

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
22 November 2007 (22.11.2007)

PCT

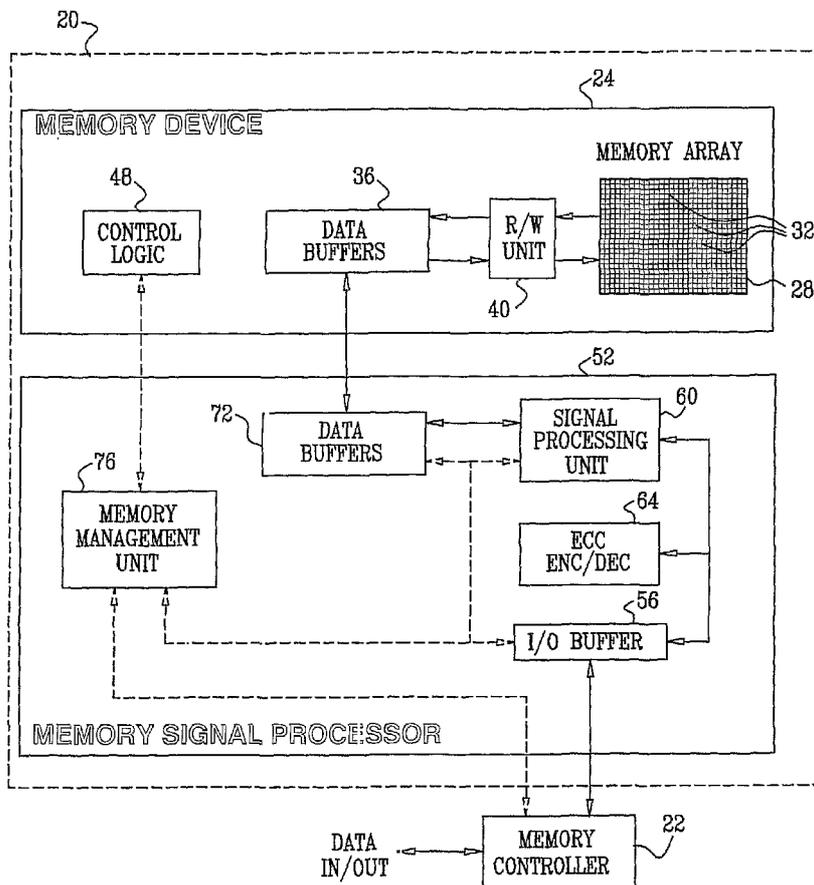
(10) International Publication Number  
**WO 2007/132456 A2**

- (51) **International Patent Classification:**  
*H04N 7114* (2006.01)
- (21) **International Application Number:**  
**PCT/IL2007/000579**
- (22) **International Filing Date:** 10 May 2007 (10.05.2007)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**

60/747,106	12 May 2006 (12.05.2006)	US
60/822,236	13 August 2006 (13.08.2006)	US
60/825,913	17 September 2006 (17.09.2006)	US
60/866,071	16 November 2006 (16.11.2006)	US
60/866,860	22 November 2006 (22.11.2006)	US
60/867,399	28 November 2006 (28.11.2006)	US
60/871,838	26 December 2006 (26.12.2006)	US
60/882,240	28 December 2006 (28.12.2006)	US
60/883,071	2 January 2007 (02.01.2007)	US
60/885,987	22 January 2007 (22.01.2007)	US
60/889,277	11 February 2007 (11.02.2007)	US
- (71) **Applicant** (for all designated States except US): **ANOBIT TECHNOLOGIES LTD.** [IL/IL]; 6 Galgalei Haplada Street, 46722 Herzlia (IL).
- (72) **Inventors; and**
- (75) **Inventors/Applicants** (for US only): **SHALVI, Ofir** [IL/IL]; 15 Hanna Senesh Street, 43362 Ra'anana (IL). **SOKOLOV, Dotan** [IL/IL]; 3 Hapamonim Street, 43391 Ra'anana (IL). **MAISLOS, Ariel** [IL/US]; 810 Durshire Way, Sunnyvale, California 94087 (US). **COHEN, Zeev** [IL/IL]; 17B Bareket Street, 30900 Zichron-yaakov (IL). **GURGI, Eyal** [IL/IL]; 10 Hameginim Street, 49442 Petah-tikva (IL). **SEMO, Gil** [IL/IL]; 16 Reuven Rubin Street, 69412 Tel-aviv (IL).
- (74) **Agents:** **SANFORD T. COLB & CO.** et al; P.O. Box 2273, 76122 Rehovot (IL).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN,

[Continued on next page]

(54) **Title:** MEMORY DEVICE WITH ADAPTIVE CAPACITY



(57) **Abstract:** A method for data storage in a memory (28) that includes a plurality of analog memory cells (32) includes estimating respective achievable storage capacities of the analog memory cells. The memory cells are assigned respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities. The data is stored in the memory cells in accordance with the respective assigned storage configurations. The achievable storage capacities of the analog memory cells are re-estimated after the memory has been installed in a host system and used for storing the data in the host system. The storage configurations are modified responsively to the re-estimated achievable capacities.

WO 2007/132456 A2



IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, 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, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, 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:**

— *without international search report and to be republished upon receipt of that report*

**(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,

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

MEMORY DEVICE WITH ADAPTIVE CAPACITY**CROSS-REFERENCE TO RELATED APPLICATIONS**

This application claims the benefit of U.S. Provisional Patent Application 60/747,106, filed May 12, 2006, U.S. Provisional Patent Application 60/822,236, filed August 13, 2006, U.S. Provisional Patent Application 60/825,913, filed September 17, 2006, U.S. Provisional Patent Application 60/866,071, filed November 16, 2006, U.S. Provisional Patent Application 60/866,860, filed November 22, 2006, U.S. Provisional Patent Application 60/867,399, filed November 28, 2006, U.S. Provisional Patent Application 60/871,838, filed December 26, 2006, U.S. Provisional Patent Application 60/882,240, filed December 28, 2006, U.S. Provisional Patent Application 60/883,071, filed January 2, 2007, U.S. Provisional Patent Application 60/885,987, filed January 22, 2007, and U.S. Provisional Patent Application 60/889,277, filed February 11, 2007, whose disclosures are incorporated herein by reference.

**FIELD OF THE INVENTION**

The present invention relates generally to memory devices, and particularly to memory devices having adaptive storage capacity.

**BACKGROUND OF THE INVENTION**

Several types of memory devices, such as Flash memories and Dynamic Random Access Memory (DRAM), use arrays of analog memory cells for storing data. Flash memory devices are described, for example, by Bez et al., in "Introduction to Flash Memory," Proceedings of the IEEE, volume 91, number 4, April, 2003, pages 489-502, which is incorporated herein by reference.

In such memory devices, each analog memory cell typically comprises a transistor, which holds a certain amount of electric charge that represents the information stored in the cell. The electric charge written into a particular cell influences the "threshold voltage" of the cell, i.e., the voltage that needs to be applied to the cell so that the cell will conduct current.

Some memory devices, commonly referred to as Single-Level Cell (SLC) devices, store a single bit of information in each memory cell. Typically, the range of possible threshold voltages of the cell is divided into two regions. A voltage value falling in one of the regions represents a "0" bit value, and a voltage belonging to the second region represents "1". Higher-density devices, often referred to as Multi-Level Cell (MLC) devices, store two or more bits per memory cell. In multi-level cells, the range of threshold voltages is divided into more than two regions, with each region representing more than one bit.

Multi-level Flash cells and devices are described, for example, by Eitan et al., in "Multilevel Flash Cells and their Trade-Offs," Proceedings of the 1996 IEEE International Electron Devices Meeting (IEDM), New York, New York, pages 169-172, which is incorporated herein by reference. The paper compares several kinds of multilevel Flash cells, such as common ground, DINOR, AND, NOR and NAND cells.

Eitan et al., describe another type of analog memory cell called Nitride Read Only Memory (NROM) in "Can NROM, a 2-bit, Trapping Storage NVM Cell, Give a Real Challenge to Floating Gate Cells?" Proceedings of the 1999 International Conference on Solid State Devices and Materials (SSDM), Tokyo, Japan, September 21-24, 1999, pages 522-524, which is incorporated herein by reference. NROM cells are also described by Maayan et al., in "A 512 Mb NROM Flash Data Storage Memory with 8 MB/s Data Rate", Proceedings of the 2002 IEEE International Solid-State Circuits Conference (ISSCC 2002), San Francisco, California, February 3-7, 2002, pages 100-101, which is incorporated herein by reference.

Other exemplary types of analog memory cells are Ferroelectric RAM (FRAM) cells, magnetic RAM (MRAM) cells and phase change RAM (PRAM, also referred to as Phase Change Memory - PCM) cells. FRAM, MRAM and PRAM cells are described, for example, by Kim and Koh in "Future Memory Technology including Emerging New Memories," Proceedings of the 24<sup>th</sup> International Conference on Microelectronics (MIEL), Nis, Serbia and Montenegro, May 16-19, 2004, volume 1, pages 377-384, which is incorporated herein by reference.

In some applications, data is stored in memory cells at varying densities. For example, U.S. Patent 6,363,008, whose disclosure is incorporated herein by reference, describes a multiple-bit-per-cell memory that includes multiple memory arrays, wherein the number of bits stored per cell is separately set for each of the memory arrays. Memory arrays that testing proves are accurate when writing, storing, and reading a larger number of bits per cell are set to store more bits per cell, and memory arrays that cannot accurately write, store, or read as many bits per cell are set to store fewer bits per cell.

In U.S. Patent 6,456,528, whose disclosure is incorporated herein by reference, a flash non-volatile memory system that normally operates its memory cells in multiple storage states is provided with the ability to operate some selected or all of its memory cell blocks in two states instead. The two states are selected to be the furthest separated of the multiple states, thereby providing an increased margin during two-state operation.

U.S. Patent 5,930,167, whose disclosure is incorporated herein by reference, describes a memory system including an array of flash Electrically Erasable and Programmable Read

Only Memory (EEPROM) cells arranged in blocks of cells that are erasable together. The individual cells store more than one bit of data as a result of operating the individual cells with more than two detectable threshold ranges or states. Any portion of the array in which data is not stored can be used as a write cache, where individual ones of the cells store a single  
5 bit of data by operating with only two detectable threshold ranges. Data coming into the memory is initially written in available blocks in two states. At a later time, in the background, the cached data is read, compressed and written back into fewer blocks of the memory in multi-state.

U.S. Patent 6,466,476, whose disclosure is incorporated herein by reference, describes  
10 multi-bit-per-cell non-volatile memory, which stores different portions of a data stream using different numbers of bits per cell. In particular, data that requires a high degree of data integrity (e.g., the header of a data frame) is stored using a relatively small number of bits per memory cell. Data that is more error-tolerant (e.g., the main data representing music, images, or video) is stored using a relatively large number of bits per memory cell.

U.S. Patent 6,643,169, whose disclosure is incorporated herein by reference, describes  
15 a method for storing data in varying numbers of bits per cell depending on the type of data involved. The number of bits may be increased per cell when fidelity is less important. The number of bits per cell may be decreased when fidelity is more important. A memory, in some embodiments, may change between storage modes on a cell by cell basis.

## 20 SUMMARY OF THE INVENTION

Embodiments of the present invention provide a method for data storage in a memory that includes a plurality of analog memory cells, the method including:

estimating respective achievable storage capacities of the analog memory cells;

25 assigning the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities;

storing the data in the memory cells in accordance with the respective assigned storage configurations; and

re-estimating the respective achievable storage capacities of the analog memory cells after the memory has been installed in a host system and used for storing the data in the host  
30 system, and modifying the storage configurations responsively to the re-estimated achievable capacities.

In some embodiments^ storing the data includes encoding the data using an Error Correcting Code (ECC), converting the encoded data to analog values selected from a set of nominal analog values and writing the analog values into the respective memory cells, and

each storage configuration specifies a respective ECC code rate and a size of a set of nominal analog values used for storing the data.

In an embodiment, for each analog memory cell, storing the data includes converting the data to analog values selected from a set of nominal analog values and writing the analog values into the memory cell, and each storage configuration specifies the nominal analog values used for storing the data in the analog memory cell.

In another embodiment, storing the data includes converting the data to analog values and writing the analog values into the respective memory cells, and estimating the achievable storage capacities includes estimating respective distortion levels affecting the analog values written into the memory cells and determining the achievable storage capacities responsively to the distortion levels.

Estimating the distortion levels may include reading the analog values from the memory cells and estimating the distortion levels based on the analog values read from the memory cells. Estimating the distortion levels may include reconstructing the data from the analog values read from the memory cells and calculating the distortion levels based on the reconstructed data. In some embodiments, calculating the distortion levels includes summing scalar functions of the analog values read from the memory cells. Summing the scalar functions may include summing squares of differences between the analog values read from the memory cells and respective expected values of the data stored in the memory cells responsively to the data. In another embodiment, calculating the distortion levels includes counting a number of memory cells in which a difference between the analog values read from the memory cells and respective expected values of the data stored in the memory cells responsively to the data exceeds a predetermined value.

In a disclosed embodiment, storing the data includes writing the analog values into the memory cells and verifying the written analog values using an iterative Program and Verify (P&V) process, and estimating the distortion levels includes calculating the distortion levels based on the analog values verified by the P&V process.

In another embodiment, storing the data includes writing the analog values into the memory cells and verifying the written analog values using an iterative Program and Verify (P&V) process, which iteratively increments the analog values by a programming step, and the storage configurations define a size of the programming step used by the iterative P&V process.

In yet another embodiment, storing the data includes encoding the data to be stored in a group of the memory cells using an Error Correcting Code (ECC) and converting the

encoded data to analog values for storage in the respective analog memory cells of the group, estimating the achievable storage capacities includes reading the analog values from the memory cells in the group and evaluating a distance metric between the read analog values and a valid codeword of the ECC that is nearest to the analog values, and modifying the storage configurations includes adapting the storage configuration of the group of the memory cells responsively to the evaluated distance metric.

In still another embodiment, re-estimating the achievable storage capacities includes reading the data stored in the memory cells, and modifying the storage configurations includes adapting the storage configurations responsively to detecting errors in the read data. Storing the data may include encoding the data using an Error Correcting Code (ECC), and adapting the storage configurations may include modifying a code rate of the ECC.

In some embodiments, estimating the achievable capacities of the memory cells includes tracking previous programming and erasure operations applied to the respective memory cells and estimating the achievable capacities responsively to the tracked previous programming and erasure operations. Estimating the achievable capacities may include calculating the achievable capacities responsively to lengths of time periods that elapsed since the previous programming and erasure operations.

In an embodiment, storing the data includes converting an initial portion of the data to analog values and writing the analog values into the respective memory cells, and subsequently storing an additional portion of the data in at least some of the memory cells by increasing the analog values written to the memory cells without erasing the cells.

In another embodiment, storing the data includes encoding the data using an Error Correcting Code (ECC) that adds redundancy bits to the data and storing the redundancy bits in some of the analog memory cells, and modifying the storage configurations includes modifying a number of the redundancy bits added by the ECC without erasing the cells. Storing the redundancy bits may include storing the redundancy bits separately from the data.

In yet another embodiment, assigning the storage configurations includes predefining a set of possible storage configurations, and modifying the storage configurations includes selecting updated storage configurations from the predefined set. The quantity of the data defined by at least one of the storage configurations may specify a non-integer number of bits per cell.

In still another embodiment, storing the data includes compressing the data before writing the data to the memory cells, and the storage configurations define respective compression ratios by which the data is to be compressed. In a disclosed embodiment,

assigning the storage configurations includes defining the storage configurations for respective groups of the memory cells.

In an embodiment, storing the data in the memory cells includes accepting a data item for storage in the memory, selecting a subset of the memory cells in which to store the data item responsively to the assigned storage configurations, and storing the data item in the selected subset of the memory cells. In some embodiments, re-estimating the achievable storage capacities is performed during idle time periods in which the data is not stored and read.

In another embodiment, assigning and modifying the storage configurations include storing the storage configurations in a configuration table. In some embodiments, initial values of the storage configurations may be stored in the configuration table before the memory has been installed in the host system. In yet another embodiment, assigning and modifying the storage configurations include maintaining a predetermined margin between the quantities of the data stored in the memory cells and the respective estimated achievable capacities. In still another embodiment, at least some of the distortion is caused by leakage current in the analog memory cells, and estimating the distortion includes estimating the leakage current.

There is additionally provided, in accordance with an embodiment of the present invention, a method for data storage in a memory that includes a plurality of analog memory cells, the method including:

estimating distortion levels in the respective analog memory cells;

estimating respective achievable storage capacities of the analog memory cells based on the estimated distortion levels;

assigning the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities; and

storing the data in the memory cells in accordance with the respective assigned storage configurations.

There is also provided, in accordance with an embodiment of the present invention, a method for data storage in a memory that includes a plurality of analog memory cells, the method including:

tracking respective achievable storage capacities of the analog memory cells while the memory is in use in a host system;

accepting data for storage in the memory;

selecting a subset of the memory cells for storing the data based on the tracked achievable capacities; and

storing the data in the memory cells of the subset.

In some embodiments, selecting the subset includes selecting the memory cells whose aggregate achievable capacity best matches a size of the accepted data. In another embodiment, accepting the data includes accepting a requested reliability level for storing the data, and selecting the subset includes selecting the memory cells responsively to the requested reliability level. Additionally or alternatively, accepting the data includes accepting a requested retention period for storing the data, and selecting the subset includes selecting the memory cells responsively to the requested retention period. Selecting the subset may include selecting the memory cells having low distortion levels with respect to other memory cells.

In another embodiment, tracking the achievable storage capacities includes tracking previous programming and erasure operations applied to the memory cells, and selecting the subset includes selecting the memory cells responsively to the previous programming and erasure operations. Selecting the subset may include distributing a number of programming and erasure operations evenly over the memory cells in the memory by selecting the memory cells having a smaller number of the previous programming and erasure operations with respect to other memory cells. Distributing the number of programming and erasure operations may include selecting the memory cells having the smaller number of the previous programming and erasure operations performed during a predetermined recent time period with respect to the other memory cells.

In a disclosed embodiment, tracking the achievable storage capacities includes calculating and reporting to the host system a size of an available memory space by summing the achievable capacities of the memory cells that are available for storing the data.

In another embodiment, the memory is divided into multiple erasure blocks, each erasure block including a group of the memory cells that are erased in a single erasure operation, and the method includes erasing a data item from the memory by:

identifying one or more of the erasure blocks in which the data item is stored;

when the erasure blocks in which the data item is stored contain stored data of another data item, identifying alternative memory cells outside the erasure blocks in which the data item is stored based on the tracked achievable capacities of the memory cells and on a size of the other data item, and copying the data of the other data item to the alternative memory cells; and

erasing the erasure blocks in which the data item is stored.

