

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 April 2009 (30.04.2009)

PCT

(10) International Publication Number
WO 2009/053962 A2

(51) International Patent Classification:
G06F 7/74 (2006.01)

(74) Agent: PEARL COHEN ZEDEK LATZER; 5 Shenkar
Street, P.O. Box 12704, Herzlia 46733 (IL).

(21) International Application Number:
PCT/IL2008/001238

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA,
CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE,
EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID,
IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK,
LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW,
MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT,
RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ,
TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM,
ZW.

(22) International Filing Date:
17 September 2008 (17.09.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/960,943 22 October 2007 (22.10.2007) US
61/071,469 30 April 2008 (30.04.2008) US

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

(71) Applicant (for all designated States except US): DENS-
BITS TECHNOLOGIES LTD. [IL/IL]; Building 30, Flr
3, Matam Industrial Park, P.O.Box 15111, 31015 Haifa
(IL).

(72) Inventor; and

(75) Inventor/Applicant (for US only): WEINGARTEN,
Hanan [IL/IL]; Hayarden 2, 46377 Herzliya (IL).

Published:

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

(54) Title: SYSTEMS AND METHODS FOR AVERAGING ERROR RATES IN NON-VOLATILE DEVICES AND STORAGE
SYSTEMS

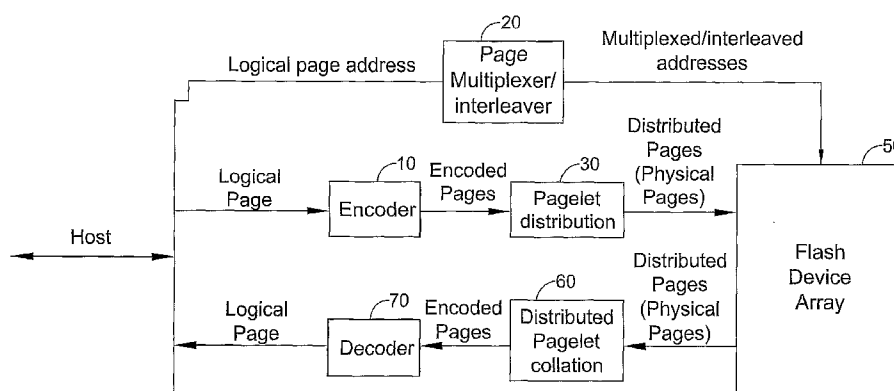


Figure 1a

(57) Abstract: A system for storing a plurality of logical pages in a set of at least one flash device, each flash device including a set of at least one erase block, the system comprising apparatus for distributing at least one of the plurality of logical pages over substantially all of the erase blocks in substantially all of the flash devices, thereby to define, for at least one logical page, a sequence of pagelets thereof together including all information on the logical page and each being stored within a different erase block in the set of erase blocks; and apparatus for reading each individual page from among the plurality of logical pages including apparatus for calling and ordering the sequence of pagelets from different erase blocks in the set of erase blocks.

SYSTEMS AND METHODS FOR AVERAGING ERROR RATES IN NON-VOLATILE DEVICES AND STORAGE SYSTEMS

5

REFERENCE TO CO-PENDING APPLICATIONS

Priority is claimed from the following co-pending applications: US Provisional Application No. 60/960,943, filed October 22, 2007 and entitled "Systems and methods
10 to reduce errors in Solid State Disks and Large Flash Devices" and US Provisional Application No. 61/071,469, filed April 30, 2008 and entitled "Systems and Methods for Averaging Error Rates in Non-Volatile Devices and Storage Systems".

Other co-pending applications include: US Provisional Application No. 60/960,207, filed September 20, 2007 and entitled "Systems and Methods for Coupling
15 Detection in Flash Memory", US Provisional Application No. 61/071,467, filed April 30, 2008 and entitled "Improved Systems and Methods for Determining Logical Values of Coupled Flash Memory Cells", US Provisional Application No. 60/996,027, filed October 25, 2007 and entitled "Systems and Methods for Coping with Variable Bit Error Rates in Flash Devices", US Provisional Application No. 61/071,466, filed April
20 30, 2008 and entitled "Systems and Methods for Multiple Coding Rates in Flash Devices", US Provisional Application No. 61/006,120, filed December 19, 2007 and entitled "Systems and Methods for Coping with Multi Stage Decoding in Flash Devices", US Provisional Application No. 61/071,464, filed April 30, 2008 and entitled "A Decoder Operative to Effect A Plurality of Decoding Stages Upon Flash Memory
25 Data and Methods Useful in Conjunction Therewith", US Provisional Application No. 61/006,385, filed January 10, 2008 and entitled "A System for Error Correction Encoder and Decoder Using the Lee Metric and Adapted to Work on Multi-Level Physical Media", US Provisional Application No. 61/064,995, filed April 8, 2008 and entitled "Systems and Methods for Error Correction and Decoding on Multi-Level Physical
30 Media", US Provisional Application No. 60/996,948, filed December 12, 2007 and entitled "Low Power BCH/RS Decoding: a Low Power Chien-Search Implementation", US Provisional Application No. 61/071,487, filed May 1, 2008 and entitled "Chien-Search System Employing a Clock-Gating Scheme to Save Power for Error Correction Decoder and other Applications", US Provisional Application No. 61/071,468, filed

-2-

April 30, 2008 and entitled "A Low Power Chien-Search Based BCH/RS Recoding System for Flash Memory, Mobile Communications Devices and Other Applications", US Provisional Application No. 60/996,782, filed December 5, 2007 and entitled "Systems and Methods for Using a Training Sequence in Flash Memory", US Provisional Application No. 61/064,853, filed March 31, 2008 and entitled "Flash Memory Device with Physical Cell Value Deterioration Accommodation and Methods Useful in Conjunction Therewith", US Provisional Application No. 61/129,608, filed July 8, 2008 and entitled "A Method for Acquiring and Tracking Detection Thresholds in Flash Devices", US Provisional Application No. 61/006,806, filed January 31, 2008 and entitled "Systems and Methods for using a Erasure Coding in Flash memory", US Provisional Application No. 61/071,486, filed May 1, 2008 and entitled "Systems and Methods for Handling Immediate Data Errors in Flash Memory", US Provisional Application No. 61/006,078, filed December 18, 2007 and entitled "Systems and Methods for Multi Rate Coding in Multi Level Flash Devices", US Provisional Application No. 61/064,923, filed April 30, 2008 and entitled "Apparatus For Coding At A Plurality Of Rates In Multi-Level Flash Memory Systems, And Methods Useful In Conjunction Therewith", US Provisional Application No. 61/006,805, filed January 31, 2008 and entitled "A Method for Extending the Life of Flash Devices", US Provisional Application No. 61/071,465, filed April 30, 2008 and entitled "Systems and Methods for Temporarily Retiring Memory Portions", US Provisional Application No. 61/064,760, filed March 25, 2008 and entitled "Hardware efficient implementation of rounding in fixed-point arithmetic", US Provisional Application No. 61/071,404, filed April 28, 2008 and entitled "Apparatus and Methods for Hardware-Efficient Unbiased Rounding", US Provisional Application No. 61/136,234, filed August 20, 2008 and entitled "A Method Of Reprogramming A Non-Volatile Memory Device Without Performing An Erase Operation", US Provisional Application No. 61/129,414, filed June 25, 2008 and entitled "Improved Programming Speed in Flash Devices Using Adaptive Programming", and several other co-pending patent applications being filed concurrently (same day).

FIELD OF THE INVENTION

The present invention relates generally to flash memory devices and more particularly to methods and systems for writing on such devices and/or for reading from such devices.

The state of the art is believed to be represented by the following documents inter alia:

US Patent Document 20070168625, entitled "Interleaving policies for flash memory"; and

United States Patent No. 6,996,004, entitled "Minimization of FG-FG coupling in flash memory".

Paulo Cappelletti, Clara Golla, Piero Olivo, Enrico Zanoni, "Flash Memories", Kluwer Academic Publishers, 1999

G. Campardo, R. Micheloni, D. Novosel, "CLSI-Design of Non-Volatile Memories", Springer Berlin Heidelberg New York, 2005

Wear leveling, according to Wikipedia, "is a technique for prolonging the service life of some kinds of erasable computer storage media, such as flash memory. The term has also been used by Western Digital to describe their hard disk preservation technique, but hard disks are not generally wear-leveled devices...".

