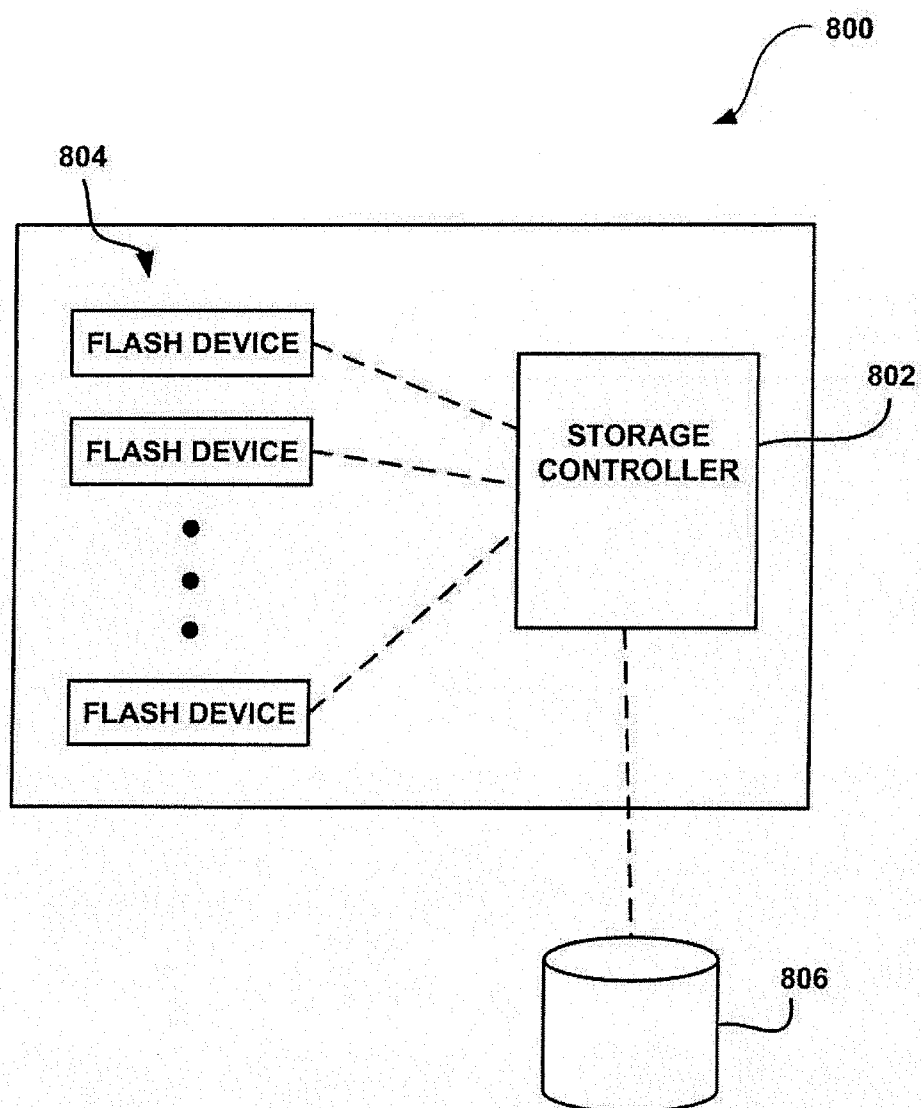


FIG. 8

9/11