In an embodiment, accepting the data for storage includes accepting the data in fixed-capacity blocks from the host system and storing the data in variable-capacity groups of the memory cells based on the tracked achievable capacities of the variable-size groups. In another embodiment, storing the data in the variable-capacity groups includes identifying two or more variable-capacity groups having capacities that are smaller than a capacity of the fixed-capacity block, and storing the data accepted in one or more of the fixed-capacity blocks in the identified two or more variable-capacity groups.

In yet another embodiment, storing the data in the variable-capacity groups includes allocating one or more of the variable-capacity groups to serve as an excess memory storage area, storing the data accepted in the fixed-capacity blocks in respective variable-capacity groups, and, when the achievable capacity of a variable-capacity group is smaller than a capacity of the fixed-capacity blocks, storing some of the data accepted in the respective fixed-capacity block in the allocated excess memory storage area. In still another embodiment, storing the data in the variable-capacity groups includes storing the data accepted in the fixed-capacity blocks sequentially in the memory cells irrespective of boundaries between the variable-capacity groups. Storing the data may include initially storing the data accepted in each fixed-capacity block in a respective variable-capacity group, and subsequently re-programming the initially-stored data sequentially, irrespective of the boundaries between the variable-capacity groups.

In some embodiments, the method includes retrieving the data from the variable-size groups, reconstructing the fixed-size blocks and outputting the data using the fixed-size blocks to the host system. The method may include reducing a number of memory access operations applied to the memory by caching at least some of the data. Caching the at least some of the data may include transferring the cached data to the memory upon detecting an immediate storage event. The immediate storage event may include at least one event selected from a group of events consisting of an approaching power failure, a time out and an End Of File (EOF) command accepted from the host system.

There is also provided, in accordance with an embodiment of the present invention, data storage apparatus, including:

an interface, which is arranged to communicate with a memory that includes a plurality of analog memory cells; and

a memory signal processor (MSP), which is arranged to estimate respective achievable storage capacities of the memory cells, to assign the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the

estimated achievable capacities, to store the data in the memory cells in accordance with the respective assigned storage configurations, to re-estimate the respective achievable storage capacities of the analog memory cells after the memory has been installed in a host system and used for storing the data in the host system, and to modify the storage configurations  
5 responsively to the re-estimated achievable capacities.

There is further provided, in accordance with an embodiment of the present invention, data storage apparatus, including:

an interface, which is arranged to communicate with a memory that includes a plurality of analog memory cells; and

10 a memory signal processor (MSP), which is arranged to estimate distortion levels in the respective analog memory cells, to estimate respective achievable storage capacities of the analog memory cells based on the estimated distortion levels, to assign the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities, and to store the data in the memory cells in  
15 accordance with the respective assigned storage configurations.

There is also provided, in accordance with an embodiment of the present invention, data storage apparatus, including:

an interface, which is arranged to communicate with a memory that includes a plurality of analog memory cells; and

20 a memory signal processor (MSP), which is arranged to track respective achievable storage capacities of the memory cells while the memory is in use in a host system, to accept data for storage in the memory, to select a subset of the memory cells for storing the data based on the tracked achievable capacities, and to store the data in the memory cells of the subset.

25 There is additionally provided, in accordance with an embodiment of the present invention, data storage apparatus, including:

a memory, which includes a plurality of analog memory cells; and

a memory signal processor (MSP), which is coupled to the memory and is arranged to estimate respective achievable storage capacities of the memory cells, to assign the memory  
30 cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities, to store the data in the memory cells in accordance with the respective assigned storage configurations, to re-estimate the respective achievable storage capacities of the analog memory cells after the memory has been installed

in a host system and used for storing the data in the host system, and to modify the storage configurations responsively to the re-estimated achievable capacities.

In some embodiments, the memory resides in a first Integrated Circuit (IC) and the MSP resides in a second IC different from the first IC. In an alternative embodiment, the memory and the MSP are integrated in a single device. Further alternatively, the MSP may be embodied in a processor of the host system. In some embodiments, the memory cells may include Flash memory cells, Dynamic Random Access Memory (DRAM) cells, Phase Change Memory (PCM) cells, Nitride Read-Only Memory (NROM) cells, Magnetic Random Access Memory (MRAM) cells or Ferroelectric Random Access Memory (FRAM) cells.

There is additionally provided, in accordance with an embodiment of the present invention, data storage apparatus, including:

a memory, which includes a plurality of analog memory cells; and

a memory signal processor (MSP), which is coupled to the memory and is arranged to track respective achievable storage capacities of the memory cells while the memory is in use in a host system, to accept data for storage in the memory, to select a subset of the memory cells for storing the data based on the tracked achievable capacities, and to store the data in the memory cells of the subset.

The present invention will be more fully understood from the following detailed description of the embodiments thereof, taken together with the drawings in which:

## BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram that schematically illustrates a memory system, in accordance with an embodiment of the present invention;

Fig. 2 is a diagram that schematically illustrates a memory cell array, in accordance with an embodiment of the present invention;

Fig. 3 is a graph showing voltage distribution in an array of multi-level memory cells, in accordance with an embodiment of the present invention;

Fig. 4 is a graph showing cell capacity distribution in a memory cell array, in accordance with an embodiment of the present invention;

Fig. 5 is a diagram showing block and page structure in a memory cell array, in accordance with an embodiment of the present invention;

Figs. 6-11 are flow charts that schematically illustrate methods for adaptively modifying memory device storage density, in accordance with embodiments of the present invention;

Fig. 12 is a flow chart that schematically illustrates a method for storing data in a variable-capacity memory device, in accordance with an embodiment of the present invention;

Fig. 13 is a flow chart that schematically illustrates a method for erasing data from a variable-capacity memory device, in accordance with an embodiment of the present invention;

5 and

Figs. 14-17 are diagrams that schematically illustrate methods for storing fixed-size data blocks in a variable-capacity memory device, in accordance with embodiments of the present invention.

## DETAILED DESCRIPTION OF EMBODIMENTS

10

### OVERVIEW

The achievable storage capacity of analog memory cell arrays, such as Flash memories, varies over time. In many cases, the achievable capacity of a memory cell array decreases over time because of repeated programming and erasure operations, aging and other effects. The capacity variations often differ from cell to cell in the array. On the other hand,  
15 memory devices are usually specified to provide a certain capacity at a certain reliability level, and to meet these specifications for a certain data retention period.

Some known memory devices and associated storage methods handle varying cell capacity by over-designing or over-specifying the size and performance of the memory cell array. In such devices, the number of memory cells and their performance are specified with  
20 margins that are sufficiently large to keep the device within specification, given the anticipated degradation over the desired retention period. Such design methods are often inefficient, increase cost and complexity and reduce the usable capacity of the memory device.

Embodiments of the present invention provide improved methods and systems for storing data in memory devices, by adaptively modifying the density in which data is stored.

25 In some embodiments, a Memory Signal Processor (MSP) stores data in a memory device that includes a plurality of analog memory cells, such as a Flash or DRAM device. The MSP continually estimates the achievable storage capacities of the memory cells. The MSP may estimate the achievable capacity based on factors such as the distortion levels in the memory cells, reliability and/or retention requirements specified for the data, errors detected  
30 in the data that is programmed in the cells and/or the history of previous programming and erasure operations the cells have gone through.

Based on the estimated achievable capacities, the MSP assigns the memory cells respective storage configurations. Each storage configuration defines the amount of data that

can be stored in the respective memory cells. For example, the storage configuration may define a number of nominal voltage levels and an Error Correction Coding (ECC) scheme used for storing data in the memory cells. In some embodiments, the MSP selects the appropriate storage configuration from a limited set of predefined configurations.

5           The MSP stores the data in the memory cells in accordance with the storage configurations assigned to the cells. The MSP adaptively modifies the storage configurations over time, throughout the lifetime of the memory device, so that the density in which data is stored in the memory cells matches the achievable capacities of the cells.

10           Unlike some known memory devices and data storage methods in which the storage density is specified *a-priori*, the methods and systems described herein track the achievable capacity of the different memory cells throughout the life cycle of the memory array and continually adapt the density in which data is stored accordingly. As a result, the achievable capacity of the memory device can be exploited optimally as it changes over time, thus reducing size and cost.

15           The MSP interacts with a memory controller or other host, i.e., accepts data for storage from the memory controller and retrieves data that is stored in memory when requested by the memory controller. In some embodiments, the MSP emulates fixed-capacity operation towards the memory controller, even though it stores the data in cell groups having variable capacity. Such methods are used, for example, when the MSP interacts with legacy memory  
20           controllers, which are designed to control memory devices having fixed capacity and fixed page size. Thus, by mediating between the fixed-capacity operation of the memory controller and the adaptive-capacity characteristics of the storage methods described herein, legacy memory controllers can be used without modification.

## SYSTEM DESCRIPTION

25           Fig. 1 is a block diagram that schematically illustrates a memory system 20, in accordance with an embodiment of the present invention. System 20 can be used in various host systems and devices, such as in computing devices, cellular phones or other communication terminals, removable memory modules (e.g. "disk-on-key" devices), digital cameras, music and other media players and/or any other system or device in which data is  
30           stored and retrieved. In a typical application, memory system 20 interacts with a memory controller 22, i.e., accepts data for storage from the memory controller and outputs data that is stored in memory to the memory controller when requested.

System 20 comprises a memory device 24, which stores data in a memory cell array 28. The memory array comprises multiple analog memory cells 32. In the context of the present patent application and in the claims, the term "analog memory cell" is used to describe any memory element that holds a continuous, analog value of a physical parameter, such as an electrical voltage or charge. Array 28 may comprise analog memory cells of any kind, such as, for example, NAND and NOR Flash cells, PCM, NROM, FRAM, MRAM and DRAM cells. The charge levels stored in the cells and/or the analog voltages written into and read out of the cells are referred to herein collectively as analog values.

Data for storage in memory device 24 is provided to the device and cached in data buffers 36. The data is then converted to analog voltages and written into memory cells 32 using a reading/writing (R/W) unit 40, whose functionality is described in greater detail below. When reading data out of array 28, unit 40 converts the electric charge, and thus the analog voltages of memory cells 32, into digital samples, each having a resolution of one or more bits. The samples produced by unit 40 are cached in buffers 36. The operation and timing of memory device 24 is managed by control logic 48.

The storage and retrieval of data in and out of memory device 24 is performed by a Memory Signal Processor (MSP) 52. MSP 52 intermediates between memory device 24 and memory controller 22 or other host. As will be shown in detail hereinbelow, MSP 52 uses novel methods for adaptively estimating the achievable capacity of memory cells 32, and for managing the storage and retrieval of data in memory array 28 based on the estimated cell capacities. MSP 52 adaptively modifies the density in which data is stored in array 28, so as to provide optimal capacity throughout the lifetime of the system.

MSP 52 can write data into memory cells 32 at different densities, depending on the achievable capacity of the cells. In the context of the present patent application and in the claims, the term "capacity" is used to describe the amount of information (usually expressed in bits-per-cell units) that a given cell or group of cells is capable of holding reliably. The term "density" is used to describe the amount of information, usually in bit-per-cell units, which is actually stored in a cell or group of cells. Thus, information can be stored reliably in a group of cells at different densities, as long as the capacity is not exceeded. Thus, the storage density of a cell or group of cells is by definition no greater than its capacity, assuming the data is stored reliably.

MSP 52 can vary the density in which data is stored in cells 32, for example by changing the number of voltage levels used for writing the data into cells 32. Using a higher number of voltage levels enables higher information density. In some embodiments, MSP 52

encodes the data stored in device 24 using an Error Correcting Code (ECC). MSP 52 comprises an encoder/decoder 64, which encodes the data to be written to device 24 and decodes the data read out of device 24. In some implementations, MSP 52 varies the information density by selecting different codes or varying the code rate of the ECC.

5 A signal processing unit 60 processes the data that is written into and retrieved from device 24. In particular, unit 60 estimates the distortion that is present in the voltages read out of cells 32. The estimated distortion is used to estimate the achievable capacity of cells 32, and to determine the desired storage density to be used. MSP 52 comprises a data buffer 72, which is used by unit 60 for storing data and serves as an interface for communicating with  
10 memory device 24. MSP 52 also comprises an Input/Output (I/O) buffer 56, which forms an interface between the MSP and the memory controller.

A memory management unit 76 manages the operation and timing of MSP 52. In particular, unit 76 tracks the estimated capacity and status of each memory cell or group of cells, so that MSP 52 stores data in each cell or group at a density that matches its estimated  
15 capacity. Signal processing unit 60 and management unit 76 may be implemented in hardware. Alternatively, unit 60 and/or unit 76 may comprise microprocessors that run suitable software, or a combination of hardware and software elements. In some embodiments, encoder/decoder 64, or parts thereof, can also be implemented in software. In some embodiments, the functionality of MSP 52 can sometimes be implemented in software  
20 and carried out by a suitable processor or other element of the host system.

The configuration of Fig. 1 is an exemplary system configuration, which is shown purely for the sake of conceptual clarity. Any other suitable configuration can also be used. Elements that are not necessary for understanding the principles of the present invention, such as various interfaces, addressing circuits, timing and sequencing circuits, data scrambling  
25 circuits and debugging circuits, have been omitted from the figure for clarity.

In some embodiments, memory device 24 and MSP 52 are implemented in two separate Integrated Circuits (ICs). In alternative embodiments, however, the memory device and MSP may be integrated in a single IC or System on Chip (SoC). In some implementations, a single MSP 52 may be connected to multiple memory devices 24.  
30 Additional architectural aspects of certain embodiments of system 20 are described in greater detail in U.S. Provisional Patent Application 60/867,399, cited above. Alternatively, some or all of the functionality of memory controller may be embodied in MSP 52, as will be explained in detail further below.

In a typical writing operation, data to be written into memory device 24 is accepted from memory controller 22 and cached in I/O buffer 56. Memory management unit 76 determines the cells into which the data is to be written. Unit 76 also determines a storage configuration (e.g., a number of voltage levels and a code rate) to be used, based on the estimated capacity of the designated cells. Encoder/decoder 64 encodes the data, and the encoded data is transferred, via data buffers 72, to memory device 24. In device 24 the data is temporarily stored in buffers 36. R/W unit 40 converts the data to analog voltage values and writes the data into the appropriate cells 32 of array 28.

In a typical reading operation, R/W unit 40 reads analog voltage values out of the appropriate memory cells 32 and converts the voltages to digital samples. The samples are cached in buffers 36 and transferred to buffers 72 of MSP 52. Blocks of data are transferred from buffers 72 to unit 60, and encoder/decoder 64 decodes the ECC of these blocks. The decoded data is transferred via I/O buffer 56 to memory controller 22.

Memory controller 22 may sometimes comprise a known device such as the PS8000 NAND Flash controller device offered by Phison Electronics Corp. (Chutung, Hsinchu, Taiwan). This device is specified in "PS8000 Controller Specification," revision 1.2, March 28, 2007, which is incorporated herein by reference. Another known memory controller is the Databahn™ Flash memory controller IP, which is offered by Denali Software, Inc. (Palo Alto, California). Details regarding this product are available at [www.denali.com/products/databahn\\_flash.html](http://www.denali.com/products/databahn_flash.html). Yet another example is the FlashFX Pro® Flash media manager offered by Datalight, Inc. (Bothell, Washington). This device is described in "FlashFX Pro 3.1 High Performance Flash Manager for Rapid Development of Reliable Products," November 16, 2006, which is incorporated herein by reference.

When memory controller 22 comprises a conventional memory controller device, it often has predefined interfaces and protocols for communicating with memory devices. These interfaces and protocols are usually designed for communicating with memory devices having fixed storage capacities. In some embodiments, MSP 52 emulates the data structures, interfaces and protocols used by memory controller 22, thus intermediating between the adaptive capacity characteristics of system 20 and the fixed capacity characteristics of memory controller 22. Several exemplary methods are described in Figs. 14-17 below.

In alternative embodiments, some of the memory management functions are implemented in MSP 52. In these embodiments, the MSP accepts data items for storage and may partition and store them in variable capacity pages, without having to emulate fixed capacity operation.

The interface between MSP 52 and memory device 24, and/or the interface between MSP 52 and memory controller 22, may conform to a known standard or protocol. For example, one or both of these interfaces may conform to the Open NAND Flash Interface (ONFI) Specification. The ONFI specification is defined in "Open NAND Flash Interface Specification," revision 1.0, December 28, 2006, which is incorporated herein by reference.

#### MEMORY ARRAY STRUCTURE AND DISTORTION MECHANISMS

Fig. 2 is a diagram that schematically illustrates memory cell array 28, in accordance with an embodiment of the present invention. Although Fig. 2 refers to Flash memory cells that are connected in a particular array configuration, the principles of the present invention are applicable to other types of memory cells and other array configurations, as well. Some exemplary cell types and array configurations are described in the references cited in the Background section above.

Memory cells 32 of array 28 are arranged in a grid having multiple rows and columns. Each cell 32 comprises a floating gate Metal-Oxide Semiconductor (MOS) transistor. A certain amount of electrical charge (electrons or holes) can be stored in a particular cell by applying appropriate voltage levels to the transistor gate, source and drain. The value stored in the cell can be read by measuring the threshold voltage of the cell, which is defined as the minimal voltage that needs to be applied to the gate of the transistor in order to cause the transistor to conduct. The read threshold voltage is proportional to the charge stored in the cell.

In the exemplary configuration of Fig. 2, the gates of the transistors in each row are connected by word lines 80. The sources of the transistors in each column are connected by bit lines 84. In some embodiments, such as in some NOR cell devices, the sources are connected to the bit lines directly. In alternative embodiments, such as in some NAND cell devices, the bit lines are connected to strings of floating-gate cells.

Typically, R/W unit 40 reads the threshold voltage of a particular cell 32 by applying varying voltage levels to its gate (i.e., to the word line to which the cell is connected) and checking whether the drain current of the cell exceeds a certain threshold (i.e., whether the transistor conducts). Unit 40 usually applies a sequence of different voltage values to the word line to which the cell is connected, and determines the lowest gate voltage value for which the drain current exceeds the threshold. Typically, unit 40 reads a group of cells from a certain row simultaneously.