EEPROM and flash memory media have individually erasable segments, each of which can be put through a finite number of erase cycles before becoming unreliable. This can be anywhere between 10,000 and 1,000,000 cycles, for example, for NAND flash devices. Erasable optical media such as CD-RW and DVD-RW are rated at up to 1,000 cycles (100,000 cycles for DVD-RAM media). Wear-leveling attempts to work around these limitations by arranging data so that erasures and re-writes are distributed evenly across the medium. In this way, no single sector prematurely fails due to a high concentration of write cycles.

"Conventional file systems like FAT, ext2 and NTFS were originally designed for magnetic disks and as such rewrite many of their data structures (such as their directories) repeatedly in place. Some file systems aggravate the problem by tracking last-access times, which can lead to file metadata being constantly rewritten in-place. There are several techniques for extending the life of the media: a. A checksum or error-correcting code can be kept for block or sector in order to detect errors or correct errors;

-4-

b. A pool of reserve space can also be kept. When a block or sector does fail, future reads and writes to it can be redirected to a replacement in that pool; c. Blocks or sectors on the media can be tracked in a least recently used queue of some sort. The data structures for the queue itself must either be stored off-device or in such a way that the space it uses is itself wear-leveled.

"On flash memory devices, such as CompactFlash and Secure Digital cards, these techniques are implemented in hardware by a built-in microcontroller. On such devices, wear-leveling is transparent and most conventional file systems can be used as-is on them. Wear-leveling can also be implemented in software by special-purpose file systems such as JFFS2 and YAFFS on flash media or UDF on optical media. All three are log-structured file systems in that they treat their media as circular logs and write to them in sequential passes. Some storage interfaces do not in themselves perform wear leveling."

U.S. Patent 6,850,443 describes wear leveling techniques for flash EEPROM systems.

The disclosures of all publications and patent documents mentioned in the specification, and of the publications and patent documents cited therein directly or indirectly, are hereby incorporated by reference.

SUMMARY OF THE INVENTION

The following terms may be construed either in accordance with any definition thereof
5 appearing in the prior art literature, or in accordance with any definition implied by any
portion of the specification, or as follows:

Bank, flash memory bank = ordered set (e.g. sequence or array) of flash memory), all
connected to a single data source, such as a PC through the USB interface sr, via a
10 unique and separate interface such as NAND Flash interface. Since different banks
each may be connected to the data source via their own interface, they may be
accessible in parallel.

Bit error rate = the average proportion of bits within a physical page, which are
15 erroneous. This may be measured through experiments in which several pages (of the
same type) are programmed and following cycling and retention are read back and
compared to their original contents. The bit error rate is then the ratio between the
average number of bit errors per page and the length of the page in bits. The bit error
rate (BER) may vary between individual devices.

20 Block = a set of flash memory device cells which must, due to physical limitations of
the flash memory device, be erased together. Also termed erase sector, erase block.
Each block typically comprises an ordered set (e.g. sequence or array) of at least one
physical page which are all erased simultaneously by each erase operation, the set
25 typically comprising a predetermined number of typically physically adjacent physical
pages.

Cell: Apparatus into which any of at least two statistically distinguishable ranges of a
physical quantity, such as electrical charge, can be induced. Each detectably distinct
30 range has a predetermined association with a digital value, also termed herein "logical
value", which is said to be "stored" in the cell. Since induction in a cell, of a target
amount of a physical quantity is typically not deterministic and instead typically
modeled by a random, Gaussian process centered about the target amount, "storing" an

-6-

individual digital value in a cell typically comprises striving to induce, in that cell, a target amount of the physical quantity which comprises the center of the range pre-associated with the individual digital value. As a result, the actual amount induced in the cell is, statistically, almost always within the range associated with the individual digital value, at least initially, hence the individual digital value can, statistically, almost always be correctly read from the cell, at least initially. In flash memory applications, a component of flash memory that stores one bit of information (in single-level cell devices) or n bits of information (in a multi-level device having $2^{\exp n}$ levels). Typically, each cell comprises a floating-gate transistor. n may or may not be an integer.

Charge level or degraded level: Amount of charge currently existing in a cell, as opposed to "program level", the amount of charge originally induced in the cell which typically exceeds the current (charge) level due to the phenomenon of gradual degradation of the amount of charge in a cell, over time.

Cycling: Repeatedly writing new data into flash memory cells and repeatedly erasing the cells between each two writing operations.

Demapping: basic cell-level reading function in which incoming digital n -type data is derived from a physical value representing a physical state in the cell having a predetermined correspondence to the incoming digital value.

Digital value: n -tuple of bits represented by a cell in flash memory capable of generating $2^{\exp n}$ distinguishable levels of a physical value such as charge, where n may or may not be an integer.

Erase cycle: The relatively slow process of erasing a block of cells, or, in certain non-flash memory devices, of erasing a single cell or the duration of so doing. An advantage of erasing cells collectively in blocks as in flash memory, rather than individually, is enhanced programming speed: Many cells are erased in a single erase cycle.

Erase-write cycle: The process of erasing a block of cells and subsequently writing new data into at least some of them.

-7-

Flash memory: Non-volatile rewritable computer memory including cells that are erased in blocks rather than individually, but are written into and read from, in smaller units e.g. individually or page by page where at least some of the blocks include more than one page each. Includes NOR-type flash memory, NAND-type flash memory, and PRAM, e.g. Samsung PRAM, inter alia. A flash memory device typically comprises ordered set (e.g. sequence or array) of flash memory cells in conjunction with reading circuitry, writing circuitry and erasing circuitry.

10 Input/output sequence = I/O sequence = transaction: Request from host to read or write a given number of logical pages from or to a flash memory system. A sequence is "short" if the number of logical pages is less than the number of banks in the flash memory system.

15 Interleaving: When writing a logical page into flash memory, distributing the data therein over disparate locations in the flash memory, so as to greatly diminish the probability that the logical page will be stored at a poor quality portion of the flash memory, such that the logical page, when read, will contain more errors than its error correction code can overcome. Distribution need not be regular or strictly alternating.

20

Mapping: basic cell-level writing function in which incoming digital n-tuple is mapped to a physical value by inducing a physical state in the cell, having a predetermined correspondence to the incoming digital value.

Page = A portion, typically 512 or 2048 or 4096 bytes in size, of a memory device such as a NAND flash memory device. Writing can be performed page by page, as opposed to erasing which can be performed only erase sector by erase sector. A few bytes, typically 16 – 32 for every 512 data bytes (termed redundancy bytes) are associated with each page (typically 16, 64 or 128 per page), for storage of error correction information. A typical block may include 32 512-byte pages or 64 2048-byte pages.

30

Precise read, soft read: Read the cell threshold voltage in precision greater than the number of the Mapping levels (2^n). The terms precise read or soft read are interchangeable.

Programming: The process of writing new data in at least some of the pages of an erase sector.

- 5 Program level (or: programmed level, originally written level): amount of charge originally induced in a cell to represent a given digital value, as opposed to "charge level".

10 Reprogrammability (N_p): An aspect of flash memory quality. This is typically operationalized by a reprogrammability parameter, also termed herein " N_p ", denoting the number of times that a flash memory or a sub-section of the memory such as a page can be re-programmed (number of erase-write cycles that the device can withstand) before the level of errors is so high as to make an unacceptably number of those errors irrecoverable with high probability given a predetermined amount of memory devoted to redundancy. Typically recoverability is investigated following a conventional aging
15 simulation process which simulates or approximates the data degradation effect that a predetermined time period e.g. a 10 year period has on the flash memory device, in an attempt to accommodate for a period of up to 10 years between writing of data in flash memory and reading of the data therefrom.

20

Retention: of original physical levels induced in the cells; retention is typically below 100% resulting in deterioration of original physical levels into current charge levels. . . .

Retention time: Flash storage time without voltage supply or specifically the elapse time
25 between page programs to page read.

Symbol: Digital value

Logical page: a portion of typically sequential data, whose amount is typically less than
30 or equal to a predetermined amount of data defined to be a pageful of data, which has typically been defined by a host (data source/destination) or user thereof, as a page, and which is sent by the host to a flash memory device for storage and is subsequently read by the host from the flash memory device.

Physical page: ordered set (e.g. sequence or array) of flash memory cells which are all written in simultaneously by each write operation, the set typically comprising a predetermined number of typically physically adjacent flash memory cells containing
5 actual data written by and subsequently read by the host, as well as, typically error correction information and back pointers used for recognizing the true address of a page.

Pagelet: portion of a logical page which is to be written onto a particular physical page.