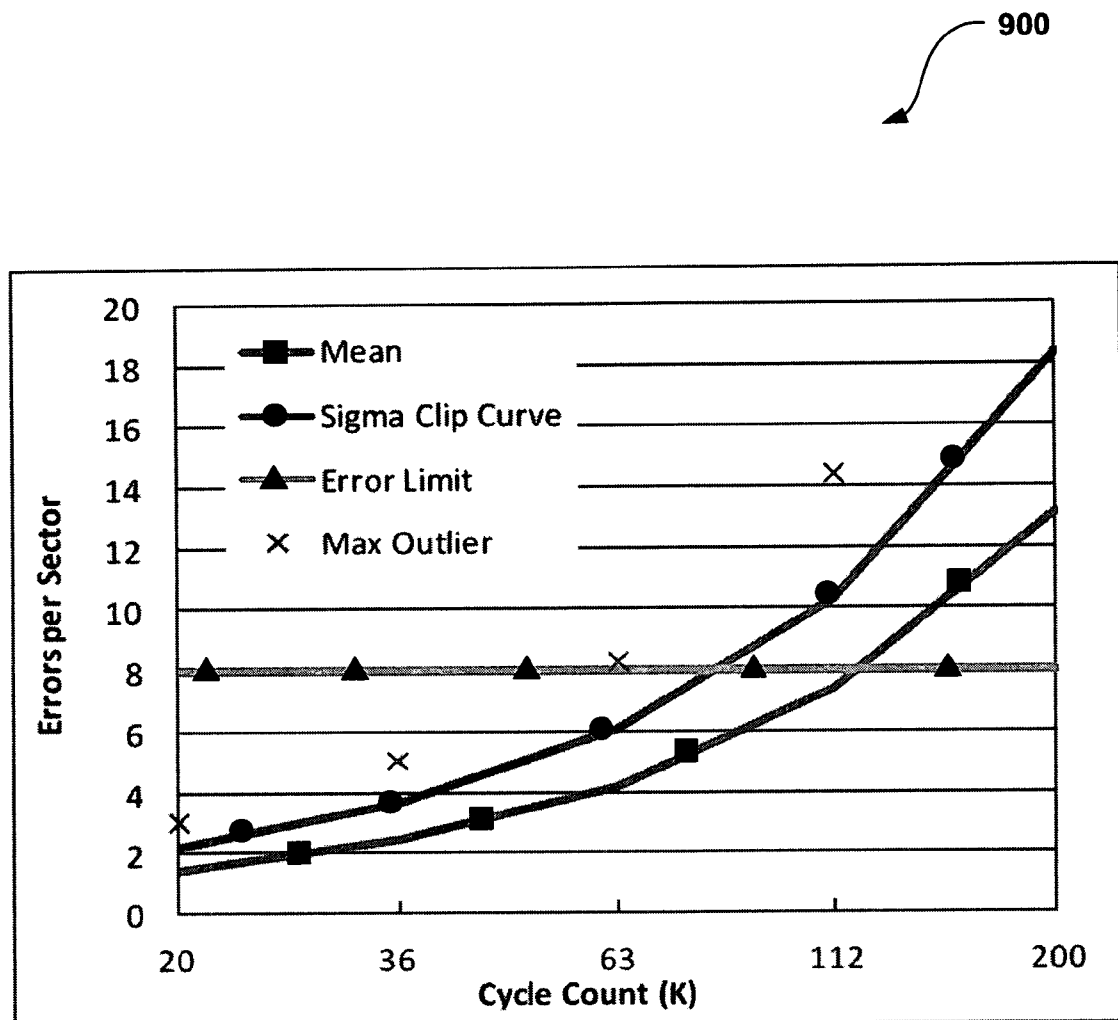
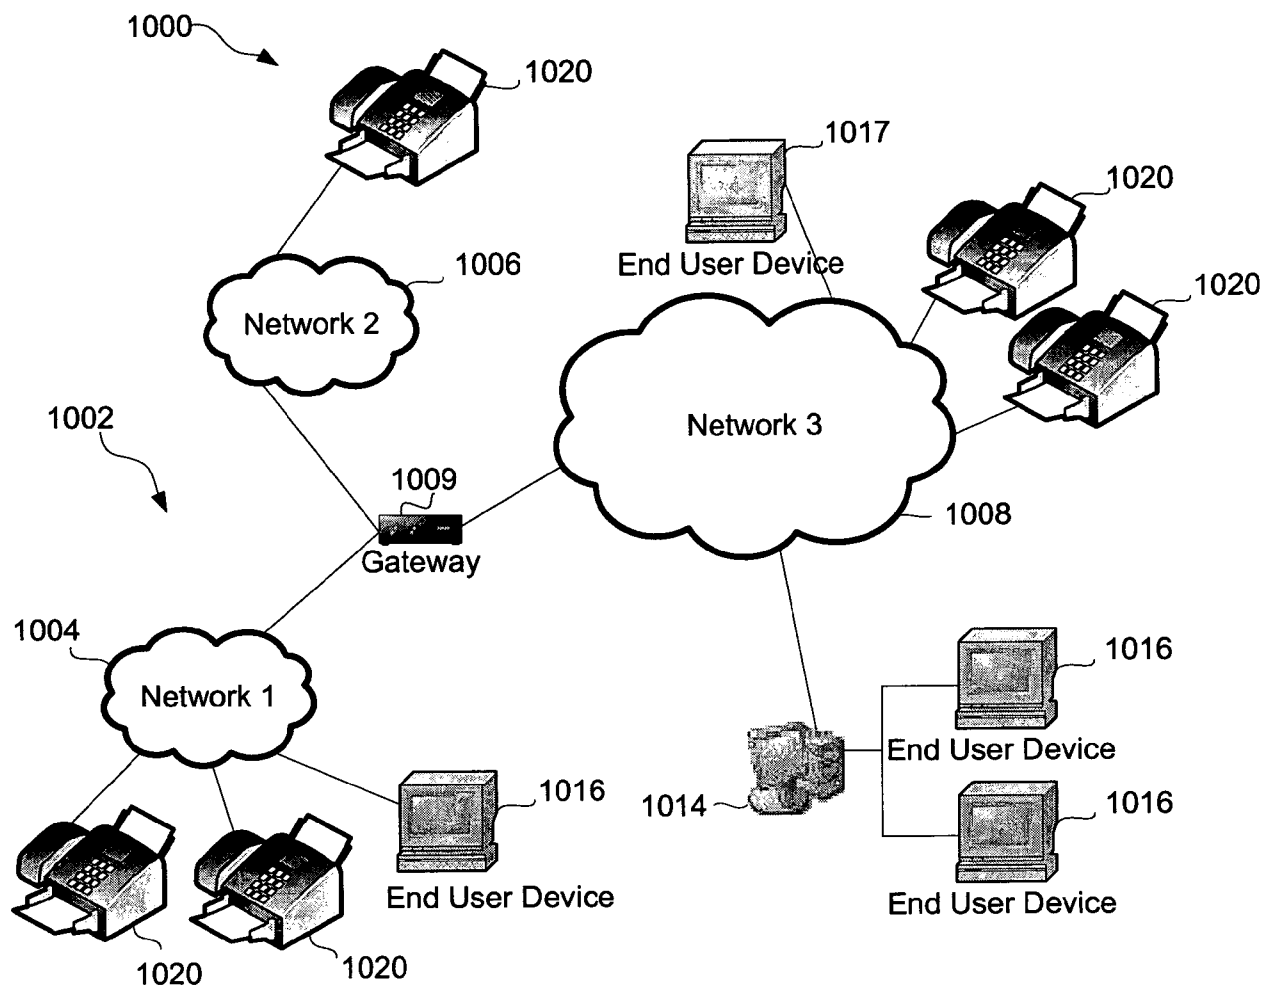