In some embodiments, unit 40 measures the drain current by pre-charging the bit line of the cell to a certain voltage level. Once the gate voltage is set to the desired value, the drain current causes the bit line voltage to discharge through the cell. Unit 40 measures the bit line voltage several microseconds after the gate voltage is applied, and compares the bit line voltage to the threshold. In some embodiments, each bit line 84 is connected to a respective sense amplifier, which amplifies the bit line current and converts it to a voltage. The voltage is compared to the threshold using a comparator.

The voltage reading method described above is an exemplary method. Alternatively, R/W unit 40 may use any other suitable method for reading the threshold voltages of cells 32. For example, unit 40 may comprise one or more Analog to Digital Converters (ADCs), which convert the bit line voltages to digital samples.

In some embodiments, entire pages (rows) are written and read in parallel. Erasing of cells is usually carried out in blocks that contain multiple pages. Typical memory devices may comprise several hundred thousand pages, each comprising several thousand cells (e.g., 16K or 32K cells per page). A typical erasure block is on the order of 128 pages, although other block sizes can also be used.

The voltages digitized by R/W unit 40 may contain various types of distortion, which are caused by different distortion mechanisms in array 28. Some distortion mechanisms affect the actual electrical charge that is stored in the cells, while other mechanisms distort the sensed voltages. For example, electrical coupling between adjacent cells in the array may modify the threshold voltage in a particular cell. This effect is referred to as interference noise. As another example, electrical charge may leak from the cells over time. As a result of this aging effect, the threshold voltage of the cells may drift over time from the initially-written value.

Another type of distortion, commonly referred to as disturb noise, is caused by a read, write or erase operation on a certain cell in the array, which causes unintended programming or erasure of other cells. As yet another example, the source-drain current of a particular cell can be affected by the charge in adjacent cells, e.g., other cells in the same NAND cell string, via an effect referred to as back pattern dependency.

### 30 DISTORTION/CAPACITY TRADE-OFF

Fig. 3 is a graph showing voltage distributions in memory cell array 28, in accordance with an embodiment of the present invention. Fig. 3 is used to demonstrate the effect of distortion on the achievable capacity of the memory cells. The memory cells whose voltage

distribution is shown in Fig. 3 are configured to each store two bits of information using four nominal voltage levels denoted 90A...90D.

In order to store two data bits in a memory cell, R/W unit 40 writes one of the four nominal voltage levels into the cell. In the present example, voltage level 90A corresponds to "11" bit values, and voltage levels 90B...90D correspond to "01", "00" and "10" bit values, respectively.

Although the R/W unit writes a particular nominal voltage level, the actual threshold voltage level of the cell usually deviates from the nominal level, because of the different distortion mechanisms and because of limited accuracy in the programming and erasing of the cells. Plots 92A...92D show an exemplary voltage distribution. Plot 92A shows the distribution of voltages in the cells that store "11" bit values. Plots 92B, 92C and 92D show the voltage distribution in the cells that store "01", "00" and "10" bit values, respectively. Note that the voltage distribution at the time of reading may differ from the voltage distribution at the time of writing, due to aging and other distortion mechanisms.

The total range of threshold voltages is divided into four intervals 96A...96D by defining three thresholds 94A...94C. When reading the threshold voltage levels of the memory cells, voltage levels that fall within interval 96A are assumed to correspond to "11" bit values. Similarly, voltage levels that fall within intervals 96B...96D are assumed to correspond to "01", "00" and "10" bit values, respectively. Intervals 96A...96D are therefore often referred to as decision intervals, and thresholds 94A...94C are referred to as decision thresholds.

When the voltage that is written in response to certain data bit values falls within the wrong decision interval, the data read out of the cell differs from the data that was written, resulting in a reading error. For a given distortion level, there is a trade-off between the error performance and the achievable capacity of the cells. When the number of nominal voltage levels increases (i.e., the storage density increases), the size of the decision intervals becomes smaller with respect to the distortion effects. As a result, the read voltages are more likely to drift or otherwise fall within erroneous decision intervals, and the error probability increases. For a given error probability, the achievable capacity of a cell drops as the level of distortion increases.

#### VARYING CELL CAPACITY ACROSS THE MEMORY CELL ARRAY

As noted above, the achievable capacity of a memory cell (i.e., the number of information bits that can be reliably programmed in the cell) depends on the level of distortion

of the particular cell. In many practical cases, the achievable capacity differs considerably from one cell to another.

Different cells 32 in array 28 often have different distortion characteristics. For example, some of the cells, such as the cells located on the external perimeter of an erasure block, have fewer neighboring cells, and therefore often have a lower level of cross-coupling noise from adjacent cells. As another example, since programming cells causes interference, disturbs and back pattern dependency errors to previously-programmed cells, cells that are programmed later may suffer from a smaller level of errors than the cells that are programmed earlier. As yet another example, a certain sense amplifier or ADC may introduce a particularly high noise level, which causes the cells along a particular bit line to have higher distortion in comparison with other bit lines. Other distortion characteristics may depend on the history of the programming and erasure operations performed in a particular cell or its environment. Because of these and other factors, the distortion in cells 32, and consequently the achievable capacity of the cells, often varies significantly across array 28.

Fig. 4 is a graph showing an exemplary cell capacity distribution in a memory cell array, in accordance with an embodiment of the present invention. A plot 100 shows the achievable capacity of the different cells. In the present example, the achievable capacity varies between approximately 4.2 and 7.3 bits per cell. A plot 104 shows a practical storage density that may be used for programming the cells, assuming a certain safety margin to guarantee reliable performance over the desired data retention period. As can be seen in the figure, when storage density is allocated individually to each memory cell in accordance with plot 104, the average density achieved is higher than five bits per cell.

If the memory cells of the array of Fig. 4 were allocated fixed storage densities, the entire array would have to be specified as having only three bits per cell, even though most of the cells are capable of achieving much higher capacities. In some known memory devices that use Bad Block Management (BBM) methods, the array could be specified as a four bit-per-cell array, with some groups of cells recognized as faulty. BBM methods are described, for example, in "Bad Block Management in NAND Flash Memories," Application note AN-1819 by STMicroelectronics (Geneva, Switzerland), November 29, 2004, which is incorporated herein by reference.

The methods and systems described herein store data in each cell or group of cells at a density that matches the achievable capacity of the cells. Cells having a lower distortion level are assigned higher capacity, while cells that suffer stronger distortion are assigned lower capacity. The methods and systems described herein track the changes in distortion and

capacity during the entire lifetime of the array, and adaptively modify the density in which data is stored in the different cells accordingly.

Cell capacity can be estimated and storage density can be assigned to individual cells or collectively to groups of cells. Assigning separate densities to small groups of cells enables a finer matching of the storage density to the achievable capacity at the expense of more complex management, and *vice versa*.

Fig. 5 is a diagram showing block and page structure in a memory cell array 110, in accordance with an embodiment of the present invention. Array 110 comprises multiple rows 114. A memory page is defined as a group of memory cells that are written (programmed) simultaneously. In some cases, each page comprises an entire row of memory cells. In other cases, each row may be divided into several pages. Data is erased from array 110 by simultaneously erasing groups of pages, referred to as erasure blocks 118. For example, a typical page may comprise 16,384 cells and a typical erasure block may comprise one hundred and twenty eight pages, although other sizes can also be used. In some embodiments, each page can be divided into several sectors (not shown in the figure).

In the description that follows, the basic capacity and density allocation unit is a page. In other words, all memory cells in a particular page are programmed using the same number of voltage levels and using the same ECC. Different pages can be assigned different densities by adapting the number of voltage levels and/or the ECC. In alternative embodiments, capacity estimation and density allocation can be performed using groups of cells having any desired granularity, such as on a page-by-page, block-by-block, sector-by-sector or even cell-by-cell basis. In some embodiments, capacity estimation and density allocation is performed jointly for cells connected to a particular bit line 84 or word line 80.

#### ADAPTIVE DENSITY ALLOCATION METHODS

The description that follows illustrates several exemplary methods for estimating the achievable capacity of groups of memory cells 32, e.g. pages, and for adaptively modifying the storage density of memory device 24, in accordance with embodiments of the present invention.

In some embodiments, management unit 76 in MSP 52 holds a storage configuration table, which holds, for each page, a set of parameters that is referred to as a storage configuration. The storage configuration of a particular page determines how data is stored in the page, and may comprise parameters such as the number of voltage levels used for storing data in the memory cells of the page, the values of these voltage levels, the ECC used for

encoding the data in the page, and/or any other suitable parameter. In some cases, such as when using Trellis Coded Modulation (TCM) or other coded modulation methods, the selection of ECC and voltage levels is combined. In addition to the storage configuration, the storage configuration table may hold parameters such as the estimated distortion level and the estimated achievable capacity of the page. In some embodiments, initial values of the storage configurations can be stored in the storage configuration table during production.

When writing data into a particular page, MSP 52 queries the storage configuration of the page. The MSP encodes the data using the appropriate ECC and maps the encoded data to the appropriate number of voltage levels, as indicated by the storage configuration of the page.

In some embodiments, when reading data from a particular page, the MSP converts the samples produced by R/W unit 40 to hard bit decisions, in accordance with the number of levels indicated by the storage configuration of the page. The MSP then configures the ECC decoder to the ECC indicated in the storage configuration, so as to decode the ECC. Alternatively, e.g., when the decoder comprises a soft decoder, the MSP decodes the samples produced by R/W unit 40 using the soft decoder without generating hard decisions.

The use of different numbers of voltage levels may have an impact on the design of R/W unit 40, such as on the design of comparators, ADCs, DACs, sense amplifiers and associated circuitry. In some embodiments, the sensitivity and resolution of the R/W unit is designed to match the maximum number of voltage levels used. In alternative embodiments, unit 40 can read the threshold voltages from a page in several iterations of increasing resolution (e.g., by using different thresholds in each iteration). Management unit 76 can determine the desired number of iterations, and the resulting resolution, based on the number of levels used for storage in the particular page.

Fig. 6 is a flow chart that schematically illustrates a method for adaptively modifying the storage density of memory device 24, in accordance with an embodiment of the present invention. The method begins with MSP 52 estimating the distortion levels in memory cells 32, at a distortion estimation step 130. MSP 52 can use different methods for estimating the cell distortion.

For example, in order to estimate the distortion in a particular target cell, the MSP can read the voltages from neighboring cells, estimate the cross-coupling ratios between the neighboring cells and the target cell, and calculate the cumulative coupling noise contributed by the neighboring cells.

In some embodiments, the MSP can estimate the distortion in a certain page following a programming operation applied to the page. Fig. 7 below shows an exemplary post-writing

estimation method. Alternatively, the MSP can estimate the distortion in a certain page following a reading operation applied to the page. Fig. 8 below shows an exemplary post-reading estimation method. Capacity can also be estimated by verifying the reliability of the programmed data, for example by applying increasingly-stronger error correcting codes until error-free storage is achieved. Such a method is described in Fig. 9 below.

Further alternatively, the distortion level can be estimated based on the history of programming and erasure operations the page has gone through. The programming and erasure history can affect the distortion level in a number of ways. Generally, older cells (i.e., cells that have recently gone through a large number of programming and erasing operations) are considered to have a higher leakage current level than younger cells that went through a smaller number of programming and erase operations. Thus, the distortion is also affected by the time that elapsed since the previous programming and erasure operations. An exemplary method for estimating cell capacity based on the history of programming and erasure operations is described in Fig. 11 below.

MSP 52 estimates the page capacity and determines the storage configuration to be used in each page, at a storage configuration calculation step 134. The MSP can use any suitable method for estimating the achievable page capacity based on the estimated distortion, and for selecting the appropriate storage configuration (e.g., number of voltage levels and ECC).

Usually, the MSP leaves a certain safety margin between the estimated capacity and the storage density it sets. The safety margin, which is also referred to as a performance margin, is set to ensure reliable operation over a desired data retention period. Several exemplary methods for determining the appropriate storage configuration are described further below. Unit 76 of MSP 52 updates the storage configuration table with the newly-calculated storage configurations. MSP 52 writes subsequent data to the pages of array 28 using the updated storage configurations, at a writing step 138.

#### P&V-ASSISTED DENSITY ALLOCATION

Fig. 7 is a flow chart that schematically illustrates a method for adaptively modifying the storage density of memory device 24 following a Program and Verify (P&V) process, in accordance with another embodiment of the present invention.

P&V processes are commonly used for programming memory cells. In a typical P&V process, a cell is programmed by applying a sequence of voltage pulses, whose voltage level increases from pulse to pulse. The programmed voltage level is read ("verified") after each

pulse, and the iterations continue until the desired level is reached or until the operation times out. P&V processes are described, for example, by Jung et al., in "A 117mm<sup>2</sup> 3.3V Only 128Mb Multilevel NAND Flash Memory for Mass Storage Applications," IEEE Journal of Solid State Circuits, (11:31), November, 1996, pages 1575-1583 and by Takeuchi et al., in "A Multipage Cell Architecture for High-Speed Programming Multilevel NAND Flash Memories," IEEE Journal of Solid-State Circuits, (33:8), August 1998, pages 1228-1238, both of which are incorporated herein by reference.

The method of fig. 7 can be combined with the normal operation of system 20. In this mode of operation, after a particular group of pages (e.g., a single page or a NAND memory block comprising several NAND strings) is programmed, its distortion level is estimated and its storage configuration is updated for use in the next programming operation. The method begins with MSP 52 programming a particular group of pages using a P&V process, at a P&V programming step 150. Following the last P&V iteration, the MSP calculates the distortion statistics, at a post P&V calculation step 154. The last P&V iteration inherently involves reading the programmed voltage levels of the different cells in the page. The MSP may use these read voltage levels to calculate the level of distortion in the page.

In some cases, the reading resolution used for programming verification is not sufficient for reliable distortion estimation. In such cases, the MSP can read the voltage levels of the programmed group of pages at a higher resolution than the reading resolution of the verify operation.

MSP 52 updates the estimated achievable capacity of the page, at a post P&V update step 158. The MSP uses the calculated distortion level to update the storage configuration of the page, i.e., the appropriate number of voltage levels and ECC. The MSP stores the updated configuration in the storage configuration table, for use in the next programming operation of the page.

In some embodiments, MSP 52 may attempt to increase the page density after programming. In these embodiments, the MSP determines whether the programmed data is sufficiently reliable (i.e., the distortion statistics calculated at step 154 indicate that the margin is sufficiently high to ensure low error probability). If sufficient performance margin exists, the MSP modifies the storage configuration of the page to reflect higher density, by increasing the number of voltage levels and/or the code rate.

When the page storage configuration is determined in conjunction with a P&V process, different trade-offs can be made between the parameters of the P&V process and the storage configuration of the page. In the P&V process, programming speed can be traded for

accuracy by varying the voltage increment, or voltage step size, between successive P&V iterations. Setting a higher P&V step size enables programming the page with fewer iterations at the expense of poorer programming accuracy, and *vice versa*. Such trade offs are described, for example, by Suh et al., in "A 3.3V 32Mb NAND Flash Memory with Incremental Step Pulse Programming Scheme," IEEE Journal of Solid-State Circuits, volume 30, number 11, 5 November, 1995, pages 1149-1156, which is incorporated herein by reference.

In some embodiments, MSP 52 sets a particular P&V step size value for each page. The P&V step size is stored as part of the storage configuration of the page in the storage configuration table. When programming a particular page, the MSP queries the storage 10 configuration of the page and uses the appropriate P&V step size.

The P&V step size, the number of voltage levels and the ECC can be selected to provide different performance trade-offs. For example, a larger P&V step size enables faster programming at the expense of a higher error probability. As another example, reduced voltage levels also enable faster programming at the expense of higher error probability, as 15 well as enabling lower disturb levels, interference levels, and wear levels. The error probability can be reduced by reducing the number of voltage levels or by introducing a stronger ECC, at the expense of lower storage density. The choice of ECC may have an effect of the MSP complexity. Thus, storage density and/or error performance can be traded for programming speed.

## 20 DECISION-DIRECTED DENSITY ALLOCATION

Fig. 8 is a flow chart that schematically illustrates a method for adaptively modifying the storage density of memory device 24 following a reading operation, in accordance with yet another embodiment of the present invention.

The method begins with MSP 52 reading a particular memory page from device 24, at 25 a page reading step 170. The MSP calculates the distortion level in the page assuming no errors are present, at a post-read calculation step 174. Typically, the MSP calculates the difference between each soft sample and the corresponding hard decision. Since the assumption is that no decoding errors are present, the difference is attributed to distortion. Note that the hard decisions may be produced by the ECC decoder. The distortion level can be 30 estimated by calculating a sum of scalar functions of the differences between the soft samples and corresponding hard decisions (e.g., a sum of squares of the differences). The distortion can also be estimated by counting the number of cells in which the difference exceeds certain thresholds. In case these thresholds are the threshold levels that are used when reading the

cells to obtain the hard decisions, then the distortion estimate comprises the number of errors that were corrected by the ECC.

The MSP updates the estimated achievable capacity of the page, at a post-read update step 178. The MSP uses the calculated distortion level to update the storage configuration of the page in the storage configuration table, for use in the next programming operation of the page. The method of fig. 8 can be combined with the normal operation of system 20. In this mode of operation, after reading the data a particular page is programmed, its distortion level is estimated and its storage density is updated for use in the next programming operation.

In some embodiments, the MSP can carry out the distortion estimation task during time periods in which the system is idle, in order not to slow down the reading operation.

In some embodiments, the MSP can estimate the distortion level blindly, i.e., without assuming that the hard decisions do not contain errors. For example, when the nominal voltage levels written to the cells are  $4V$  and  $3V_5$  if a large number of the read voltages are around  $2V$ , the MSP may conclude that the level of distortion is high even without decoding the data.

#### DENSITY ADAPTATION BASED ON THE RELIABILITY OF THE PROGRAMMED DATA

In some embodiments, MSP 52 estimates the achievable capacity of a given page by assessing the reliability of the data that is written in the page. In some cases, the MSP adapts the storage configuration of a page without erasing the cells. In these cases, additional data can be written into an existing page in a subsequent programming cycle without erasing the page. The change in density is achieved by modifying the number of voltage levels and/or by modifying the ECC.

Fig. 9 is a flow chart that schematically illustrates an exemplary method for adaptively modifying the storage configuration of memory device 24, in accordance with an embodiment of the present invention. In the present example, the data is programmed using a systematic ECC (i.e., an ECC that adds redundancy bits based on the uncoded information bits, without modifying the uncoded bits). The storage density and error performance are modified by adapting the number of redundancy bits.

The method begins with MSP 52 programming a particular page without coding, at an uncoded programming step 190. Typically but not necessarily, the number of voltage levels is selected so that the storage density slightly exceeds the predefined performance margin, i.e., the page may contain some errors.