10 Reliability: Reliability of a flash memory device may be operationalized as the probability that a worst-case logical page written and stored in that device for a predetermined long time period such as 10 years will be successfully read i.e. that sufficiently few errors, if any, will be present in the physical page/s storing each logical
15 page such that the error code appended to the logical page will suffice to overcome those few errors.

The applicability of certain embodiments of the present invention includes but is not limited to the following applications:

20 Solid State Disk: A Disk drive built with no moving parts and made of silicon based chips that store the information, such as NAND/NOR Flash, NROM and phase change memories).

Solid State storage applications: Applications which use silicon based chips to store large amounts of memory such as but not limited to USB drives, SD cards, and
25 MP3 players.

Certain embodiments of the invention seek to provide application of interleaving to combat reliability problems in Flash memories.

Certain embodiments of the invention seek to provide breaking of an encoded
30 page into pagelets, typically non-overlapping, which are programmed across an Erase Block and several devices.

Certain embodiments of the invention seek to provide interleaving with different sizes of I/O transactions.

Certain embodiments of the present invention mitigate the bit error rate (BER) problem by distributing each "logical" page across several physical pages in an erase block and/or across several devices. Thus, each "logical" page only suffers from an average bit error rate (BER) and program-erase performance is derived from the average bit error rate (BER) instead of the worst case BER.

According to certain embodiments of the present invention, interleaved programming of the "logical" pages is used, e.g. by breaking them up into pagelets which are programmed into different pages across an erase block and on different devices (if applicable). A read operation of a "logical" page will suffer only from an average bit error rate (BER) and the reliability of the entire system will be a function of the average bit error rate (BER) instead of the worst case BER.

There is thus provided, in accordance with at least one embodiment of the present invention, a flash memory accessory device operative in conjunction with a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of physical pages, the accessory system comprising logical page distributing apparatus for distributing data from each logical page in the stream between at least two of the multiplicity of physical pages.

Further in accordance with at least one embodiment of the present invention, the multiplicity of physical pages is partitioned into a plurality of erase sectors each having physical pages, wherein the physical pages within an individual erase sector can only be erased in unison, wherein the apparatus for distributing distributes data from each logical page in the stream between physical pages within a single erase sector.

Still further in accordance with at least one embodiment of the present invention, the multiplicity of physical pages is partitioned into at least first and second erase sectors each having physical pages, wherein the physical pages within each erase sector can only be erased in unison, wherein the apparatus for distributing distributes information from each logical page in the stream between at least a first physical page within the first erase sector and at least a second physical page within the second erase sector.

Further in accordance with at least one embodiment of the present invention, the flash memory system comprises at least first and second flash memory banks which are accessible in parallel and wherein the apparatus for distributing distributes information from each logical page in the stream between at least a first physical page within the

-11-

first flash memory bank and at least a second physical page within the second flash memory bank.

Also provided, in accordance with at least one embodiment of the present invention, is a flash memory accessory device operative in conjunction with a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of flash memory cells, wherein the multiplicity of flash memory cells can be divided into several subsets of flash memory cells each having a characteristic range of bit error rates which differs from the characteristic range of bit error rates of at least one other subset of flash memory cells, the accessory device comprising bit error rate (BER) distributing apparatus operative to distribute data from each logical page in the stream between at least two of the subsets.

Further in accordance with at least one embodiment of the present invention, the multiplicity of flash memory cells can be partitioned into several orthogonal partitions of flash memory cells, each partition defining several subsets into which the multiplicity of flash memory cells is divided, each subset having a characteristic range of bit error rates which differs from the characteristic range of bit error rates of at least one other subset of flash memory cells defined by the same partition, wherein the bit error rate (BER) distributing apparatus is operative to distribute data from each logical page in the stream between at least two subsets defined by each of at least two of the partitions.

Further in accordance with at least one embodiment of the present invention, one of the orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into physical pages.

Additionally in accordance with at least one embodiment of the present invention, the physical pages include at least first and second layers and wherein the multiplicity of flash memory cells comprises a multiplicity of multi-level cells each storing at least two bits including a most significant bit (MSB) and a least significant bit (LSB) thereby to define a multiplicity of MSBs and a multiplicity of LSBs and wherein the apparatus for distributing is also operative to distribute at least a first MSB from among the multiplicity of MSBs and a first LSB from among the multiplicity of LSBs to the first layer and to distribute at least a second MSB from among the multiplicity of MSBs and a second LSB from among the multiplicity of LSBs to the second layer.

-12-

Still further in accordance with at least one embodiment of the present invention, one of the orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into erase sectors.

5 Additionally in accordance with at least one embodiment of the present invention, one of the orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into flash memory devices.

Further in accordance with at least one embodiment of the present invention, one of the orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into flash memory banks.

10 Still further in accordance with at least one embodiment of the present invention, the apparatus for distributing is operative to distribute data from each logical page in the stream between at least two subsets defined by each of the partitions.

15 Additionally in accordance with at least one embodiment of the present invention, the logical page distributing apparatus is operative to distribute at least a pagelet within at least a logical page onto a physical page determined by computing a predetermined function of a logical address received from a controller.

Further in accordance with at least one embodiment of the present invention, the first and second physical pages are read from simultaneously when it is desired to access therefrom, a logical page stored therewithin.

20 Additionally in accordance with at least one embodiment of the present invention, the first and second physical pages are written to simultaneously when it is desired to store a logical page therewithin.

25 Additionally in accordance with at least one embodiment of the present invention, the page-distributing apparatus comprises apparatus for partitioning, for at least one set of logical pages, each of several logical pages in the set into pagelets and apparatus for grouping pagelets from the several logical pages together using a predetermined pagelet grouping scheme to form a multi-origin physical page-ful of data containing pagelets originating from more than one logical page.

30 Further in accordance with at least one embodiment of the present invention, the pagelet grouping scheme comprises a pagelet permutation scheme.

Still further in accordance with at least one embodiment of the present invention, the at least one set of logical pages comprises a plurality of sets of logical pages and

-13-

wherein the pagelet permutation scheme is characterized in that not all of the plurality of sets of logical pages are permuted in the same way.

Additionally in accordance with at least one embodiment of the present invention, the device also comprises logical page distributing apparatus for distributing data from each logical page in the stream between at least two of the multiplicity of physical pages, and wherein the bit error rate (BER) distributing apparatus is operative to distribute the two physical pages in two of the subsets respectively.

Also provided, in accordance with at least one embodiment of the present invention, is a system for storing a plurality of logical pages in a set of at least one flash device, each flash device including a set of at least one erase block, the system comprising apparatus for distributing each of the plurality of logical pages over substantially all of the erase blocks in substantially all of the flash devices, thereby to define, for each logical page, a sequence of pagelets thereof together including all information on the logical page and each being stored within a different erase block in the set of erase blocks; and apparatus for reading each individual page from among the plurality of logical pages including apparatus for calling and ordering the sequence of pagelets from different erase blocks in the set of erase blocks.

Further in accordance with at least one embodiment of the present invention, the set of flash devices includes a set of at least one bank of flash devices, each bank in the set of banks including at least one flash device.

Further provided, in accordance with at least one embodiment of the present invention, is a method for facilitating interaction between a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of physical pages, the method comprising distributing data from each logical page in the stream between at least two of the multiplicity of physical pages.

Also provided, in accordance with at least one embodiment of the present invention, is a method for facilitating interaction between a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of flash memory cells, wherein the multiplicity of flash memory cells can be divided into several subsets of flash memory cells each having a characteristic range of bit error rates which differs from the characteristic range of bit error rates of at least one other subset of flash memory cells, the method comprising distributing data from each logical page in the stream between at least two of the subsets.

-14-

Additionally provided, in accordance with at least one embodiment of the present invention, is a method for storing a plurality of logical pages in a set of at least one flash device, each flash device including a set of at least one erase block, the method comprising distributing each of the plurality of logical pages over substantially all of the erase blocks in substantially all of the flash devices, thereby to define, for each logical page, a sequence of pagelets thereof together including all information on the logical page and each being stored within a different erase block in the set of erase blocks; and reading each individual page from among the plurality of logical pages including apparatus for calling and ordering the sequence of pagelets from different erase blocks in the set of erase blocks.