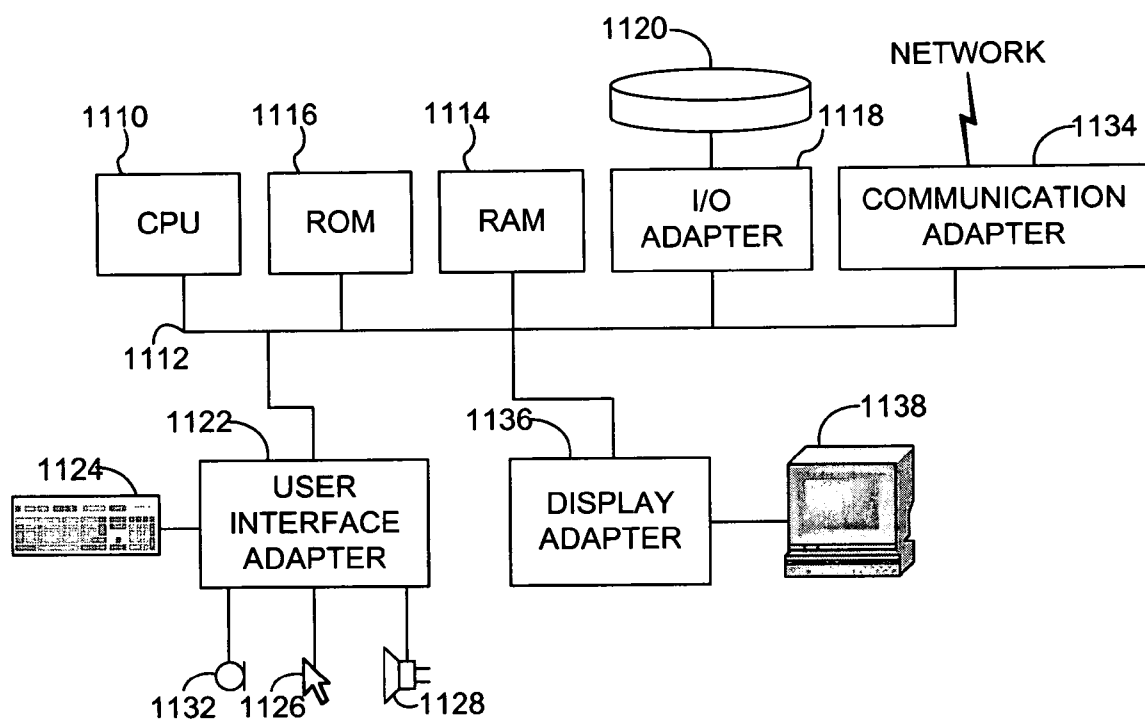


FIG. 9

10/11

**FIG. 10**

11/11

**FIG. 11**