The MSP reads and optionally decodes the data from the page, at a decoding step 194, and checks the distortion level in the memory cells, at a distortion level checking step 198. The MSP can use different methods and criteria to determine whether the distortion too high, too low or acceptable with respect to the storage configuration that is currently assigned to the page. For example, the MSP may check whether or not the decoded page contains errors that were not corrected by the ECC.

Alternatively, the MSP can evaluate a distance metric between the set of voltages read from the memory cells of the page and valid codewords of the ECC. The MSP may compare the distance from the read voltages to the correct codeword (i.e., the codeword that was actually written to the cells) to the distance from the other codewords. Any suitable distance metric, such as a Hamming metric or a Euclidean metric, can be used for this purpose.

The MSP adapts the storage configuration assigned to the page based on the distortion assessed at step 198. If the distortion level is acceptable, the MSP maintains the current storage configuration and the method terminates, at a termination step 202. If the distortion level is too high, the MSP increases the redundancy level of the ECC (i.e., adds redundancy bits), at a redundancy increasing step 204. Consequently, the density of the page is reduced. The MSP updates the storage configuration table accordingly, at a configuration updating step 206. The method then loops back to step 194 above. If the distortion level is too low (i.e., the storage density can be increased without unacceptably degrading performance), the MSP reduces the redundancy of the ECC, at a redundancy reduction step 208. The MSP updates the storage configuration table at step 206, and the method loops back to step 194 above. The iterations continue until the method converges to the appropriate number of redundancy bits that matches the current distortion level.

Checking the reliability of the programmed data and, if necessary, increasing the number of redundancy bits can be performed occasionally along the lifetime of device 24, such as using the method of Fig. 8 above. As distortion levels change over time, the storage configuration can be adapted accordingly to maintain the desired error performance.

In some embodiments, the storage configuration can be adapted using the method of fig. 9 without erasing the cells. For example, when the ECC comprises a systematic code, the redundancy bits can be stored separately from the uncoded bits. The number of redundancy bits can be increased or decreased without erasing or re-programming the cells that store the uncoded bits

An alternative technique for increasing the storage density of a group of cells (e.g., a page) without erasing cells or programming additional cells involves performing additional

programming operations on the already-programmed cells. Consider, for example, a memory cell array in which the cells can be programmed to levels  $0 \dots (2n-1)$ , wherein the stored charge level corresponding to level  $i$  is smaller than the charge level associated with level  $i+1$ .

In some embodiments,  $m$  memory cells are initially programmed to even-order levels, i.e., to levels  $0, 2, \dots, 2\lceil m/2 \rceil - 2$ . After estimating the distortion level in the cells, the MSP encodes additional bits of information using a binary encoder, which generates  $m$  encoded bits. The code-rate of the encoder depends on the estimated distortion. The MSP re-programs (without erasing) the cells for which the encoded bit is T to odd-order levels, i.e., to levels  $1, 3, \dots, 2\lceil m/2 \rceil - 1$ . For example, the new cell level can be written as  $y(m) = x(m) + b(m)$ , wherein  $y(m)$  denotes the new level of cell  $m$ ,  $x(m)$  denotes the previous level of the cell, and  $b(m)$  denotes the value of the  $m^{\text{th}}$  encoded bit (O' or T). This scheme can be regarded as an adaptive-rate Trellis Coded Modulation (TCM) that is adapted responsively to the estimated distortion.

### FINITE SETS OF STORAGE CONFIGURATIONS

In some embodiments, MSP 52 selects an appropriate storage configuration (e.g., number of levels and ECC scheme) for each page from a predetermined set of storage configurations. The use of a finite set of storage configurations reduces the complexity of the MSP, and in particular simplifies the storage configuration table in unit 76.

Fig. 10 is a flow chart that schematically illustrates a method for adaptively modifying the storage density of memory device 24 using a finite set of storage configurations, in accordance with an embodiment of the present invention. The method begins by predefining a set of storage configurations, at a definition step 210. Each storage configuration specifies a particular number of voltage levels and a particular ECC scheme. Thus, each storage configuration defines a certain storage density, i.e., a certain bit-per-cell value.

Note that the number of bits per cell need not necessarily be integer, since (1) the mapping of data bits to voltage levels can be performed jointly on groups of cells and not on individual cells, and (2) the number of voltage levels may not be a power of two. Exemplary methods for implementing fractional bit-per-cell storage density are described further below. The following table shows an exemplary set of storage configurations:

Storage configuration	Bits/cell	Levels/cell
0	N/A	N/A
1	1	2

2	~1.5	3
3	2	4
4	~2.5	6
5	3	8
6	~3.5	12
7	4	16
8	~4.25	20
9	~4.5	24

Storage configuration 0 identifies bad pages, which are not used. Configurations 1-9 provide nine different density levels, between 1 and 4.5 bits per cell. In alternative embodiments, any other suitable set of storage configurations can be used.

Typically, the storage configuration table in management unit 76 is initialized with a certain default storage configuration from the predefined set. When writing a particular page, MSP 52 encodes and maps the data using the ECC and number of levels specified by the storage configuration of the page, at a page programming step 214.

Following a programming operation or following a reading operation, the MSP estimates the performance margin of the page, at a margin estimation step 218. The MSP may estimate the performance margin by measuring the distortion in the voltages read from the page or by assessing the reliability of the data read from the page, as described above. The margin may be expressed in terms of Signal-to-Noise Ratio (SNR) or using any other suitable metric.

The MSP uses the performance margin to determine whether the storage configuration specified for the page matches the current capacity of the page, or whether the storage density should be increased or decreased. For example, the MSP may compare the estimated margin to a threshold.

If the estimated margin indicates that the storage configuration should be modified, the MSP selects a different configuration from the predefined set, at a configuration updating step 222. The MSP updates the configuration selection in the storage configuration table, so that subsequent programming operations on this page will use the updated storage configuration.

In some embodiments, MSP 52 compresses the data to be stored in a particular page before ECC encoding. Any compression method known in the art, which may be lossless or lossy, may be used for this purpose. The combination of compression and ECC encoding can produce different storage densities, i.e., different bit-per-cell values.

In some embodiments, the total number of bits stored in each page (denoted M) remains constant in the different storage configurations, and the number of net information bits (denoted K) is variable. For example, assume that a page comprises 4,224 memory cells, and that the cells are programmed using sixteen voltage levels (i.e., each cell stores four bits).

5 The total number of bits stored in a page is  $4,224 \times 4 = 16,896$ . With an ECC coding rate of  $7/8$ , the number of net information bits that can be stored in a page is 14,784. With a lower ECC coding rate of  $13/16$ , the number of net information bits that can be stored in a page is 13,728. The second storage configuration is, however, more robust due to the stronger ECC.

10 Alternatively, the number of voltage levels per memory cell can vary from one storage configuration to another, i.e., both K and M are varied. For example, the following table shows three storage configurations for use in an array of 4,224-cell pages:

Voltage levels per cell	Stored bits per cell	ECC rate	K	M	Net info bits per cell
16	4	7/8	14,784	16,896	3.5
12	3.5	13/14	13,728	14,784	3.25
12	3.5	6/7	12,672	14,784	3

15 In some embodiments, each page is divided into sectors having a constant number of net information bits. The number of sectors per page is varied. For example, the following table shows three storage configurations for use in an array of 4,224-cell pages. In the present example, - each sector contains 1,412 net information bits that are encoded using a rate  $2/3$  ECC to produce 2,118 stored bits.

Sectors per page	K	M	Levels per cell	Stored bits per cell
8	11,264	16,896	16	4
7	9,856	14,784	12	3.5
6	8,448	12,672	8	3

CAPACITY ESTIMATION USING PROGRAM/ERASE HISTORY

20 As noted above, the level of distortion in a particular memory cell may depend on the history of programming and erasure operations the cell has gone through. Thus, MSP typically assigns storage configurations having lower storage densities to older cells, and *vice versa*.

In some cases, distortion is caused by the leakage current of the cells. The effects of leakage current are described, for example, by Mielke et al., in "Recovery Effects in the Distributed Cycling of Flash Memories," IEEE Annual International Reliability Physics Symposium, San Jose, California, March, 2006, pages 29-35, which is incorporated herein by reference. According to this paper, dielectric damage caused by programming and erasure cycling partially recovers during the delays between cycles. The paper characterizes the effect of these delays on the charge-detrapping data-retention mechanism. Thus, the level of distortion that is associated with leakage current depends on the length of time that that elapsed from the occurrence of previous programming and erasure operations. Other types of distortion may also be influenced by charge trapping, and thus by recent programming and erasure operations. For example, the level of disturb errors may increase in memory cells having a high number of charge traps.

Fig. 11 is a flow chart that schematically illustrates a method for adaptively modifying the storage density of memory device 24 based on the history of programming and erasure operations, in accordance with an embodiment of the present invention. The method begins with management unit 76 of MSP 52 recording the times in which programming and erasure operations were performed on the different memory cells 32 of array 28. The MSP may maintain a running counter that tracks the current time in order to time-tag the programming and erasure operations. Alternatively, in some cases the time tags are provided by memory controller 22.

Since programming operations are typically performed page-by-page and erase operations are performed on entire erasure blocks, the MSP can record the programming and erasure times per page. The information can be stored as part of the storage configuration of each page in the storage configuration table. Erasure times can be stored either per page or per erasure block.

MSP 52 estimates the leakage current in a particular page, at a leakage estimation step 234. The MSP can use various estimators and criteria to estimate the leakage current based on the recorded programming and erasure times. For example, the leakage current can be estimated by the function

$$I_{Leakage} \approx a_0 \cdot CurrentTime + \sum_{n=J}^K a_i ProgramTime(n)$$

wherein  $ProgramTime(n)$  denotes the time of the  $n$ 'th programming operation performed on the page, as recorded at step 230 above.  $CurrentTime$  denotes the time in which the estimate is

calculated.  $\alpha_i, \dots, K$  denote weighting coefficients, which typically assign a higher weight to more recent programming operations. When  $\alpha_i=0$  for  $i>1$ , the estimate is based only on the last programming operation.

In alternative embodiments, the leakage current can be estimated based on the number of erase operations performed in a certain time interval (e.g., in the last day, hour or minute). Further alternatively, the leakage current can be estimated iteratively, such as using the function

$$I_{Leakage}(n+1) = (1 - \delta) \cdot I_{Leakage}(n) + \delta \cdot [ProgramTime(n+1) - ProgramTime(n)]$$

wherein  $0 < \delta < 1$  denotes a weighting coefficient or "forgetting factor." The iterative estimation enables MSP 52 to record only the times of last two programming operations and the value of the previous estimate, rather than storing a list of multiple time tags.

Alternatively, the MSP can use any other estimate of method for estimating the leakage current based on the recorded history of programming and erase operations.

In some embodiments, the MSP estimates the cell distortion based on other factors, at a complementary distortion estimation step 238. The distortion estimation can be based on measuring cross-coupling ratios from neighboring cells or on any other suitable process.

The MSP combines the leakage current estimation performed at step 234 above with the distortion estimation performed at step 238 above and estimates the capacity of the page, at a capacity estimation step 242. For example, the capacity can be estimated using the function

$$Capacity = Q_0 + \beta \cdot I_{leakage} + \gamma \cdot DistortionEstimate$$

wherein  $Q_0$  denoted a baseline capacity value, and  $\beta$  and  $\gamma$  denote weighting coefficients that assign the desired relative weights to the leakage current and distortion estimates, in some embodiments, step 238 may be omitted and the MSP may estimate the page capacity based only on the leakage current estimate, such as by setting  $\gamma=0$ .

MSP 52 updates the storage configuration table with the updated page capacity values. The MSP may determine the page density and the appropriate storage configuration (e.g., ECC and number of voltage levels) for this page using any suitable method, such as the methods described hereinabove.

DENSITY ALLOCATION WITH FRACTIONAL BIT-PER-CELL VALUES

In some embodiments, MSP 52 can define storage configurations having a non-integer number of bits per cell. For example, in some cases it is possible to increase the storage density and approach the achievable capacity of a cell by using a number of voltage levels that is not a power of two, MSP 52 can use various methods for mapping data bits to voltage levels, when the number of bits per cell is not integer.

Let  $K$  denote the number of voltage levels used for storing data in a particular page.  $K$  is not necessarily a power of two, thus the storage density is not necessarily integer. The maximum number of bits that can be represented by the  $K$  levels is given by  $n_{opt} = \log_2(K)$ , which is generally a non-integer value. In order to reaching this optimal storage density, an infinite number of data bits would need to be stored in an infinite number of memory cells, and the process would involve infinite delay. In practical scenarios, however, the deviation from the optimal storage density can be traded for processing delay (i.e., the number of cells used for storing the data).

Let  $n_c$  denote the number of memory cells used to store the data, and  $n_{tb}$  denote the total number of bits stored in the  $n_c$  cells.  $n_c$  and  $n_{tb}$  are both integer values. The storage density is thus given by  $n_{bpc} = n_{tb}/n_c$ .  $n_{tb}$  and  $n_c$  should be chosen so that the deviation from optimal density, given by  $ngap = n_{opt} - n_{bpc}$  is non-negative and minimized. Increasing  $n_c$  enables  $n_{bpc}$  to approach  $n_{opt}$ , but increases the mapping complexity and the processing delay.

Having selected the desired values of  $n_c$  and  $n_{tb}$ , MSP 52 maps each group of  $n_{tb}$  data bits to  $n_c$  voltages, which are in turn written into  $n_c$  memory cells. Each of the  $n_{tb}$  voltages can take one of  $K$  possible levels. The  $n_{tb}$  data bits are denoted  $b_1, \dots, b_{n_{tb}} \in \{0, 1\}$ .

In some embodiments, MSP 52 maps the  $n_{tb}$  data bits to the  $n_c$  voltages by representing the  $n_{tb}$ -bit binary number expressed by the bits in  $K$  basis. In other words, MSP 52 determines  $n_c$  coefficients  $k_1, \dots, k_{n_c} \in \{0, \dots, K-1\}$ , such that

$$b_{n_{tb}} \cdot 2^{n_{tb}-1} + b_{n_{tb}-1} \cdot 2^{n_{tb}-2} + \dots + b_2 \cdot 2 + b_1 = k_{n_c} \cdot K^{n_c-1} + k_{n_c-1} \cdot K^{n_c-2} + \dots + k_2 \cdot K + k_1$$

MSP 52 then programs the  $n_c$  memory cells with the values  $k_1, \dots, k_{n_c}$ .

For example, let  $X=11$ . The optimal storage density in this case is  $n_{opt} \approx 4 \log_2(1.1) \approx 3.4594$ . Assuming data is written in groups of  $n_c=3$  cells, the deviation from optimal density is  $n_{gap} \approx 0.1261$  and  $n_{tb}=10$ . Thus, sequences of ten data bits are mapped to groups of three memory cells, each using eleven voltage levels. For example, the bit sequence "1101100100" is mapped to  $k_1=7$ ,  $k_2=1$  and  $k_3=10$ .

#### STORING AND ERASING DATA IN A VARIABLE-CAPACITY MEMORY DEVICE

Fig. 12 is a flow chart that schematically illustrates a method for storing data in a variable-capacity memory device, in accordance with an embodiment of the present invention. The method begins with MSP 52 accepting a data item for storage from memory controller 22, in an input step 250.

In some embodiments, the memory controller specifies a certain reliability requirement for storing the data item. The memory controller may use any suitable method or form for specifying the requested storage reliability, such as by explicitly specifying a maximum tolerated error probability or by selecting a reliability level from a set of predefined levels. Often but not necessarily, the reliability requirement is associated with the error-tolerance of the application that generated the data item. For example, a file containing audio or video media may be stored at a relatively low reliability level. Files containing program code and operating system data files, on the other hand, will usually be specified with a high reliability level.

Additionally or alternatively, the memory controller may specify a certain retention requirement for storing the data item. The retention requirement indicates the time period over which the data item is expected to be stored. For example, temporary operating system files can be specified with a low retention period.

MSP 52 selects a set of one or more pages in array 28 in which to store the data item, at a page selection step 254. Typically, management unit 76 tracks the status of each memory page. Each page may have a status of "erased" (i.e., available for programming), "contains information," "faulty" or "ready for erasing." As noted above, unit 76 also holds the storage configuration (e.g., the number of voltage levels and ECC) used in each page, which defines the amount of data that can be stored in the page.

Unit 76 selects a set of one or more "erased" pages whose aggregate size is sufficient for storing the data item. In some embodiments, unit 76 scans through the "erased" pages in

sequential order and selects pages until their aggregate size is sufficient for storing the data item. Alternatively, unit 76 may apply various policies or criteria for selecting pages.

For example, unit 76 may select pages that have gone through a relatively low number of programming and erasure cycles, in order to achieve better wear leveling across array 28.

5 As another example, unit 76 may search for a set of pages whose aggregate size provides the closest match to the size of the data item (i.e., exceeds the size of the data item by as little as possible), in order to minimize the number of unused memory cells. As yet another example, unit 76 may give preference to high-capacity pages, in order to split the data item among a minimal number of pages and reduce programming time, wearing and management overhead.

10 Alternatively, unit 76 may use any other suitable criterion or combination of criteria when selecting pages for storage.

In some embodiments, unit 76 modifies the storage configuration (and thus the storage density) of the pages based on the reliability and/or retention requirements specified for the data item. For example, when the data item has a relatively low reliability and/or retention

15 requirement, unit 76 may store the item in cells having higher distortion and/or at a higher density with respect to the distortion level of the cells.

MSP 52 stores the data item in the selected pages, at a storage step 258. Unit 76 sets the status of the pages used for storage to "contains information," and updates the storage configuration of the pages if necessary. Unit 76 also records the identities of the pages used

20 for storing the data item, so that these pages can be addressed when system 20 is later requested to retrieve the data item.

In some embodiments, system 20 can be requested by memory controller 22 to report the available memory size. In these embodiments, unit 76 sums the capacities of the pages marked as "erased" and reports the result as the available memory size.

25 Fig. 13 is a flow chart that schematically illustrates a method for erasing data from a variable-capacity memory device, in accordance with an embodiment of the present invention. The method begins with MSP 52 accepting a request from memory controller 22 to delete a certain data item from memory device 24, at a delete requesting step 270. Management unit 76 in MSP 52 identifies the pages used for storing the data item and marks them as "ready for

30 erasing," at a deletion step 274.

The pages that store the data item may be located in one or more erasure blocks. Before these blocks can be erased, any pages in these blocks that are in use (i.e., have "containing information" status) should be relocated to other erasure blocks.