Any suitable processor, display and input means may be used to process, display, store and accept information, including computer programs, in accordance with some or all of the teachings of the present invention, such as but not limited to a conventional personal computer processor, workstation or other programmable device or computer or electronic computing device, either general-purpose or specifically constructed, for processing; a display screen and/or printer and/or speaker for displaying; machine-readable memory such as optical disks, CDROMs, magnetic-optical discs or other discs; RAMs, ROMs, EPROMs, EEPROMs, magnetic or optical or other cards, for storing, and keyboard or mouse for accepting. The term "process" as used above is intended to include any type of computation or manipulation or transformation of data represented as physical, e.g. electronic, phenomena which may occur or reside e.g. within registers and /or memories of a computer.

The above devices may communicate via any conventional wired or wireless digital communication means, e.g. via a wired or cellular telephone network or a computer network such as the Internet.

The apparatus of the present invention may include, according to certain embodiments of the invention, machine readable memory containing or otherwise storing a program of instructions which, when executed by the machine, implements some or all of the apparatus, methods, features and functionalities of the invention shown and described herein. Alternatively or in addition, the apparatus of the present invention may include, according to certain embodiments of the invention, a program as above which may be written in any conventional programming language, and optionally a machine for executing the program such as but not limited to a general purpose

-15-

computer which may optionally be configured or activated in accordance with the teachings of the present invention.

Any trademark occurring in the text or drawings is the property of its owner and occurs herein merely to explain or illustrate one example of how an embodiment of the
5 invention may be implemented.

BRIEF DESCRIPTION OF THE DRAWINGS

Certain embodiments of the present invention are illustrated in the following drawings:

5 Fig. 1A is a simplified block diagram illustration of a system for distributing logical pages evenly within flash device memory of uneven quality, which is constructed and operative in accordance with certain embodiments of the present invention;

10 Fig. 1B is a variation on the system of Fig. 1A in that some of the functional units of Fig. 1A are implemented within a Flash ucontroller which controls a NAND flash memory array, all in accordance with certain embodiments of the present invention;

15 Fig. 2 is a prior art scheme for storing logical pages within flash device memory in which each logical page is stored in one or more adjacent physical pages within a single erase sector, device and bank, such that the reliability of the device is determined by the worst page's bit error rate (BER);

20 Fig. 3 is a scheme, constructed and operative in accordance with certain embodiments of the present invention, for storing logical pages within a flash memory device including a plurality of banks e.g. 4 banks, in which each logical page is divided e.g. partitioned into a plurality of pagelets which are respectively distributed between the plurality of banks;

25 Fig. 4 is a scheme, constructed and operative in accordance with certain embodiments of the present invention, for storing logical pages within flash device memory including at least one erase sector of uneven quality, in which each logical page is divided e.g. partitioned into a plurality of pagelets which are interleaved through at least one erase sector;

 Fig. 5 is a simplified prior art block diagram illustration of an example of a flash memory having L banks, each bank having M devices, each device having K erase blocks, each erase block having T physical pages;

30 Fig. 6A is a simplified flowchart illustration of a method for writing on a flash memory device of uneven quality, using the system of Fig. 1A for distributing logical pages evenly within the flash device memory, wherein the demux/mux and interleaver

of Fig. 1A are each operative in accordance with certain embodiments of the present invention;

Fig. 6B is a simplified flowchart illustration of a method for writing on a flash memory device of uneven quality, using the system of Fig. 1A for distributing logical pages evenly within the flash device memory, wherein the demux/mux of Fig. 1A is
5 operative in accordance with certain embodiments of the present invention;

Fig. 6C is a simplified flowchart illustration of a method for writing on a flash memory device of uneven quality, using the system of Fig. 1A for distributing logical pages evenly within the flash device memory, wherein the interleaver of Fig. 1A is
10 operative in accordance with certain embodiments of the present invention;

Fig. 7 is a simplified flowchart illustration of a method for reading from a flash memory device of uneven quality, in which the system of Fig. 1A has distributed logical pages evenly, in accordance with certain embodiments of the present invention;

Figs. 8 – 12 are simplified flowchart illustrations and tables useful in
15 understanding certain embodiments of the present invention; and

Fig. 13 is a simplified flowchart illustration of a method useful in handling short I/O requests, in accordance with certain embodiments of the present invention.

DETAILED DESCRIPTION OF CERTAIN EMBODIMENTS

Today's Flash memory devices store information as charge in either a floating gate transistor or an NROM transistor. Multi-Level Cells (MLC) store several bits by setting the amount of charge in the cell. The amount of charge is then measured by a detector, e.g. by a threshold voltage of the cell's transistor gate. Due to inaccuracies during the multiple erase and programming procedures and because of charge loss due to time and temperature (also known as retention), the measured levels suffer from a random distortion. As a result, some errors appear when the device is read. This reliability is then measured by the error probability or by a bit error rate (BER).

Flash devices are organized into pages. Each page contains a section allocated for data (512bytes-4Kbytes) and a small amount of bytes (16-32 bytes for every 512 data bytes) allocated for redundancy and back pointers. The redundancy bytes are used to store error correcting information, for correcting errors which may have occurred e.g. during the page Read. Each Read and Program operation is performed on an entire page. A number of pages are grouped together to form an Erase Block (erase block). A page cannot be erased unless the entire erase block which contains it is erased.

An important measure of Flash device quality is the number of times (N_p) the device is guaranteed to be able to be reprogrammed before irrecoverable errors occur i.e. before the allocation of 16-32 bytes of redundancy per 512 bytes of data bytes becomes insufficient to correct errors. The higher the number of program-erase cycles, the higher the bit error rate (BER). Thus, today's multi-level cell devices can perform around $N_p=1000$ cycles or less before the allocation of 16-32 bytes of redundancy per 512 bytes of data bytes becomes insufficient to correct errors.

In some types of Flash memory, the bit error rate varies across pages within the erase block. In practically all types of Flash device the bit error rate (BER) varies between individual devices. Thus, in order to assure a given error performance or a given number of program-erase cycles, it is conventional to compute the N_p (number of program-erase cycles) to accommodate the worst case device and the worst case page within each erase block.

One application of Flash devices is solid state disks (SSD) where an array of Flash devices is used as the storage media of a computer hard drive, thus, enjoying the fast Read and Access times of Flash chips. In a solid-state disk (SSD), several Flash

chips are programmed and read simultaneously to increase the Read and program speeds. For this purpose, SSDs are arranged into multiple "Banks" that allow parallel read/write operation, each Bank typically comprising a number of Flash chips. The Read/Program performance is multiplied by the number of "Banks", compared with a single Flash chip. Each "Bank" may comprise multiple chips and each chip may comprise multiple erase sectors. Each solid-state disk may, for example, comprise 4 or 8 banks. Each erase sector or erase block may for example comprise 128 pages.

The incoming pages are coded and de-multiplexed into multiple "banks". This increases the programming speed. The error probability across multiple devices or across an erase block is not uniform and causes non-uniform error probability among devices and EBs. Simple allocation of entire logical pages to banks, as shown in prior art Fig. 2, reduces the reliability to the worst case "bank" and worst case page within the erase block. The term "spare" in Fig. 2 refers to the fact that typically, each page in NAND Flash is made up of data, say 2K bytes, plus "spare": additional spare bytes, e.g. 128 bytes, used for storing error correction information and management data also termed herein meta-data.

Advanced wear leveling is used to control error probability variances due to different program-erase cycles by mapping "logical" pages ("logical" pages contain a sequence of data sent/read by the host) into "physical" pages ("physical" pages contain the actual data being programmed onto the Flash devices, including, on top of the data read/written by the host, error correction information and back pointers used for recognizing the true address of a page) in a way that evens-out the program-erase cycles over all Erase Blocks, thus attempting to even-out error probability across the data blocks. But even such methods cannot guarantee even error probability due to the variance across devices and across pages in an erase block. The EBs and devices are physically different and the wear leveling does not address these differences.

A top-level block diagram of a system for distributing logical pages evenly within flash device memory of uneven quality, according to one embodiment of the invention, is described with reference to Fig. 1A. Embodiments of the various components of the system of Fig. 1A (or Fig. 1B) are next described with reference to the logical page distribution schemes of Figs. 3 – 4 as compared to the logical page distribution scheme of prior art Fig. 2. One possible structure for flash memory is described in Fig. 5. Three alternative methods for writing on that structure according to

certain embodiments of the present invention, using the system of Fig. 1A (or Fig. 1B), are presented in Figs. 6A – 6C. A method for reading from that structure according to certain embodiments of the present invention, using the system of Fig. 1A (or Fig. 1B), is presented in Fig. 7.

5 Referring to Fig. 1A, according to certain embodiments of the invention, the system comprises an Encoder 10, Page Demux/Mux 20, physical page mapper+Interleaver 30, a “free” page/erase block mapper and locator 40 and a Flash array 50 which may comprise several banks of commercially available flash memory devices. A Host 60, which may for example comprise a personal computer having an
10 operating system and application programs such as Word or a digital entertainment device, writes to the flash array 50, through a Flash controller such as the controllers distributed by STMicroelectronics and SMSC. The host subsequently reads from the flash array 50 a sequence of Logical pages. For example, in USB drives, the Host 60 may send an address to be read from or written to, per each pagelet of 512 bytes. These
15 may be aggregated into pages by the NAND controller.

Alternatively, functional units of Fig. 1A may be implemented within a Flash uController which controls the NAND flash memories, as shown in Fig. 1B. The uController 100 may be based on commercially available uController systems such as those distributed by Samsung, Eureka Technology, Denali, Arasan and Barco. The
20 uController 100 translates read/write commands generated by the host through a standard interface, such as but not limited to a Universal Serial Bus (USB), Secure Disk (SD) interface, to NAND Flash read/program page commands and Erase Block commands. The uController 100 may be regarded as implementing a level of virtualization, concealing the fact that the information is being written and read from a
25 NAND flash e.g. by concealing the fact that data must be written to pages of a certain size and that a page may not be rewritten unless the entire erase block that contains it is erased.

Certain embodiments of the Encoder 10, Decoder 70, Pagelet distribution unit 30, Distributed pagelet collation unit 60 and the page multiplexer 30 are each described
30 in detail below. However generally, these functional units may be added to (or integrated with) the conventional functional units occurring in a conventional NAND Flash controller 80 so as to provide an additional layer of virtualization in accordance with certain embodiments of the present invention. The additional virtual layer presents

-21-

to the Flash controller a NAND flash with larger pages than the actual physical pages residing within the actual NAND Flash devices; these larger pages are referred to herein as logical pages. For example, each logical page may be 4 times the size of the physical pages. The encoder 10 and decoder 70 each work on pages of the size stored in the flash. Units 30 and 60 mix several encoded pages together and the page multiplexer then spread several (say, 4) pages over different locations in the flash device. Various pages may be distributed within a block or across blocks or both.

The above functional units preferably interact as follows when writing occurs:

a. The logical pages are encoded by encoder 10 and appended with redundancy bytes to form Encoded Pages. The logical pages may be divided into several pagelets (such as for example 4 pagelets) the size of a physical page and each physical page may be encoded separately.

b. The Encoded pages are broken down into pagelets by page multiplexer 20 and then grouped together with pagelets of other encoded pages to form physical pages.

c. "Page Multiplexer/Interleaver" 20 interleaves the pages between several devices, several erase blocks, and/or several locations within an erase block. A particular advantage of certain embodiments of the present invention is that pagelets (portions of pages) assigned to several banks, can be written to, and read from, the various banks simultaneously to increase write performance and read performance respectively.

d. 'free' page/erase block mapper 40 translates a logical address sent by the host into a "physical" address, the actual Erase Block address on each bank of the array. A "free" page is one which has not been programmed hence is in an erased state. The "page multiplexer" 20 then addresses pages addressed with respect to the beginning of each Erase Block.

A read operation typically comprises the same steps in reverse order i.e. first (d) then (c) then (b) then (a).

Certain embodiments of the Pagelet distributor 30 of Fig. 1A are now described. Typically, this component takes J "Encoded" pages and breaks down each page into J pagelets. The pagelets may then be grouped together into physical pages such that each physical page contains J sub-pagelets, each sub-pagelet from a different "Encoded" page. It is appreciated that throughout the present specification, a sub-pagelet may be a pagelet.

It is appreciated that J may be any suitable integer such as 4. J may be chosen to be sufficiently large as to average the BER over several pages and to obtain approximately the same bit error rates for all “logical addresses” (those defined by the NAND controller).

5 The following notation is used herein to identify pagelets: Epaglet_{ij} = the j'th sub-pagelet of the i'th encoded page. Ppagelet_{ij} denotes the j'th sub-pagelet of the i'th physical page.

Any of the following schemes may be pre-selected for ordering the pagelets within the “physical” pages e.g. as follows (i=1..J, j=1..J):

10 Scheme a: Ppagelet_{ij} = Epaglet_{ji}. For example, if a sequence of 4 encoded pages is provided, each encoded page may be broken down to 4 Epaglets. Then, Epaglet 1 of each of the encoded pages 1, 2, 3 and 4 are all grouped together to construct physical page #1. Similarly, physical pages 2, 3, 4 may be constructed with Epaglets 2, 3, and 4 respectively.

15 Scheme b: Ppagelet_{ij} = Epaglet_{[(i+j-2 modulo J) + 1]j}. For example, given a sequence of 4 encoded pages, each encoded page may be broken down to 4 Epaglets. Then Epaglets 1,2,3 and 4 of encoded pages 1 2 3 and 4 respectively are grouped together, in order, to form Ppagelet 1. Epaglets 2,3,4 and 1 of encoded pages 1 2 3 and 4 respectively are grouped together, in order, to form Ppagelet 2. Epaglets 3,4,1 and 2
20 of encoded pages 1 2 3 and 4 respectively are grouped together, in order, to form Ppagelet 3 and Epaglets 4,1,2 and 3 of encoded pages 1 2 3 and 4 respectively are grouped together, in order, to form Ppagelet 4.

More generally, a permutation function Prm(k) may be defined from the set {0,1,...,J*J-1} to {0,1,...,J*J-1} and then assign Ppagelet_{ij} may be Epaglet_{[Prm(i-1+(j-1)*J) modulo J + 1] [Prm((i-1)+(j-1)*J) div J + 1]}.

25 A particular advantage of using a permutation function as above, rather than always sequencing pagelets in the following order: 1, 2, 3, ...J and always sequencing logical pages in the following order: 1, 2, 3, ...J e.g. as shown in Figs. 3 - 4, is that even if there is a regularity in the variation of flash memory quality, e.g. such that the initial portions of each erase block tend to be of poor quality, or such that the last erase block in each device tends to be of poor quality, the logical data is generally evenly distributed over the high-quality flash memory portions and the low-quality flash memory portions. If there is no regularity in the variation of flash memory quality, i.e. if

the variation is generally random, then even without permutation, the logical data is evenly distributed over the randomly occurring high-quality low-quality flash memory portions.

Fig. 3 illustrates an interleaving scheme which breaks down the “encoded” pages into several e.g. 4 pagelets and programs them across typically the same number e.g. 4 memory banks according to the above scheme (b). Fig. 4 illustrates an interleaving scheme which breaks down the “encoded” pages into 4 (say) pagelets and programs them across an erase block in a single memory bank according to the above scheme (b).

The above discussion assumes that the physical pages are the same size as the encoded pages. In some applications, the physical pages are larger, or smaller, than the encoded pages such that the physical page contains P pagelets. Thus J encoded Pages are mapped into $J \cdot J/P$ physical pages. For example, given a sequence of 4 encoded pages (one “logical” page), each encoded page may be broken down to $J=4$ Epagelets and each physical page may contain $P=8$ Ppagelets (the same size of an Epagelet), in which case Epagelets 1 and 2 of encoded pages 1, 2, 3 and 4 respectively are grouped together to form physical page 1, whereas Epagelets 2 and 3 of encoded pages 5, 6, 7 and 8 are grouped together to form physical page 2. Thus, the 4 encoded pages (one logical page) are stored over 2 physical pages.

Fig. 8 illustrates a method for partitioning encoded pages into pagelets and then spreading the pagelets over physical pages (each capable of storing P pagelets). The method supports two possible options for ordering the Epagelets into the physical pages: according to a predefined permutation or according to a formula.

Certain embodiments of the Physical Page interleaver 30 of Fig. 1A are now described.

Physical Page interleaver 30 obtains J “physical-pagefuls of data and maps them into the flash array. If the Flash array contains L banks, an L -fold increase in performance may be obtained if each physical-page-ful of data is written into all L banks simultaneously. The mapping of the “physical-page-fuls may be done using several strategies such as but not limited to the following:

Strategy a: Each physical-page-ful of data is written to a different bank, so as to effectively average bit error rate (BER) over banks. In this embodiment, physical-page-ful j may be programmed into physical page $j-1 \bmod L + 1$ of bank whose serial

-24-

number is: $\text{floor}((j-1)/L)+1$, where $\text{Floor}(x)$ is the largest integer which is still smaller or equal to x . For example, consider the case of a sequence of 8 “physical” pages and an array with 4 banks. Pages 1 and 5 in the sequence are inserted to the first bank. Pages 2 and 6 are inserted to the second bank and pages 3,7 and 4,8 are inserted to the 3rd and 4th banks respectively.