Unit 76 checks whether the erasure blocks that contain the pages identified at step 274 above contain in-use pages, at an in-use checking step 278. If the erasure blocks do not contain any in-use pages, unit 76 erases the blocks and updates the status of the pages in these blocks to "erased," at an erasure step 282. If, on the other hand, unit 76 detects that some  
5 pages in the blocks intended for erasure are in use, it relocates the data stored in these pages to other pages, outside the blocks intended for erasure. (In some cases, unit 76 may decide not to erase the block, such as when an erasure block contains a single page belonging to the erased data item and all other pages are already erased.)

Unit 76 selects a set of one or more pages, which are located outside the blocks  
10 intended for erasure and whose aggregate size is sufficient for storing the data in the in-use pages, at a relocation selection step 286. Unit 76 can select the pages to which data is relocated in accordance with any suitable policy or criteria. For example, the pages can be selected out of the available pages in sequential order. Alternatively, unit 76 may select the pages so as to improve wear leveling or to reduce the number of unused memory cells, as  
15 described in page selection step 254 of the method of Fig. 12 above. Further alternatively, unit 76 can use any other suitable criterion.

Unit 76 copies the data stored in the in-use pages to the pages selected at step 286 above, at a relocation step 290. Note that the relocation operation does not necessarily preserve the number of pages or the partitioning of data into pages. Since the pages in device  
20 24 have different capacities, the relocated data can be copied to a different number of pages and/or partitioned among the pages in a different manner with respect to the original storage. Once the data is relocated, the blocks identified at step 274 above are erased, at erasure step 282.

#### DATA STORAGE MANAGEMENT IN VARIABLE-CAPACITY MEMORY DEVICES

As noted above, MSP 52 interacts with memory controller 22, i.e., accepts data for  
25 storage from the memory controller and outputs data that is stored in memory to the memory controller when requested. System 20 can interact with the memory controller or generally with the host system in several ways, which differ from one another in the level of management functionality carried out by the MSP.

At one extreme, system 20 interacts with a legacy memory controller, which is  
30 designed to control fixed-capacity memory devices. The memory controller in this case maintains a fixed-capacity page table, indicating which fixed-capacity pages are used for storing each data item. The fixed-capacity page table may also hold the status of each fixed-

capacity page or block, and indicate whether a certain erasure block is considered to be a bad block.

In these embodiments, MSP 52 mediates between the adaptive capacity characteristics of system 20 and the fixed capacity characteristics of memory controller 22, by maintaining two parallel indexing schemes. Towards the memory controller, the MSP communicates using the indexing scheme of the fixed-capacity pages, so that system 20 appears to be a fixed-capacity device to the memory controller. Internally in system 20, MSP 52 stores the data items efficiently in variable-capacity pages. The MSP maintains a separate indexing scheme of variable-capacity pages, which indicates which variable-capacity pages are used to store each data item. The MSP is able to translate between the two indexing schemes when writing and reading data.

In some cases, memory controller 22 also carries out functions such as ECC encoding and decoding, detection of errors in the data read from memory and/or bad block management. Using the techniques described herein, known memory controllers can be used to control memory device 24 without having to modify their data structures or protocols.

At the other extreme, the functionality of memory controller 22 is substantially carried out by MSP 52 itself. In these embodiments, the MSP accepts data items of arbitrary sizes from the host system (e.g., from a computer CPU), and stores the data items in the variable-capacity pages of array 28. Since the memory management functions are carried out by the MSP, there is no need to emulate fixed-capacity indexing, and only one indexing scheme is maintained by the MSP.

Further alternatively, MSP 52 may support any other suitable division of memory management functions between system 20 and the host system. A single MSP device may support different interfacing and indexing schemes in order to interact with different types of memory controllers or hosts.

Figs. 14-17 are diagrams that schematically illustrate methods for storing fixed-size data blocks in a variable-capacity memory device, in accordance with embodiments of the present invention.

Fig. 14 illustrates a method for interacting with a memory controller, which is designed to store data in memory devices having erasure blocks having a fixed number of cells, and which carries out a Bad Block Management (BBM) function. In some known BBM methods, a certain erasure block is classified as a bad block even if only a single page in the block is faulty. Such methods are extremely inefficient, since most of the memory cells in a

bad block are still usable. Using the method of Fig. 14, on the other hand, system 20 stores data efficiently in the usable memory cells of blocks classified as bad.

In a typical writing operation, memory controller 22 accepts a data item 300 for storage in a memory array 304. Array 304 comprises multiple memory blocks 308. Memory controller 22 divides the data item into one or more fixed-size blocks, which are referred to as controller blocks 312. Each controller block 312 is destined for storage in a certain available memory block 308. Memory controller 22 provides controller blocks 312 to MSP 52 for storage.

Some of blocks 308 may have reduced capacity (i.e., capacity that is smaller than the size of controller blocks 312) due to distortion effects, manufacturing faults or for any other reason. Using known BBM methods, such blocks would have been classified as bad blocks and remained unused, even though the vast majority of their memory cells may still be usable.

MSP 52, on the other hand, regards all memory blocks 308 as variable-capacity blocks and does not classify blocks as good or bad. The MSP extracts the data from the controller blocks and re-partitions the data into multiple fragments 316, which may be of the same size or of different sizes. The MSP identifies multiple memory blocks 308 whose capacities match the sizes of fragments 316, and stores each fragment in a respective memory block. In the present example, the MSP divides the data item into three fragments and stores the fragments in the memory blocks denoted #3, #6 and #10.

MSP 52 may use various methods and criteria to determine the blocks in which to store each data item. For example, the MSP may search for a pair of memory blocks having reduced capacity, whose aggregated capacity is sufficient to store the data of a single controller block. The MSP then maps the controller block to the pair of memory blocks. As another example, the MSP may search the available memory blocks and determine a set of two or more blocks 308, whose aggregated capacity best matches the size of data item 300. Using this method, the amount of wasted memory is minimized.

Alternatively, the MSP may search for a set having a minimum number of blocks in which the data item can be stored, in order to reduce management overhead. Further alternatively, the MSP may select blocks that have gone through a minimum number of recent programming and erasure cycles, in order to improve wear leveling. Any other suitable method or criterion can also be used. In some embodiments, a single memory block 308 can be used to store fragments belonging to more than one data item.

The partitioning of data item 300 into fragments 316 can be performed either independently of or in connection with the selection of memory blocks 308. For example, the

MSP may first fragment the data item into fixed-size fragments, regardless of the number and identity of the memory blocks in which the fragments will ultimately be stored in. Alternatively, the MSP may first identify the memory blocks in which the data item is to be stored, and then fragment the data item according to the number and size of these memory  
5 blocks.

In some embodiments, MSP 52 emulates storage in fixed-capacity blocks toward the memory controller, such that the fragmentation performed by the MSP is transparent to the memory controller and its BBM function. For example, the MSP may maintain two indexing tables, i.e., an indexing table of fixed-size controller blocks and a corresponding indexing  
10 table of variable-capacity memory blocks. Whenever a data item is sent by the memory controller for storage in a certain set of controller blocks, the MSP indicates which variable-capacity memory blocks are used for storing the item. When the data item is requested by the memory controller, the MSP retrieves the item from the variable-capacity memory blocks in which it was stored, and sends the item to the memory controller, as if it was stored in the  
15 fixed-size controller blocks known to the controller.

As can be appreciated, the method of Fig. 14 increases the achievable capacity of the memory device considerably with respect to known BBM methods, since it enables data to be stored in memory blocks that would have otherwise been classified as bad blocks.

Although the description of Fig. 14 referred to a storage granularity of entire erasure  
20 blocks, the method can also be implemented using finer granularity, i.e., page granularity. In other words, the MSP can determine a suitable set of variable-capacity memory pages in which the data item can be stored. Fragmenting and storing data items on a page-by-page basis enables higher storage density at the expense of relatively high management overhead.

In an exemplary implementation, when the nominal page size used by the memory  
25 controller is  $P$  bits, the MSP classifies the memory pages in the memory device into reduced-capacity memory pages whose capacity is smaller than  $P$ , and nominal-capacity memory pages whose capacity is greater than or equal to  $P$ . The MSP can then store each controller page having  $P$  bits either in a nominal-capacity memory page, or in a pair of reduced-capacity memory pages. The MSP typically maintains a table that maps each controller page to the  
30 corresponding nominal-capacity memory page or pair of reduced-capacity memory pages. Thus, memory pages whose capacity falls below the nominal size of the controller page can still be used for storing data.

Fig. 15 illustrates another method for efficiently storing fixed-size controller pages having  $P$  bits in variable-capacity memory pages, in accordance with an embodiment of the

present invention. A memory array 320 comprises multiple memory pages 324. A memory page is defined as a group of memory cells that are programmed and read simultaneously.

Although the memory pages of array 320 typically have the same number of cells, their capacity (i.e. the amount of information that they can reliably store) may vary from page to page and over time. Generally, the capacity of each memory page 324 may be smaller than, equal to or greater than  $P$ . MSP 52 allocates one or more of the memory pages to function as excess capacity pages 328, which are used for storing excess bits of other memory pages whose capacity is smaller than  $P$ . In the present example, array 320 comprises thirty-two memory pages, of which one page serves as an excess capacity page.

When the MSP accepts a controller page having  $P$  bits for storage, it assesses the capacity of the next-available memory page 324, denoted  $C_n$ . If  $C_n > P$  (i.e., the capacity of the next-available memory page is sufficient for storing the  $P$  bits of the controller page), the MSP stores the  $P$  bits in the memory page. If the capacity of the next-available memory page is not sufficient for storing the  $P$  bits of the controller page, i.e., when  $C_n \leq P$ , the MSP stores  $C_n$  bits out of the  $P$  bits of the controller page in the memory page, and stores the remaining  $P - C_n$  bits in unused memory cells of excess capacity pages 328.

Note that writing or reading a page using the method described above may involve two page write or read operations, in some embodiments, the MSP may cache some or all of the excess capacity pages, or portions of these pages, in a RAM or other cache memory. By caching excess capacity pages, the MSP can reduce the number of dual-read and dual-write operations.

When reading data out of the memory device, the MSP can cache the last read excess capacity bits. When retrieving a memory page, the MSP checks whether the excess bits of the page are present in the cached data. If the desired excess bits are cached, they can be appended to the retrieved memory page without having to physically read the excess capacity page. When writing data into the memory device, the MSP can cache the excess capacity bits and physically store them in the memory device less frequently.

In order to prevent the cached data from being lost in case of voltage failure, the MSP may comprise circuitry that identifies voltage irregularities and rapidly stores the cached data in the memory device when necessary. The memory controller may also instruct the MSP to store the cached data in the memory device using an End Of File (EOF) attribute or other suitable command. The memory controller will typically use this command before switching off power and/or at the end of writing a data file. Further alternatively, the MSP can also store the cached data in the memory after a certain time out.

Fig. 16 illustrates an alternative method for efficiently storing fixed-size controller pages in variable-capacity memory pages, in accordance with an embodiment of the present invention. In the method of fig. 16, MSP 52 stores controller pages (denoted CTRL PAGE #1...#5), each having  $P$  bits, in variable-capacity memory pages 334 of a memory array 330.

5 The capacity of each memory page 334 can be smaller than, equal to or greater than  $P$ .

MSP 52 stores the controller pages one following the other, regardless of the boundaries of memory pages 334. Thus, a certain memory page may contain data belonging to one or more controller pages, and the data of a certain controller page may be stored in one or more memory pages. The MSP records the location of each controller page in the array, such  
10 as by storing the start and end addresses of each controller page. This storage method utilizes the entire capacity of the memory pages, at the expense of some management overhead.

The storage method is typically transparent to the memory controller. When the memory controller requests a certain controller page, the MSP reads the requested controller page from the appropriate memory location and sends the data to the controller.

15 In some cases, such as when storing large data files or digital images, the memory controller sends sequences of successive controller pages for storage. In such cases, the MSP can reduce the number of physical page programming operations by buffering the data accepted from the memory controller. The MSP physically writes the data to the memory array, for example, when the buffer is full, when explicitly instructed by the memory  
20 controller (e.g., responsively to an EOF command) or when the memory controller sends a controller page that is not part of the sequence. An exemplary command that can be used for this purpose is the "Page Cache Program" command, which is defined on page 61 of the ONFI specification, cited above. The MSP can also reduce the number of page-read operations required for reading a large amount of data that is stored sequentially in memory by buffering  
25 the data read from the memory.

Fig. 17 illustrates yet another method for efficiently storing fixed-size controller pages in variable-capacity memory pages, in accordance with another embodiment of the present invention. In the method of Fig. 17, the MSP writes the controller pages in two stages, in order to increase programming speed.

30 The memory pages of a memory array 340 are divided into a temporary storage area 344 and a permanent storage area 348. When the MSP accepts the controller pages from the memory controller, it initially writes them into separate memory pages in area 344, such that each controller page is written into a single memory page.

At a later point in time, typically when communication with the memory controller is idle, the MSP re-programs the data into permanent storage area 348. The data in area 348 is stored in the closely-packed, capacity-efficient method of Fig. 16 above. Pages in area 344 that have been re-programmed are freed and can be reused for temporary storage of subsequent controller pages.

#### WEAR LEVELING CONSIDERATIONS

The level of distortion present in a particular page often depends on the number of programming and erasure cycles the page has gone through. Typically, the performance and capacity of a page deteriorates as it goes through an increasing number of programming and erasure cycles. Thus, it is often advantageous to distribute the programming of memory pages, so that performance degradation is distributed evenly across the memory cell array. This feature is referred to as wear leveling, and is described, for example, in "Wear Leveling in Single Level Cell NAND Flash Memories," Application note AN-1822 by STMicroelectronics (Geneva, Switzerland), February, 2007, which is incorporated herein by reference.

When MSP 52 selects a page or group of pages in which to write data, the MSP may select a page or group of pages that have gone through a relatively small number of programming and erasure cycles. Such wear leveling considerations can be used in any of the methods described above.

In some cases, the MSP can select a page based on a metric that combines wear leveling with distortion estimation, such as

$$Metric = \gamma \cdot NumberWriteEraseCycles + \delta \cdot DistortionEstimation$$

wherein  $\gamma$  and  $\delta$  denote weighting coefficients. *DistortionEstimation* denotes the estimated distortion level. For example, *DistortionEstimation* can be derived by counting the memory cells whose voltage has dropped significantly. Another factor that can be added to the definition of *Metric* above is the time that passed since the last erase operation of the block (similarly to the equation that defines *ILeakage* in the description of Fig. 11 above).

Although the embodiments described herein mainly refer to adapting the storage density of multi-level memory cells by modifying the number of voltage levels and ECC scheme, the methods and systems described herein can also be used in single-level cells, by varying only the ECC.

Although the embodiments described herein mainly address retrieving data from solid-state memory devices, the principles of the present invention can also be used for storing and

retrieving data in Hard Disk Drives (HDD) and other data storage media and devices. In some storage devices such as HDD and digital cassettes, the storage media is not pre-divided into discrete cells before the data is stored. Rather, the data is stored in regions of the continuous media, and the locations of these regions within the continuous media are defined as part of the storage process itself. In the context of the present patent application and in the claims, such regions, which are defined during the storage process in media that is *a-priori* continuous, are also considered as "analog memory cells".

It will thus be appreciated that the embodiments described above are cited by way of example, and that the present invention is not limited to what has been particularly shown and described hereinabove. Rather, the scope of the present invention includes both combinations and sub-combinations of the various features described hereinabove, as well as variations and modifications thereof which would occur to persons skilled in the art upon reading the foregoing description and which are not disclosed in the prior art.

## CLAIMS

1. A method for data storage in a memory that includes a plurality of analog memory cells, the method comprising:
  - estimating respective achievable storage capacities of the analog memory cells;
  - 5 assigning the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities;
  - storing the data in the memory cells in accordance with the respective assigned storage configurations; and
  - re-estimating the respective achievable storage capacities of the analog memory cells
- 10 after the memory has been installed in a host system and used for storing the data in the host system, and modifying the storage configurations responsively to the re-estimated achievable capacities.
2. The method according to claim 1, wherein storing the data comprises encoding the data using an Error Correcting Code (ECC), converting the encoded data to analog values
- 15 selected from a set of nominal analog values and writing the analog values into the respective memory cells, and wherein each storage configuration specifies a respective ECC code rate and a size of a set of nominal analog values used for storing the data.
3. The method according to claim 1, wherein, for each analog memory cell, storing the data comprises converting the data to analog values selected from a set of nominal analog
- 20 values and writing the analog values into the memory cell, and wherein each storage configuration specifies the nominal analog values used for storing the data in the analog memory cell.
4. The method according to claim 1, wherein storing the data comprises converting the data to analog values and writing the analog values into the respective memory cells, and
- 25 wherein estimating the achievable storage capacities comprises estimating respective distortion levels affecting the analog values written into the memory cells and determining the achievable storage capacities responsively to the distortion levels.
5. The method according to any of claims 1-4, wherein estimating the distortion levels comprises reading the analog values from the memory cells and estimating the distortion
- 30 levels based on the analog values read from the memory cells.

6. The method according to claim 5, wherein estimating the distortion levels comprises reconstructing the data from the analog values read from the memory cells and calculating the distortion levels based on the reconstructed data.
7. The method according to claim 5, wherein calculating the distortion levels comprises  
5 summing scalar functions of the analog values read from the memory cells.
8. The method according to claim 7, wherein summing the scalar functions comprises summing squares of differences between the analog values read from the memory cells and respective expected values of the data stored in the memory cells responsively to the data.
9. The method according to claim 6, wherein calculating the distortion levels comprises  
10 counting a number of memory cells in which a difference between the analog values read from the memory cells and respective expected values of the data stored in the memory cells responsively to the data exceeds a predetermined value.
10. The method according to claim 4, wherein storing the data comprises writing the analog values into the memory cells- and verifying the written analog values using an iterative  
15 Program and Verify (P&V) process, and wherein estimating the distortion levels comprises calculating the distortion levels based on the analog values verified by the P&V process.
11. The method according to claim 4, wherein storing the data comprises writing the analog values into the memory cells and verifying the written analog values using an iterative  
20 Program and Verify (P&V) process, which iteratively increments the analog values by a programming step, and wherein the storage configurations define a size of the programming step used by the iterative P&V process.
12. The method according to any of claims 1-4, wherein storing the data comprises encoding the data to be stored in a group of the memory cells using an Error Correcting Code (ECC) and converting the encoded data to analog values for storage in the respective analog  
25 memory cells of the group, wherein estimating the achievable storage capacities comprises reading the analog values from the memory cells in the group and evaluating a distance metric between the read analog values and a valid codeword of the ECC that is nearest to the analog values, and wherein modifying the storage configurations comprises adapting the storage configuration of the group of the memory cells responsively to the evaluated distance metric.
- 30 13. The method according to any of claims 1-4, wherein re-estimating the achievable storage capacities comprises reading the data stored in the memory cells, and wherein

modifying the storage configurations comprises adapting the storage configurations responsively to detecting errors in the read data.