Strategy b: Each physical-page-ful of data is written to a different physical page within an Erase Block (within a device) so as to effectively average bit error rate (BER) over pages. In this embodiment, if the erase block contains M pages then “physical” page j is programmed into page $(j-1)*S + S_0$ modulo $M + 1$ with respect to the beginning of the erase block where S and S_0 are the step size and initial step respectively (in pages). For example, consider an array with only one bank, and assume the Erase blocks contain 16 physical pages, the step size is 5 physical pages and the initial step is 1 physical page. Furthermore, assume that there is a sequence of 8 “physical” pages to be programmed within one “logical page” ($j=8$ here by the host). Pages 1,2,3,4 in the sequence are programmed respectively to pages 1, 6, 11, 16 in the erase block. Pages 5, 6, 7, 8 in the sequence are programmed to pages 5, 10, 15, 3 respectively, in the erase block.

Strategy c: Each physical-page-ful of data is written to a different bank and to a differently positioned page within each bank so as to average the bit error rate (BER) both over devices and over physical pages within an erase block. Strategies a and b are orthogonal hence may be combined. For example, consider an array with 4 banks, each including an erase block with 16 pages, a step size of 4 physical pages, an initial step of 1 physical page, and a sequence of 4 “physical” pages to be programmed in each I/O operation. Page 1 in the sequence are programmed to page 1 of the first bank. Page 2 in the sequence are programmed to page 5 in Bank 2. Pages 3, 4 in the sequence are programmed to page 9 in bank 3, and page 13 in bank 4 respectively.

Certain embodiments of the Free page/erase block mapper and locator 40 of Fig. 1B are now described. A conventional NAND flash controller 80 typically comprises a free page/erase block mapper and locator 40 as part of its memory management unit. This conventional functional unit, 40, within a conventional NAND flash controller 80 allocates empty Erase blocks to be programmed or finds the appropriate erase blocks on the array which contain the appropriate addresses to be read. Typically, unit 40

translates between virtual addresses generated by the controller which "thinks" that the NAND flash has larger pages than it actually does.

Several methods are known in the art for allocating erase blocks. According to one such method, a predefined address computation unit is used. In this case there is a one to one mapping between the "logical" address and the actual address of the "physical" page on the device. If a certain page has been previously programmed due to a previous programming of the same "logical" page, the entire Erase block is erased and reprogrammed with the new data. As an example, assume that a sequence of 4 "logical" pages are to be programmed to addresses 33, 34, 35 and 36. The array contains 4 banks, each bank containing devices with erase block size of 16 pages. A possible one to one mapping may be: $\text{erase block\#1} = \text{floor}(33/4/16) + 1$. If pages 65, 66, 67 and 68 are to be programmed, all physical pages may be mapped to erase block #2 of all devices. The Interleaver may then rename the addresses 1,2,3 and 4 relative to the beginning of erase block #1.

Alternatively, the mapper may contain a lookup block table as shown in Fig. 9 where each Block in the Flash memory is intended to contain a consecutive set of virtual addresses, say 256 addresses. These sets of addresses need not be ordered within the Flash and may be programmed to non consecutive Blocks. The table also contains a page usage vector which determines which page is being used and which is not.

In order to save time, a block which was previously used may not be actually erased but instead merely marked as erased, whereas in fact, a different Erase Block is used in its stead. In this case the mapper finds a free erase block for each bank in the array and the pages are programmed to that free erase block. One way of implementing this embodiment is to maintain a table of all free EBs currently available and to update the table each time an erase block has been used or erased, e.g. as shown in Figs. 9 - 10.

Certain embodiments of the flash array 50 of Fig. 1A are now described. As shown in Fig. 5, flash array 50, according to an embodiment of the present invention, comprises L banks of Flash memory devices. Each Bank comprises M Flash memory devices. Each memory device contains K Erase Blocks and each Erase Block contains T pages. At each bank, the EBs are numbered serially from 1 to MK. In each erase block, the pages are numbered from 1 to T, resetting at the beginning of each erase block.

Optionally, special treatment is allotted to short Input/output Sequences. This is because, in order to fully benefit from certain embodiments shown and described

-26-

herein, all input/output sequences may need to contain at least as many “logical” pages as the number of banks (i.e. $J=L$). The NAND Flash controller 80 may append zeros to sequences which were sent by the host and which are shorter than the logical page. Thus, a memory device may be selected whose number of banks equals the number of pages in a typical I/O transaction in a given application. Even if a particular application is characterized by a central tendency toward a transaction (input/output sequence) of a particular length, some transactions may be shorter. If these transactions are padded with blank pages to obtain a total of J pages, this wastes both time and space during the sequence programming. In order to enhance handling of short I/O requests, which are recognized as such by the flash device's NAND controller 80, the method of Fig. 13 may be performed. The method of Fig. 13 may include some or all of the following steps, suitably ordered e.g. as shown:

Step 1310: Set an adjustable J in the Page Demux/Mux, Physical Page Mapper and interleaver blocks, to accommodate shorter transactions. Each time a transaction is processed, J is adjusted accordingly, depending on the size of the sequence.

Step 1320: Rank the pages in the Flash array according to their bit error rate (BER), either at the bank-level, or at the device-level, or at the erase-block level, or even within each erase-block.

Step 1330: Program only the J highest ranking pages e.g. according to the ranking in step 1320, e.g. using the interleaving methods defined earlier.

Another option is to select a subset of pages such that their average bit error rate (BER) is above or equal to that of the average bit error rate (BER), averaged over typical number of pages.

A particular advantage of handling short I/O sequences as described above with reference to Fig. 13, is that this ensures that the reliability of the shorter transactions is not determined by the bit error rates of the worst pages, but rather by bit error rates of “above average” pages or even best available pages. However, since short I/O sequences are uncommon in many applications, the overall reliability may hardly be affected by their occurrence even without special handling of short I/O sequences.

Special handling of short sequences e.g. as described above may involve modification of the control of the NAND Flash interface to enable both short “logical pages” and long ones to be sent. Furthermore, the memory management typically is modified to map pages of the size of the physical pages and not of the “logical” pages

-27-

such that the NAND controller is able to take note of whether these pages were already written to the Flash device or not. In addition the NAND controller mapper is typically able to distinguish between pages that were programmed using a short sequence and those programmed using a long sequence, .e.g. using a suitable table such as that shown in Fig. 11. In the table of Fig. 11, the page usage vector pertains to physical page size and not logical pages. Also stored is a short/long vector which identifies pages that were programmed as part of short sequences as well as pages programmed as part of long sequences. This is typically useful when the information is read from the device.

Fig. 12 illustrates a method for distributing pagelets into physical pages when the long sequences contain J encoded pages and the short ones contain J/2 pages. The pages are distributed into addresses as before, e.g. in the page multiplexer, by translating the address given by the controller as follows:

Physical Block Address = floor(Logical page address / M)

i = Logical Address within a block = Logical page address mod M

Physical Page Address within a page =

$$\text{floor}(i/J) + \begin{cases} \frac{2M}{J}i & i = 0 \dots J/2 - 1 \\ \frac{2M}{J}(i - J/2) + \frac{M}{J} & i = J/2 \dots J - 1 \end{cases}$$

Fig. 6A is a simplified flowchart illustration of a method for writing on the flash memory device of Fig. 1A, in accordance with certain embodiments of the present invention. In step 610, host 60 requests i/o operation: "write J logical pages into memory device whose physical pages are of size KJ". In step 620, encoder 10 encodes the J logical pages including appending redundancy bits to each and sends J encoded logical pages to mux/demux 20. In step 630, mux/demux 20 partitions each of the J encoded logical pages into J pagelets, and groups pagelets from different logical pages together using a predetermined pagelet grouping scheme (with or without predetermined permutation) to form multi-origin physical page-fuls each containing pagelets originating from more than one logical page.

In step 640, interleaver 30 maps the multi-origin physical page-fuls into flash device array 50 using free pages identified by mapper/locator 40 and stores physical page location information conventionally to facilitate eventual accessing of each physical page-ful responsive to future read requests. Selection of page to map each

page-ful to, from among free pages identified by unit 40, is as per pre-determined distribution scheme e.g. distribution of adjacent multi-origin physical page-fuls over different erase blocks, over different devices, over different banks, or any combination thereof; and/or even distribution of adjacent multi-origin physical page-fuls to the various pages of an individual erase block, or to the various erase blocks of an individual device, or to the various devices of an individual bank, or any combination thereof. In step 650, writing circuitry of flash device array 50 writes multi-origin physical page-fuls into flash device array 50 in accordance with mapping determined by interleaver 30.

A particular advantage of certain embodiments of the method of Fig. 6A is that individual logical pages are less closely associated with specific low-quality local portions of the flash memory.

Fig. 6B is a simplified flowchart illustration of a method for writing on a flash memory device of uneven quality, using the system of Fig. 1A for distributing logical pages evenly within the flash device memory, wherein the demux/mux of Fig. 1A is operative in accordance with certain embodiments of the present invention. Fig. 6B is similar to Fig. 6A except that conventional step 660 replaces step 640 of Fig. 6A which is an inventive step. Fig. 6C is a simplified flowchart illustration of a method for writing on a flash memory device of uneven quality, using the system of Fig. 1A for distributing logical pages evenly within the flash device memory, wherein the interleaver of Fig. 1A is operative in accordance with certain embodiments of the present invention. Fig. 6C is similar to Fig. 6A except that step 670 replaces step 630 of Fig. 6A.

Fig. 7 is a simplified flowchart illustration of a method for reading from the flash memory device of Fig. 1A, in accordance with certain embodiments of the present invention. In step 710, host 60 requests i/o operation: "read J logical pages from memory device" (physical pages of memory device are of size KJ). In step 720, interleaver 30 identifies the physical pages in the flash device array 50 on which reside the multi-origin physical page-fuls corresponding to the J logical pages, using physical page location information stored in step 640 of Fig. 6. In step 730, reading circuitry of flash device array 50 accesses the physical pages identified by interleaver 30 and provides them to the mux/demux 20.

-29-

In step 740, mux/demux 20 partitions each of the physical pages into pagelets, and groups pagelets from each individual logical page together based on the predetermined pagelet grouping scheme (with or without predetermined permutation) which was used in Fig. 6, step 630. In step 750, encoder/decoder 10 decodes the J logical pages using redundancy bits appended to each and sends J decoded logical pages to host 60.

Included in the scope of the present invention, inter alia, are electromagnetic signals carrying computer-readable instructions for performing any or all of the steps of any of the methods shown and described herein, in any suitable order; machine-readable instructions for performing any or all of the steps of any of the methods shown and described herein, in any suitable order; program storage devices readable by machine, tangibly embodying a program of instructions executable by the machine to perform any or all of the steps of any of the methods shown and described herein, in any suitable order; a computer program product comprising a computer useable medium having computer readable program code having embodied therein, and/or including computer readable program code for performing, any or all of the steps of any of the methods shown and described herein, in any suitable order; any technical effects brought about by any or all of the steps of any of the methods shown and described herein, when performed in any suitable order; any suitable apparatus or device or combination of such, programmed to perform, alone or in combination, any or all of the steps of any of the methods shown and described herein, in any suitable order; information storage devices or physical records, such as disks or hard drives, causing a computer or other device to be configured so as to carry out any or all of the steps of any of the methods shown and described herein, in any suitable order; a program pre-stored e.g. in memory or on an information network such as the Internet, before or after being downloaded, which embodies any or all of the steps of any of the methods shown and described herein, in any suitable order, and the method of uploading or downloading such, and a system including server/s and/or client/s for using such; and hardware which performs any or all of the steps of any of the methods shown and described herein, in any suitable order, either alone or in conjunction with software.

Certain operations are described herein as occurring in the microcontroller internal to a flash memory device. Such description is intended to include operations which may be performed by hardware which may be associated with the

-30-

microcontroller such as peripheral hardware on a chip on which the microcontroller may reside. It is also appreciated that some or all of these operations, in any embodiment, may alternatively be performed by the external, host-flash memory device interface controller including operations which may be performed by hardware which may be associated with the interface controller such as peripheral hardware on a chip on which the interface controller may reside. Finally it is appreciated that the internal and external controllers may each physically reside on a single hardware device, or alternatively on several operatively associated hardware devices.

Any data described as being stored at a specific location in memory may alternatively be stored elsewhere, in conjunction with an indication of the location in memory with which the data is associated. For example, instead of storing page- or erase-sector-specific information within a specific page or erase sector, the same may be stored within the flash memory device's internal microcontroller or within a microcontroller interfacing between the flash memory device and the host, and an indication may be stored of the specific page or erase sector associated with the cells.

It is appreciated that the teachings of the present invention can, for example, be implemented by suitably modifying, or interfacing externally with, flash controlling apparatus. The flash controlling apparatus controls a flash memory array and may comprise either a controller external to the flash array or a microcontroller on-board the flash array or otherwise incorporated therewithin. Examples of flash memory arrays include Samsung's K9XXG08UXM series, Hynix's HY27UK08BGFM Series, Micron's MT29F64G08TAAWP or other arrays such as but not limited to NOR or phase change memory. Examples of controllers which are external to the flash array they control include STMicroelectronics's ST7265x microcontroller family, STMicroelectronics's ST72681 microcontroller, and SMSC's USB97C242, Traspan Technologies' TS-4811, Chipsbank CBM2090/CBM1190. Example of commercial IP software for Flash file systems are: Denali's Spectra™ NAND Flash File System, Aarsan's NAND Flash Controller IP Core and Arasan's NAND Flash File System. It is appreciated that the flash controller apparatus need not be NAND-type and can alternatively, for example, be NOR-type or phase change memory-type.

Flash controlling apparatus, whether external or internal to the controlled flash array, typically includes the following components: a Memory Management/File

-31-

system, a NAND interface (or other flash memory array interface), a Host Interface (USB, SD or other), error correction circuitry (ECC) typically comprising an Encoder and matching decoder, and a control system managing all of the above.

5 The present invention may for example interface with or modify, as per any of the embodiments described herein, one, some or all of the above components and particularly the NAND interface component and, in some embodiments, with the memory management component.

10 It is appreciated that software components of the present invention including programs and data may, if desired, be implemented in ROM (read only memory) form including CD-ROMs, EPROMs and EEPROMs, or may be stored in any other suitable computer-readable medium such as but not limited to disks of various kinds, cards of various kinds and RAMs. Components described herein as software may, alternatively, be implemented wholly or partly in hardware, if desired, using conventional techniques.

15 Features of the present invention which are described in the context of separate embodiments may also be provided in combination in a single embodiment. Conversely, features of the invention which are described for brevity in the context of a single embodiment may be provided separately or in any suitable subcombination.

CLAIMS

1. A flash memory accessory device operative in conjunction with a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of physical pages, the accessory system comprising:

logical page distributing apparatus for distributing data from at least one logical page in said stream between at least two of said multiplicity of physical pages.

2. A device according to claim 1 wherein said multiplicity of physical pages is partitioned into a plurality of erase sectors each having physical pages, wherein the physical pages within an individual erase sector can only be erased in unison, wherein said apparatus for distributing distributes data from at least one logical page in said stream between physical pages within a single erase sector.

3. A device according to claim 1 wherein said multiplicity of physical pages is partitioned into at least first and second erase sectors each having physical pages, wherein the physical pages within each erase sector can only be erased in unison, wherein said apparatus for distributing distributes information from at least one logical page in said stream between at least a first physical page within the first erase sector and at least a second physical page within the second erase sector.

4. A device according to claim 1 wherein said flash memory system comprises at least first and second flash memory banks which are accessible in parallel and wherein said apparatus for distributing distributes information from at least one logical page in said stream between at least a first physical page within the first flash memory bank and at least a second physical page within the second flash memory bank.

5. A flash memory accessory device operative in conjunction with a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of flash memory cells, wherein said multiplicity of flash memory cells can be divided into several subsets of flash memory cells each having a characteristic range of bit error rates which differs from the characteristic range of bit error rates of at least one other subset of flash memory cells,

-33-

the accessory device comprising bit error rate (BER) distributing apparatus operative to distribute data from at least one logical page in said stream between at least two of said subsets.

5 6. A device according to claim 5 wherein said multiplicity of flash memory cells can be partitioned into several orthogonal partitions of flash memory cells, each partition defining several subsets into which the multiplicity of flash memory cells is divided, each subset having a characteristic range of bit error rates which differs from the characteristic range of bit error rates of at least one other subset of flash memory
10 cells defined by the same partition, wherein said bit error rate (BER) distributing apparatus is operative to distribute data from at least one logical page in said stream between at least two subsets defined by each of at least two of the partitions.