14. The method according to claim 13, wherein storing the data comprises encoding the data using an Error Correcting Code (ECC), and wherein adapting the storage configurations  
5 comprises modifying a code rate of the ECC.

15. The method according to any of claims 1-4, wherein estimating the achievable capacities of the memory cells comprises tracking previous programming and erasure operations applied to the respective memory cells and estimating the achievable capacities responsively to the tracked previous programming and erasure operations.

10 16. The method according to claim 15, wherein estimating the achievable capacities comprises calculating the achievable capacities responsively to lengths of time periods that elapsed since the previous programming and erasure operations.

17. The method according to any of claims 1-4, wherein storing the data comprises converting an initial portion of the data to analog values and writing the analog values into the  
15 respective memory cells, and subsequently storing an additional portion of the data in at least some of the memory cells by increasing the analog values written to the memory cells without erasing the cells.

18. The method according to any of claims 1-4, wherein storing the data comprises encoding the data using an Error Correcting Code (ECC) that adds redundancy bits to the data  
20 and storing the redundancy bits in some of the analog memory cells, and wherein modifying the storage configurations comprises modifying a number of the redundancy bits added by the ECC without erasing the cells.

19. The method according to claim 18, wherein storing the redundancy bits comprises storing the redundancy bits separately from the data.

25 20. The method according to any of claims 1-4, wherein assigning the storage configurations comprises predefining a set of possible storage configurations, and wherein modifying the storage configurations comprises selecting updated storage configurations from the predefined set.

21. The method according to any of claims 1-4, wherein the quantity of the data defined by  
30 at least one of the storage configurations specifies a non-integer number of bits per cell.

22. The method according to any of claims 1-4, wherein storing the data comprises compressing the data before writing the data to the memory cells, and wherein the storage configurations define respective compression ratios by which the data is to be compressed.

23. The method according to any of claims 1-4, wherein assigning the storage configurations comprises defining the storage configurations for respective groups of the memory cells.

24. The method according to any of claims 1-4, wherein storing the data in the memory cells comprises accepting a data item for storage in the memory, selecting a subset of the memory cells in which to store the data item responsively to the assigned storage configurations, and storing the data item in the selected subset of the memory cells.

25. The method according to any of claims 1-4, wherein re-estimating the achievable storage capacities is performed during idle time periods in which the data is not stored and read.

26. The method according to any of claims 1-4, wherein assigning and modifying the storage configurations comprise storing the storage configurations in a configuration table.

27. The method according to claim 26, and comprising storing initial values of the storage configurations in the configuration table before the memory has been installed in the host system.

28. The method according to any of claims 1-4, wherein assigning and modifying the storage configurations comprise maintaining a predetermined margin between the quantities of the data stored in the memory cells and the respective estimated achievable capacities.

29. The method according to any of claims 1-4, wherein at least some of the distortion is caused by leakage current in the analog memory cells, and wherein estimating the distortion comprises estimating the leakage current.

30. A method for data storage in a memory that includes a plurality of analog memory cells, the method comprising:

estimating distortion levels in the respective analog memory cells;

estimating respective achievable storage capacities of the analog memory cells based on the estimated distortion levels;

assigning the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities; and

storing the data in the memory cells in accordance with the respective assigned storage configurations.

31. A method for data storage in a memory that includes a plurality of analog memory cells, the method comprising:

5 tracking respective achievable storage capacities of the analog memory cells while the memory is in use in a host system;

accepting data for storage in the memory;

selecting a subset of the memory cells for storing the data based on the tracked achievable capacities; and

10 storing the data in the memory cells of the subset.

32. The method according to claim 31, wherein selecting the subset comprises selecting the memory cells whose aggregate achievable capacity best matches a size of the accepted data.

33. The method according to claim 31, wherein accepting the data comprises accepting a  
15 requested reliability level for storing the data, and wherein selecting the subset comprises selecting the memory cells responsively to the requested reliability level.

34. The method according to claim 31, wherein accepting the data comprises accepting a requested retention period for storing the data, and wherein selecting the subset comprises selecting the memory cells responsively to the requested retention period.

20 35. The method according to claim 31, wherein selecting the subset comprises selecting the memory cells having low distortion levels with respect to other memory cells.

36. The method according to claim 31, wherein tracking the achievable storage capacities comprises tracking previous programming and erasure operations applied to the memory cells, and wherein selecting the subset comprises selecting the memory cells responsively to the  
25 previous programming and erasure operations.

37. The method according to claim 36, wherein selecting the subset comprises distributing a number of programming and erasure operations evenly over the memory cells in the memory by selecting the memory cells having a smaller number of the previous programming and erasure operations with respect to other memory cells.

30 38. The method according to claim 37, wherein distributing the number of programming and erasure operations comprises selecting the memory cells having the smaller number of the

previous programming and erasure operations performed during a predetermined recent time period with respect to the other memory cells.

39. The method according to claim 31, wherein tracking the achievable storage capacities comprises calculating and reporting to the host system a size of an available memory space by summing the achievable capacities of the memory cells that are available for storing the data.

40. The method according to claim 31, wherein the memory is divided into multiple erasure blocks, each erasure block comprising a group of the memory cells that are erased in a single erasure operation, and comprising erasing a data item from the memory by

identifying one or more of the erasure blocks in which the data item is stored;

10 when the erasure blocks in which the data item is stored contain stored data of another data item, identifying alternative memory cells outside the erasure blocks in which the data item is stored based on the tracked achievable capacities of the memory cells and on a size of the other data item, and copying the data of the other data item to the alternative memory cells; and

15 erasing the erasure blocks in which the data item is stored.

41. The method according to claim 31, wherein accepting the data for storage comprises accepting the data in fixed-capacity blocks from the host system and storing the data in variable-capacity groups of the memory cells based on the tracked achievable capacities of the variable-size groups.

20 42. The method according to claim 41, wherein storing the data in the variable-capacity groups comprises identifying two or more variable-capacity groups having capacities that are smaller than a capacity of the fixed-capacity block, and storing the data accepted in one or more of the fixed-capacity blocks in the identified two or more variable-capacity groups.

25 43. The method according to claim 41, wherein storing the data in the variable-capacity groups comprises allocating one or more of the variable-capacity groups to serve as an excess memory storage area, storing the data accepted in the fixed-capacity blocks in respective variable-capacity groups, and, when the achievable capacity of a variable-capacity group is smaller than a capacity of the fixed-capacity blocks, storing some of the data accepted in the respective fixed-capacity block in the allocated excess memory storage area.

30 44. The method according to claim 41, wherein storing the data in the variable-capacity groups comprises storing the data accepted in the fixed-capacity blocks sequentially in the memory cells irrespective of boundaries between the variable-capacity groups.

45. The method according to claim 44, wherein storing the data comprises initially storing the data accepted in each fixed-capacity block in a respective variable-capacity group, and subsequently re-programming the initially-stored data sequentially, irrespective of the boundaries between the variable-capacity groups.

5 46. The method according to claim 41, and comprising retrieving the data from the variable-size groups, reconstructing the fixed-size blocks and outputting the data using the fixed-size blocks to the host system.

47. The method according to claim 31, and comprising reducing a number of memory access operations applied to the memory by caching at least some of the data.

10 48. The method according to claim 47, wherein caching the at least some of the data comprises transferring the cached data to the memory upon detecting an immediate storage event.

49. The method according to claim 48, wherein the immediate storage event comprises at least one event selected from a group of events consisting of an approaching power failure, a time out and an End OfFile (EOF) command accepted from the host system.

15 50. Data storage apparatus, comprising:  
an interface, which is arranged to communicate with a memory that includes a plurality of analog memory cells; and  
a memory signal processor (MSP), which is arranged to estimate respective achievable  
20 storage capacities of the memory cells, to assign the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities, to store the data in the memory cells in accordance with the respective assigned storage configurations, to re-estimate the respective achievable storage capacities of the analog memory cells after the memory has been installed in a host system and  
25 used for storing the data in the host system, and to modify the storage configurations responsively to the re-estimated achievable capacities.

51. The apparatus according to claim 50, wherein the MSP is arranged to encode the data using an Error Correcting Code (ECC), to convert the encoded data to analog values selected from a set of nominal analog values and to write the analog values into the respective memory  
30 cells, and wherein each storage configuration specifies a respective ECC code rate and a size of a set of nominal analog values used for storing the data.

52. The apparatus according to claim 50, wherein, for each analog memory cell, the MSP is arranged to store the data by converting the data to analog values selected from a set of nominal analog values and writing the analog values into the memory cell, and wherein each storage configuration specifies the nominal analog values used for storing the data in the analog memory cell.
53. The apparatus according to claim 50, wherein the MSP is arranged to convert the data to analog values and to write the analog values into the respective memory cells, to estimate respective distortion levels affecting the analog values written into the memory cells and to determine the achievable storage capacities responsively to the distortion levels.
54. The apparatus according to any of claims 50-53, wherein the MSP is arranged to read the analog values from the memory cells and to estimate the distortion levels based on the analog values read from the memory cells.
55. The apparatus according to claim 54, wherein the MSP is arranged to reconstruct the data from the analog values read from the memory cells, and to calculate the distortion levels based on the reconstructed data.
56. The apparatus according to claim 55, wherein the MSP is arranged to calculate the distortion levels by summing scalar functions of the analog values read from the memory cells.
57. The apparatus according to claim 56, wherein the MSP is arranged to sum squares of differences between the analog values read from the memory cells and respective expected values of the data stored in the memory cells responsively to the data.
58. The apparatus according to claim 56, wherein the MSP is arranged to calculate the distortion levels by counting a number of memory cells in which a difference between the analog values read from the memory cells and respective expected values of the data stored in the memory cells responsively to the data exceeds a predetermined value.
59. The apparatus according to claim 53, wherein the MSP is arranged to write the analog values into the memory cells and verify the written analog values using an iterative Program and Verify (P&V) process, and to calculate the distortion levels based on the analog values verified by the P&V process.
60. The apparatus according to claim 53, wherein the MSP is arranged to write the analog values into the memory cells and verify the written analog values using an iterative Program and Verify (P&V) process, which iteratively increments the analog values by a programming

step, and wherein the storage configurations define a size of the programming step used by the iterative P&V process.

61. The apparatus according to any of claims 50-53, wherein the MSP is arranged to encode the data to be stored in a group of the memory cells using an Error Correcting Code (ECC), to convert the encoded data to analog values for storage in the respective analog memory cells of the group, to read the analog values from the memory cells in the group, to evaluate a distance metric between the read analog values and a valid codeword of the ECC that is nearest to the analog values and to adapt the storage configuration of the group of the memory cells responsively to the evaluated distance metric.
62. The apparatus according to any of claims 50-53, wherein the MSP is arranged to read the data stored in the memory cells and to modify the storage configurations responsively to detecting errors in the read data.
63. The apparatus according to claim 62, wherein the MSP is arranged to encode the data using an Error Correcting Code (ECC), and to modify the storage configurations by modifying a code rate of the ECC.
64. The apparatus according to any of claims 50-53, wherein the MSP is arranged to track previous programming and erasure operations applied to the memory cells and to estimate the respective achievable capacities responsively to the tracked previous programming and erasure operations.
65. The apparatus according to claim 64, wherein the MSP is arranged to calculate the achievable capacities responsively to lengths of time periods that elapsed since the previous programming and erasure operations.
66. The apparatus according to any of claims 50-53, wherein the MSP is arranged to convert an initial portion of the data to analog values and write the analog values into the respective memory cells, and to subsequently store an additional portion of the data in the memory cells by increasing the analog values written to at least some of the memory cells without erasing the cells.
67. The apparatus according to any of claims 50-53, wherein the MSP is arranged to encode the data using an Error Correcting Code (ECC) that adds redundancy bits to the data and storing the redundancy bits in some of the analog memory cells, and to modify the storage configurations by modifying a number of the redundancy bits added by the ECC without erasing the cells.

68. The apparatus according to claim 67, wherein the MSP is arranged to store the redundancy bits separately from the data.

69. The apparatus according to any of claims 50-53, wherein the MSP is arranged to modify the storage configurations by selecting updated storage configurations from a predefined set of possible storage configurations.

70. The apparatus according to any of claims 50-53, wherein the quantity of the data defined by at least one of the storage configurations specifies a non-integer number of bits per cell.

71. The apparatus according to any of claims 50-53, wherein the MSP is arranged to compress the data before writing the data to the memory cells, and wherein the storage configurations define respective compression ratios by which the data is to be compressed.

72. The apparatus according to any of claims 50-53, wherein the MSP is arranged to define the storage configurations for respective groups of the memory cells.

73. The apparatus according to any of claims 50-53, wherein the MSP is arranged to accept a data item for storage in the memory, to select a subset of the memory cells in which to store the data item responsively to the assigned storage configurations, and to store the data item in the selected subset of the memory cells.

74. The apparatus according to any of claims 50-53, wherein the MSP is arranged to re-estimate the achievable storage capacities during idle time periods in which the data is not stored and read.

75. The apparatus according to any of claims 50-53, wherein the MSP is arranged to store the storage configurations in a configuration table.

76. The apparatus according to any of claims 50-53, wherein the MSP is arranged to maintain a predetermined margin between the quantities of the data stored in the memory cells and the respective estimated achievable capacities.

77. The apparatus according to any of claims 50-53, wherein at least some of the distortion is caused by leakage current in the analog memory cells, and wherein the MSP is arranged to estimate the leakage current.

78. Data storage apparatus, comprising:  
an interface, which is arranged to communicate with a memory that includes a plurality of analog memory cells; and

a memory signal processor (MSP), which is arranged to estimate distortion levels in the respective analog memory cells, to estimate respective achievable storage capacities of the analog memory cells based on the estimated distortion levels, to assign the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities, and to store the data in the memory cells in accordance with the respective assigned storage configurations.

79. Data storage apparatus, comprising:

an interface, which is arranged to communicate with a memory that includes a plurality of analog memory cells; and

a memory signal processor (MSP), which is arranged to track respective achievable storage capacities of the memory cells while the memory is in use in a host system, to accept data for storage in the memory, to select a subset of the memory cells for storing the data based on the tracked achievable capacities, and to store the data in the memory cells of the subset.

80. The apparatus according to claim 79, wherein the MSP is arranged to select the memory cells whose aggregate achievable capacity best matches a size of the accepted data.

81. The apparatus according to claim 79, wherein the MSP is arranged to accept a requested reliability level for storing the data and to select the memory cells responsively to the requested reliability level.

82. The apparatus according to claim 79, wherein the MSP is arranged to accept a requested retention period for storing the data and to select the memory cells responsively to the requested retention period.

83. The apparatus according to claim 79, wherein the MSP is arranged to select the memory cells having low distortion levels with respect to other memory cells.

84. The apparatus according to claim 79, wherein the MSP is arranged to track previous programming and erasure operations applied to the memory cells and to select the memory cells responsively to the previous programming and erasure operations.

85. The apparatus according to claim 84, wherein the MSP is arranged to distribute a number of programming and erasure operations evenly over the memory cells in the memory by selecting the memory cells having a smaller number of the previous programming and erasure operations with respect to other memory cells.

86. The apparatus according to claim 85, wherein the MSP is arranged to select the memory cells having the smaller number of the previous programming and erasure operations performed during a predetermined recent time period with respect to the other memory cells.

5 87. The apparatus according to claim 79, wherein the MSP is arranged to calculate and report to the host system a size of an available memory space by summing the achievable capacities of the memory cells that are available for storing the data.

88. The apparatus according to claim 79, wherein the memory is divided into multiple erasure blocks, each erasure block comprising a group of the memory cells that are erased in a single erasure operation, and wherein the MSP is arranged to erase a data item from the  
10 memory by:

identifying one or more of the erasure blocks in which the data item is stored;

when the erasure blocks in which the data item is stored contain stored data of another data item, identifying alternative memory cells outside the erasure blocks in which the data item is stored based on the tracked achievable capacities of the memory cells and on a size of  
15 the other data item, and copying the data of the other data item to the alternative memory cells; and

erasing the erasure blocks in which the data item is stored.

89. The apparatus according to claim 79, wherein the MSP is arranged to accept the data from the host system in fixed-capacity blocks and to store the data in variable-capacity groups  
20 of the memory cells based on the tracked achievable capacities of the variable-size groups.

90. The apparatus according to claim 89, wherein the MSP is arranged to identify two or more variable-capacity groups having capacities that are smaller than a capacity of the fixed-capacity block, and to store the data accepted in one or more of the fixed-capacity blocks in the identified two or more variable-capacity groups.

25 91. The apparatus according to claim 89, wherein the MSP is arranged to allocate one or more of the variable-capacity groups to serve as an excess memory storage area, to store the data accepted in the fixed-capacity blocks in respective variable-capacity groups, and, when the achievable capacity of a variable-capacity group is smaller than a capacity of the fixed-capacity blocks, to store some of the data accepted in the respective fixed-capacity block in  
30 the allocated excess memory storage area.

92. The apparatus according to claim 89, wherein the MSP is arranged to store the data accepted in the fixed-capacity blocks sequentially in the memory cells irrespective of boundaries between the variable-capacity groups.

5 93. The apparatus according to claim 92, wherein the MSP is arranged to initially store the data accepted in each fixed-capacity block in a respective variable-capacity group, and to subsequently re-program the initially-stored data sequentially, irrespective of the boundaries between the variable-capacity groups.

10 94. The apparatus according to claim 89, wherein the MSP is arranged to retrieve the data from the variable-size groups, to reconstruct the fixed-size blocks and to output the data using the fixed-size blocks to the host system.

95. The apparatus according to claim 79, wherein the MSP is arranged to reduce a number of memory access operations applied to the memory by caching at least some of the data.

96. The apparatus according to claim 95, wherein the MSP is arranged to transfer the cached data to the memory upon detecting an immediate storage event.

15 97. The apparatus according to claim 96, wherein the immediate storage event comprises at least one event selected from a group of events consisting of an approaching power failure, a time out and an End OfFile (EOF) command accepted from the host system.

98. Data storage apparatus, comprising:

a memory, which comprises a plurality of analog memory cells; and

20 a memory signal processor (MSP), which is coupled to the memory and is arranged to estimate respective achievable storage capacities of the memory cells, to assign the memory cells respective storage configurations defining quantities of data to be stored in the memory cells based on the estimated achievable capacities, to store the data in the memory cells in accordance with the respective assigned storage configurations, to re-estimate the respective  
25 achievable storage capacities of the analog memory cells after the memory has been installed in a host system and used for storing the data in the host system, and to modify the storage configurations responsively to the re-estimated achievable capacities.

99. The apparatus according to claim 98, wherein the memory resides in a first Integrated Circuit (IC) and wherein the MSP resides in a second IC different from the first IC.

30 100. The apparatus according to claim 98, wherein the memory and the MSP are integrated in a single device.

101. The apparatus according to claim 98, wherein the MSP is embodied in a processor of the host system.
102. The apparatus according to claim 98, wherein the memory cells comprise Flash memory cells.
- 5 103. The apparatus according to claim 98, wherein the memory cells comprise Dynamic Random Access Memory (DRAM) cells.
104. The apparatus according to claim 98, wherein the memory cells comprise Phase Change Memory (PCM) cells.
105. The apparatus according to claim 98, wherein the memory cells comprise Nitride  
10 Read-Only Memory (NROM) cells.
106. The apparatus according to claim 98, wherein the memory cells comprise Magnetic Random Access Memory (MRAM) cells.
107. The apparatus according to claim 98, wherein the memory cells comprise Ferroelectric Random Access Memory (FRAM) cells.
- 15 108. Data storage apparatus, comprising:  
a memory, which comprises a plurality of analog memory cells; and  
a memory signal processor (MSP), which is coupled to the memory and is arranged to track respective achievable storage capacities of the memory cells while the memory is in use in a host system, to accept data for storage in the memory, to select a subset of the memory  
20 cells for storing the data based on the tracked achievable capacities, and to store the data in the memory cells of the subset.
109. The apparatus according to claim 108, wherein the memory resides in a first Integrated Circuit (IC) and wherein the MSP resides in a second IC different from the first IC.
110. The apparatus according to claim 108, wherein the memory and the MSP are  
25 integrated in a single device.
111. The apparatus according to claim 108, wherein the MSP is embodied in a processor of the host system.
112. The apparatus according to claim 108, wherein the memory cells comprise Flash memory cells.

113. The apparatus according to claim 108, wherein the memory cells comprise Dynamic Random Access Memory (DRAM) cells.

114. The apparatus according to claim 108, wherein the memory cells comprise Phase Change Memory (PCM) cells.

5 115. The apparatus according to claim 108, wherein the memory cells comprise Nitride Read-Only Memory (NROM) cells.

116. The apparatus according to claim 108, wherein the memory cells comprise Magnetic Random Access Memory (MRAM) cells.

10 117. The apparatus according to claim 108, wherein the memory cells comprise Ferroelectric Random Access Memory (FRAM) cells.

FIG. 1

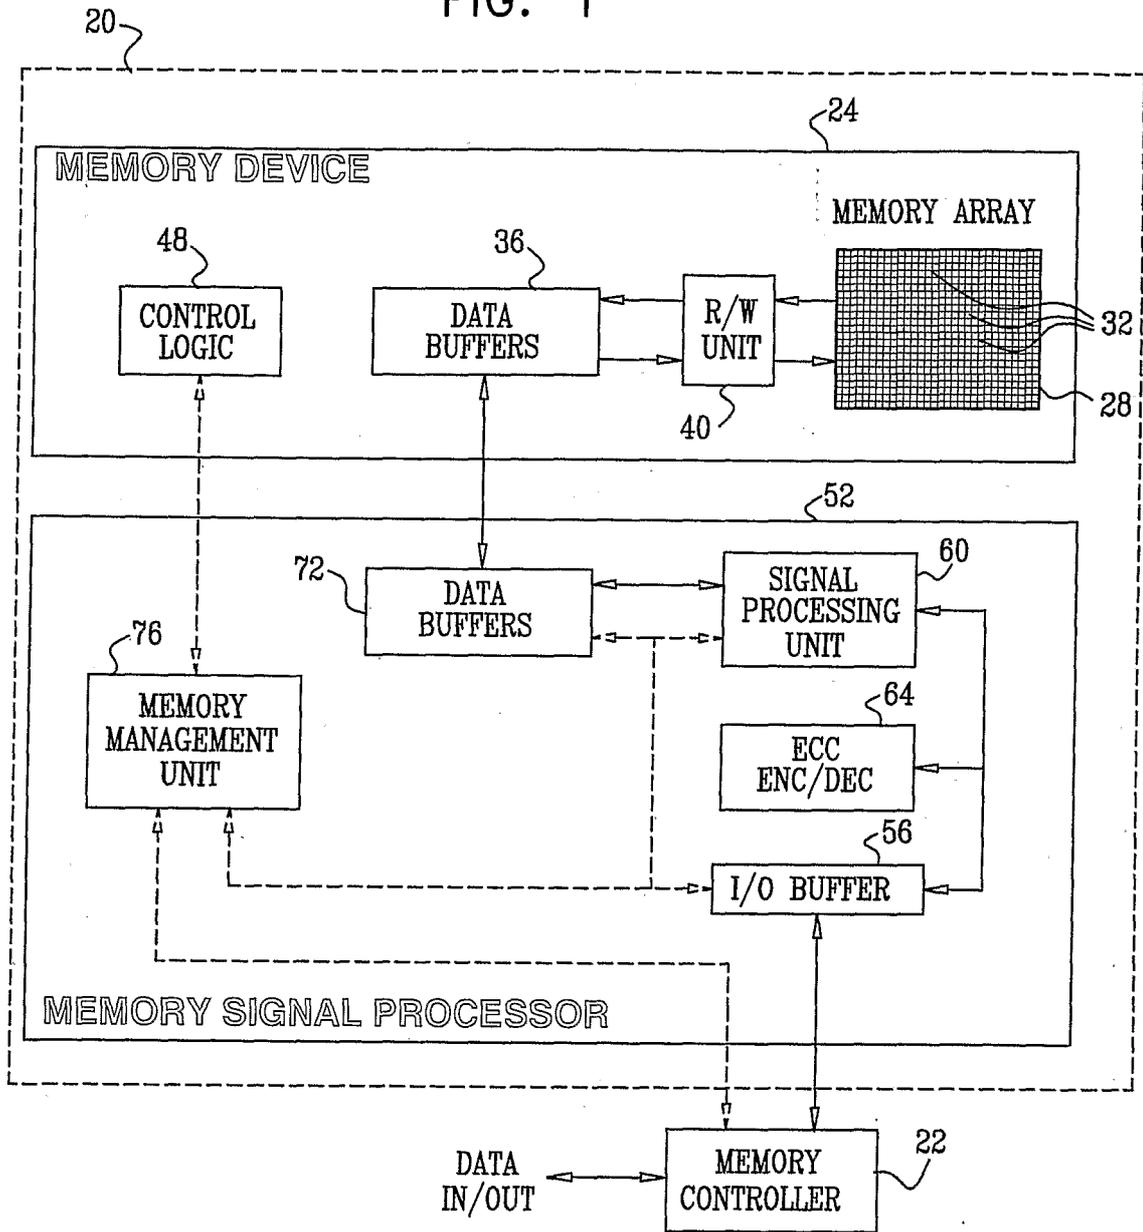


FIG. 2

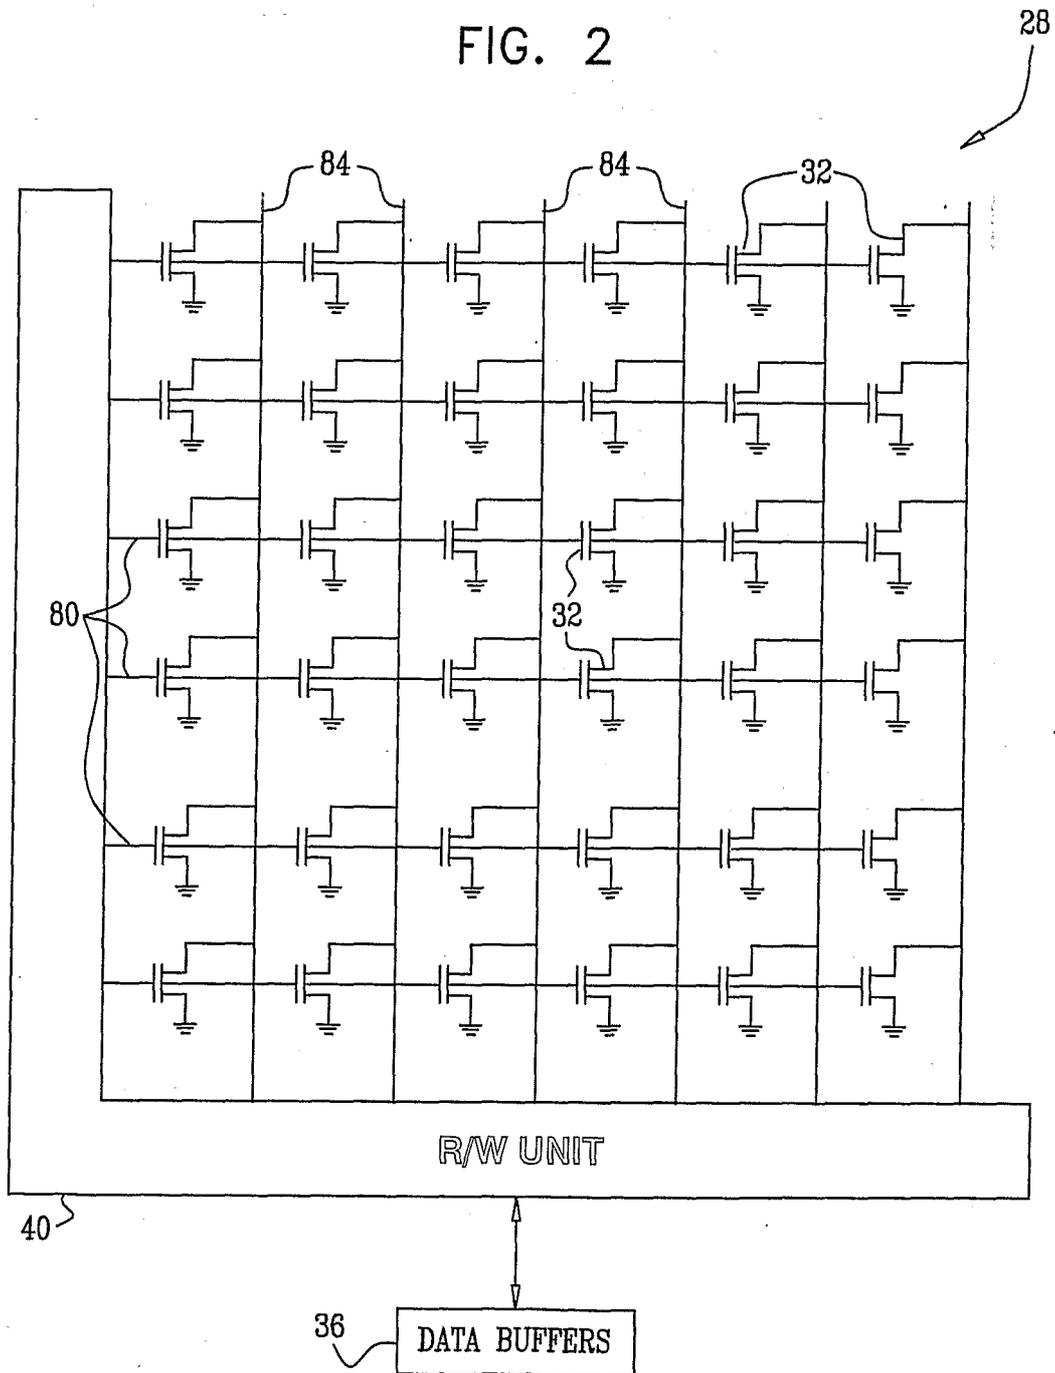


FIG. 3

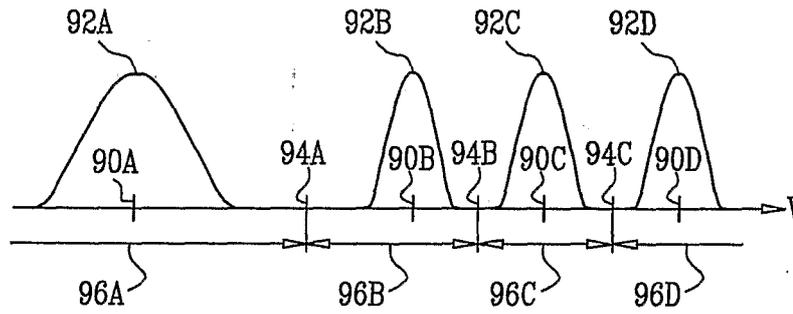
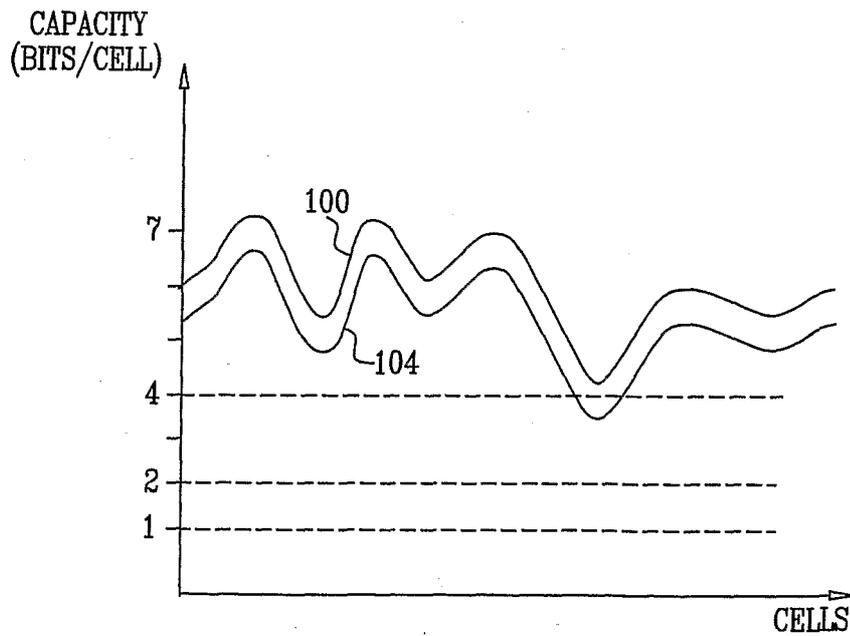


FIG. 4



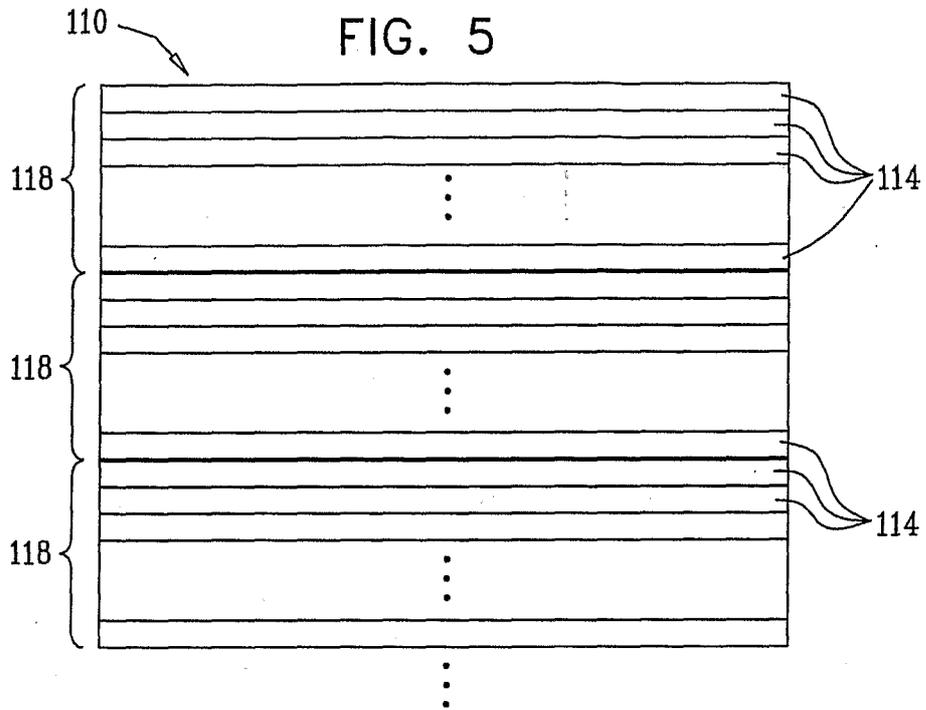
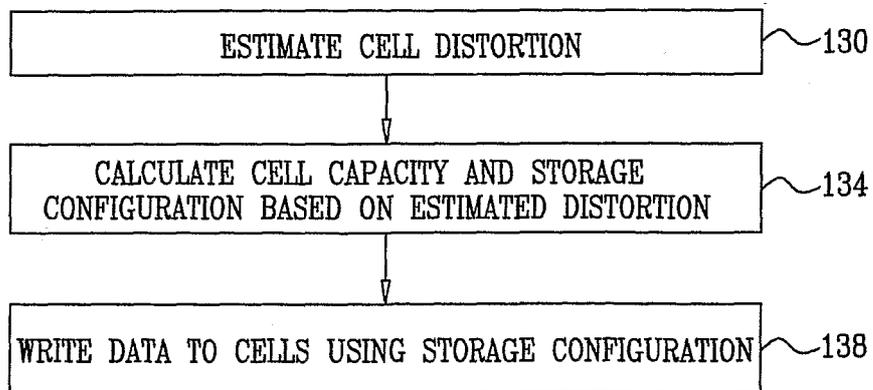


FIG. 6



5/10

FIG. 7

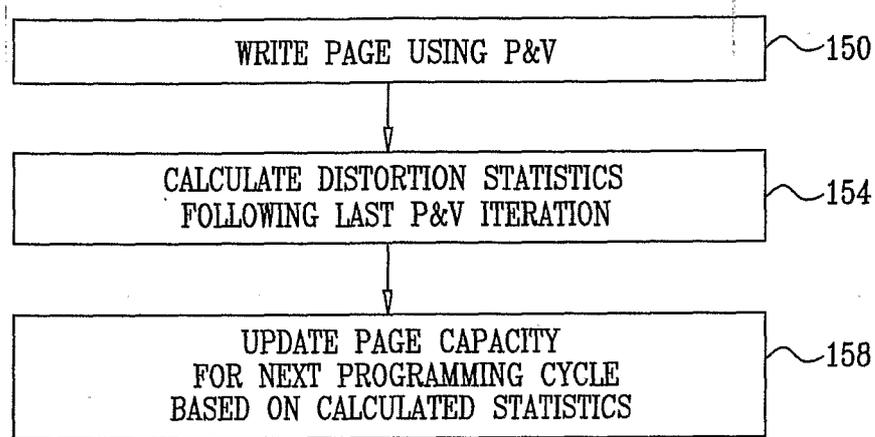
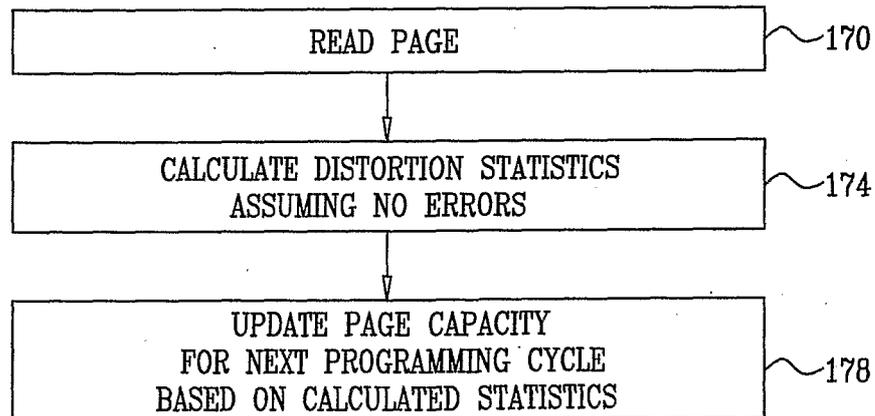


FIG. 8



6/10

FIG. 9

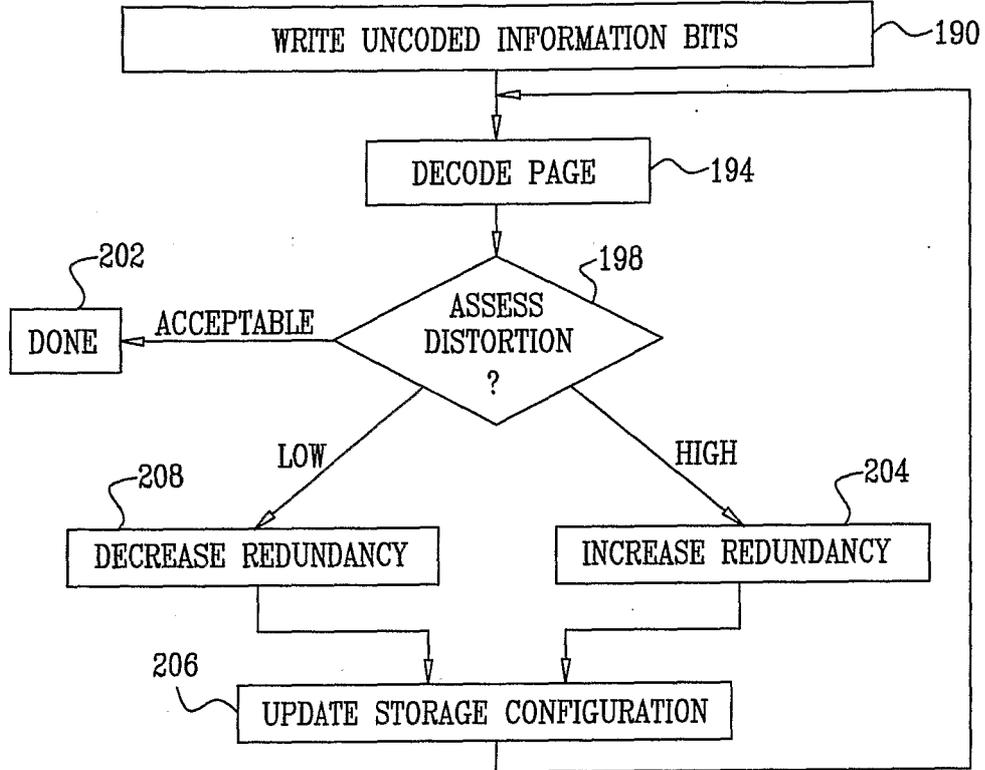
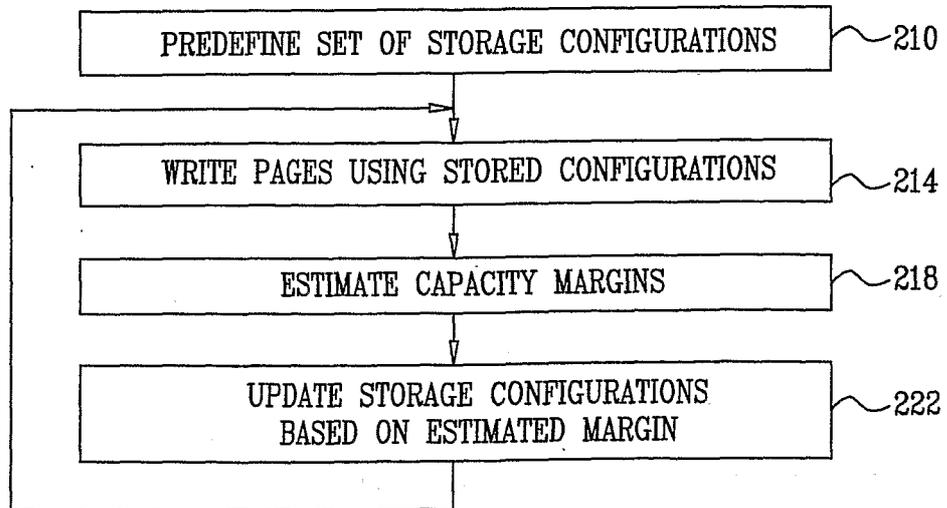


FIG. 10



7/10

FIG. 11

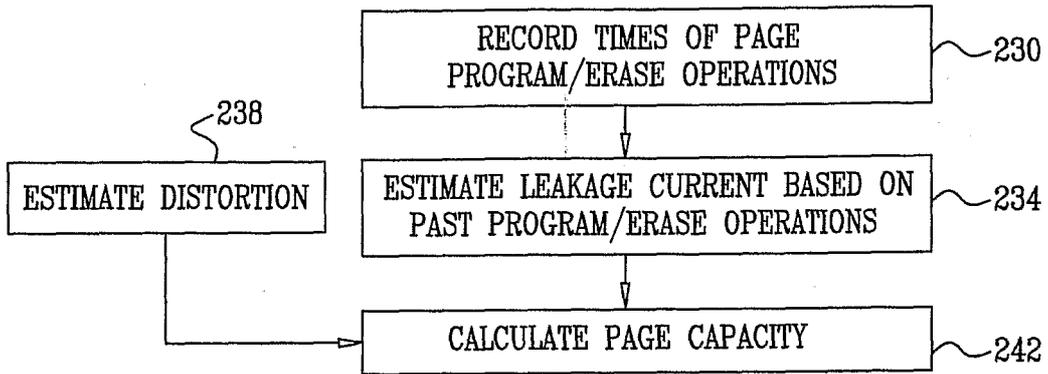


FIG. 12

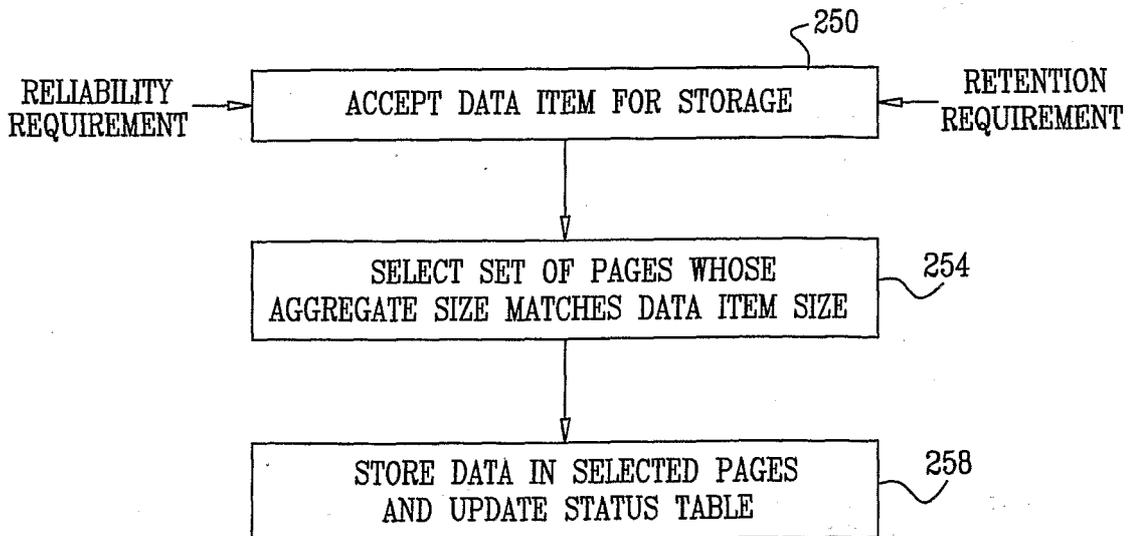
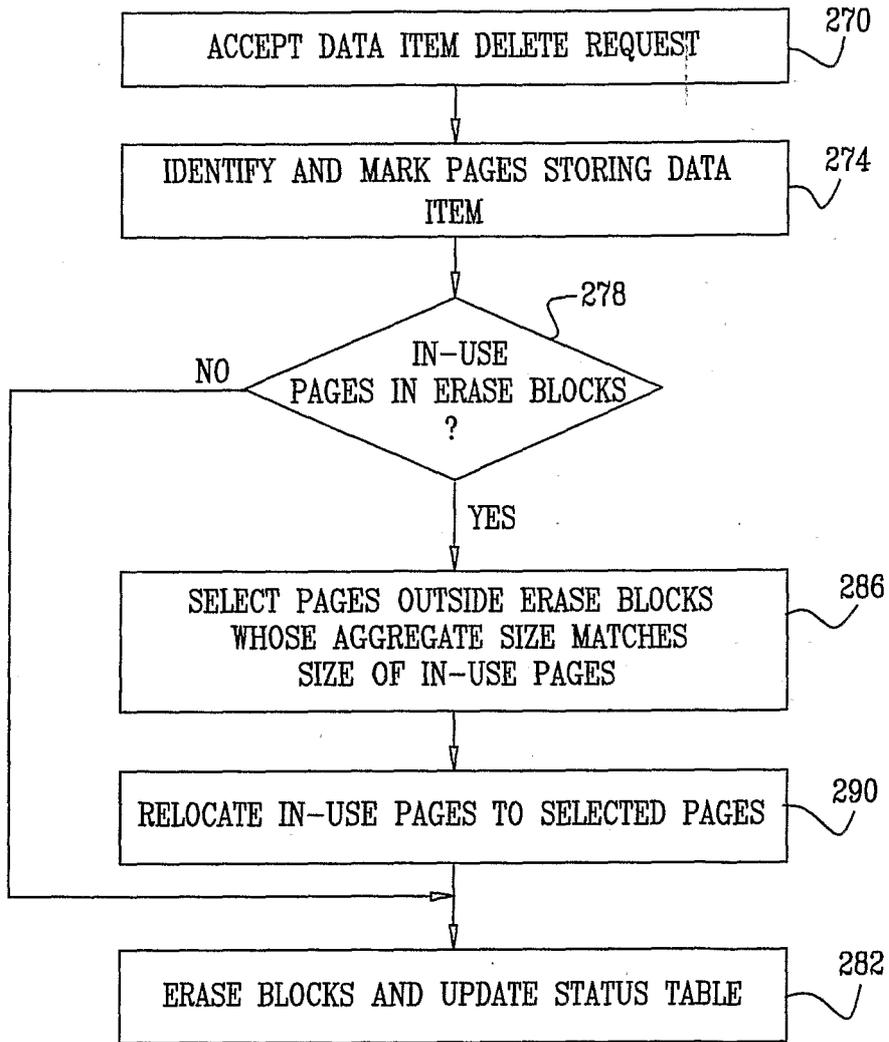


FIG. 13



9/10

FIG. 14

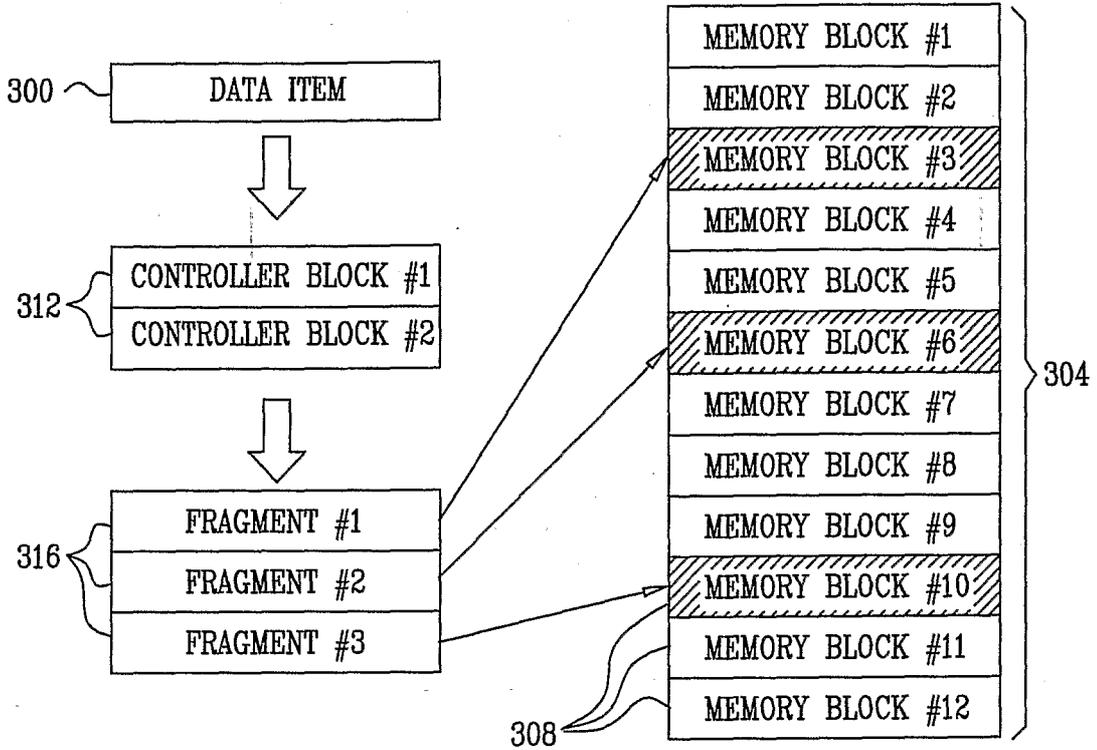
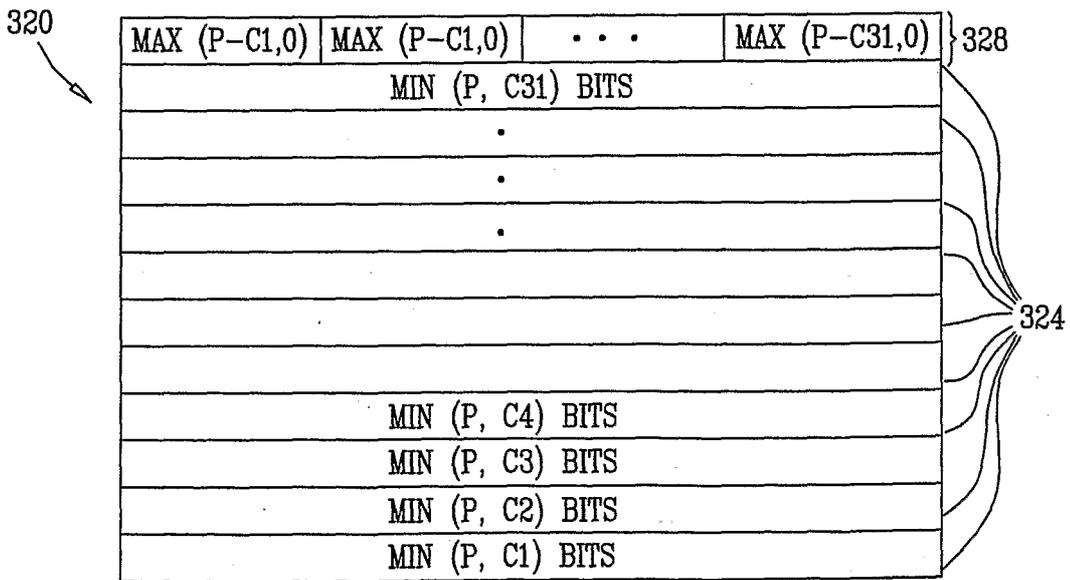


FIG. 15



10/10

FIG. 16

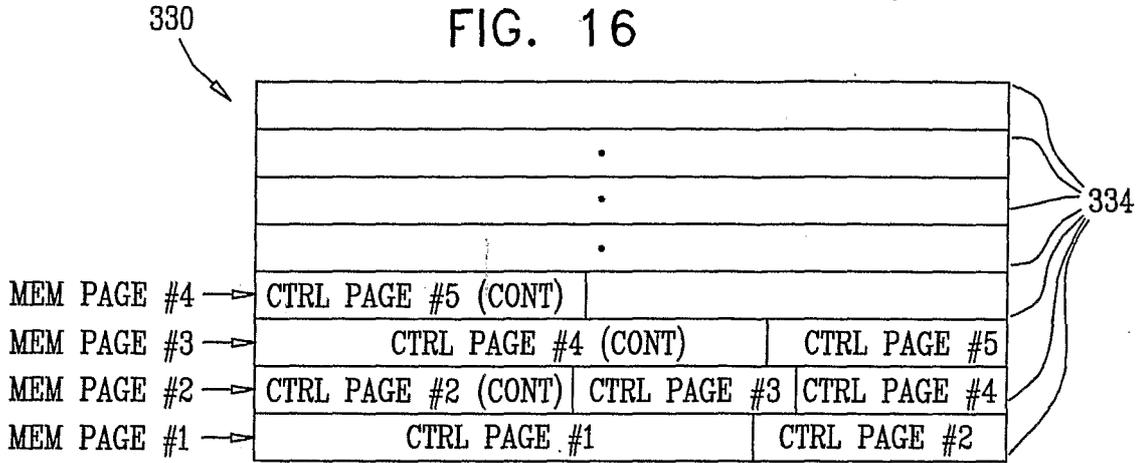


FIG. 17