7. A device according to claim 6 wherein one of said orthogonal partitions
15 comprises partitioning of the multiplicity of flash memory cells into physical pages.

8. A device according to claim 6 wherein said physical pages include at least first and second layers and wherein said multiplicity of flash memory cells comprises a multiplicity of multi-level cells each storing at least two bits including a most
20 significant bit (MSB) and a least significant bit (LSB) thereby to define a multiplicity of MSBs and a multiplicity of LSBs and wherein said apparatus for distributing is also operative to distribute at least a first MSB from among said multiplicity of MSBs and a first LSB from among said multiplicity of LSBs to the first layer and to distribute at least a second MSB from among said multiplicity of MSBs and a second LSB from
25 among said multiplicity of LSBs to the second layer.

9. A device according to claim 6 wherein one of said orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into erase sectors.

30 10. A device according to claim 6 wherein one of said orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into flash memory devices.

-34-

11. A device according to claim 6 wherein one of said orthogonal partitions comprises partitioning of the multiplicity of flash memory cells into flash memory banks.

5 12. A device according to claim 6 wherein said apparatus for distributing is operative to distribute data from at least one logical page in said stream between at least two subsets defined by each of the partitions.

10 13. A device according to claim 1 wherein said logical page distributing apparatus is operative to distribute at least a pagelet within at least a logical page onto a physical page determined by computing a predetermined function of a logical address received from a controller

15 14. A device according to claim 4 wherein said first and second physical pages are read from simultaneously when it is desired to access therefrom, a logical page stored therewithin.

20 15. A device according to claim 4 wherein said first and second physical pages are written to simultaneously when it is desired to store a logical page therewithin.

25 16. A device according to claim 1 wherein said page-distributing apparatus comprises apparatus for partitioning, for at least one set of logical pages, each of several logical pages in said set into pagelets and apparatus for grouping pagelets from said several logical pages together using a predetermined pagelet grouping scheme to form a multi-origin physical page-ful of data containing pagelets originating from more than one logical page.

30 17. A device according to claim 16 wherein said pagelet grouping scheme comprises a pagelet permutation scheme.

18. A device according to claim 17 wherein said at least one set of logical pages comprises a plurality of sets of logical pages and wherein said pagelet permutation

-35-

scheme is characterized in that not all of said plurality of sets of logical pages are permuted in the same way.

19. An accessory device according to claim 5 and also comprising
5 logical page distributing apparatus for distributing data from at least one logical page in said stream between at least two of said multiplicity of physical pages, and wherein said bit error rate (BER) distributing apparatus is operative to distribute said two physical pages in two of said subsets respectively.

10 20. A system for storing a plurality of logical pages in a set of at least one flash device, each flash device including a set of at least one erase block, the system comprising:

apparatus for distributing at least one of said plurality of logical pages over substantially all of said erase blocks in substantially all of said flash devices, thereby to define, for at least one logical page, a sequence of pagelets thereof together including all
15 information on the logical page and each being stored within a different erase block in said set of erase blocks; and

apparatus for reading each individual page from among said plurality of logical pages including apparatus for calling and ordering said sequence of pagelets from different erase blocks in said set of erase blocks.

20 21. Apparatus according to claim 1 wherein said set of flash devices includes a set of at least one bank of flash devices, each bank in said set of banks including at least one flash device.

25 22. A method for facilitating interaction between a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of physical pages, the method comprising:

distributing data from at least one logical page in said stream between at least two of said multiplicity of physical pages.

30 23. A method for facilitating interaction between a data source supplying a stream of logical pages including data and a flash memory system including a multiplicity of flash memory cells, wherein said multiplicity of flash memory cells can be divided into

-36-

several subsets of flash memory cells each having a characteristic range of bit error rates which differs from the characteristic range of bit error rates of at least one other subset of flash memory cells,

the method comprising distributing data from at least one logical page in said
5 stream between at least two of said subsets.

24. A method for storing a plurality of logical pages in a set of at least one flash device, each flash device including a set of at least one erase blocks, the method comprising:

distributing at least one of said plurality of logical pages over substantially all of
10 said erase blocks in substantially all of said flash devices; thereby to define, for at least one logical page, a sequence of pagelets thereof together including all information on the logical page and each being stored within a different erase block in said set of erase blocks; and

reading each individual page from among said plurality of logical pages
15 including apparatus for calling and ordering said sequence of pagelets from different erase blocks in said set of erase blocks.

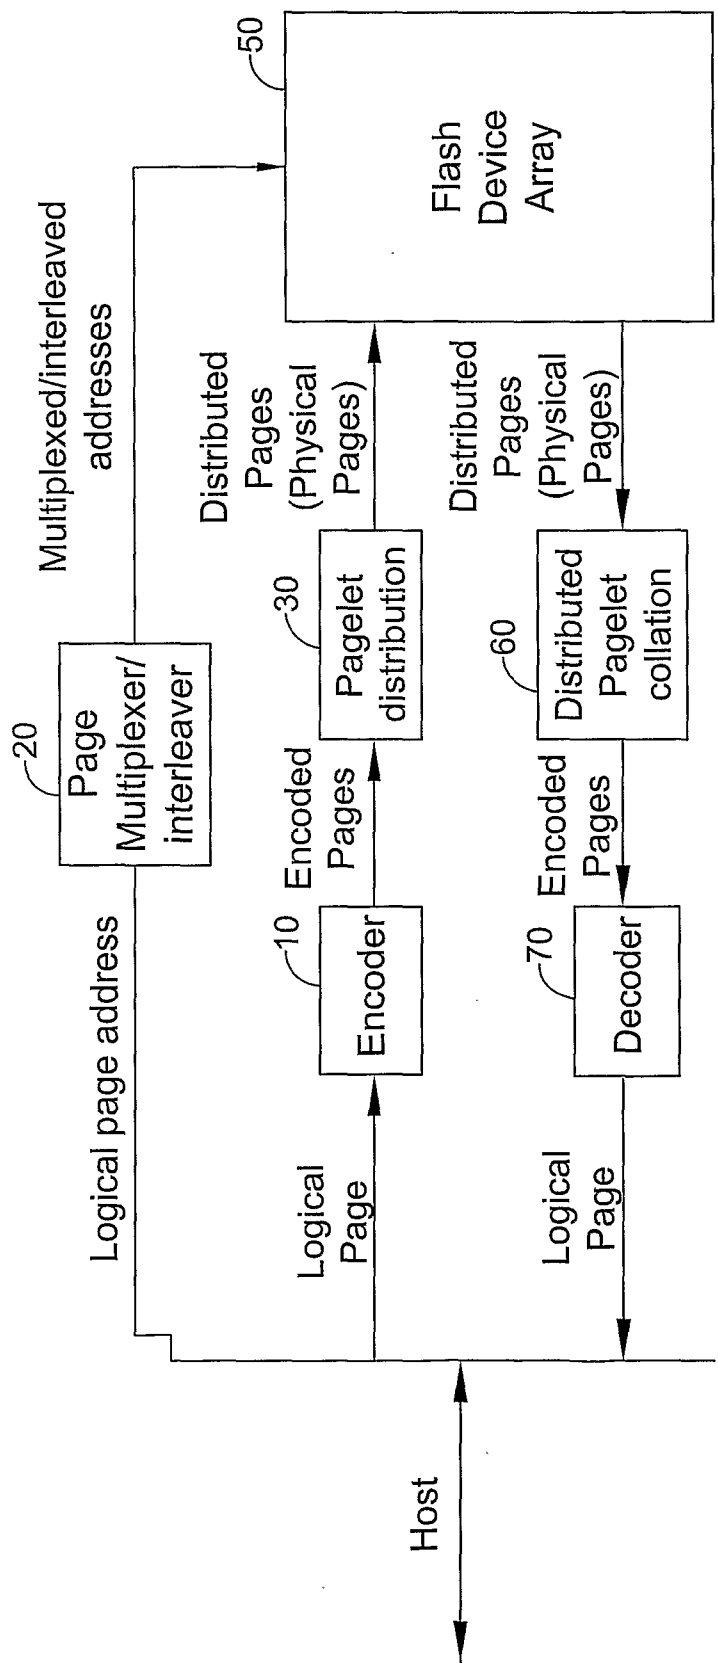


Figure 1a

2/14

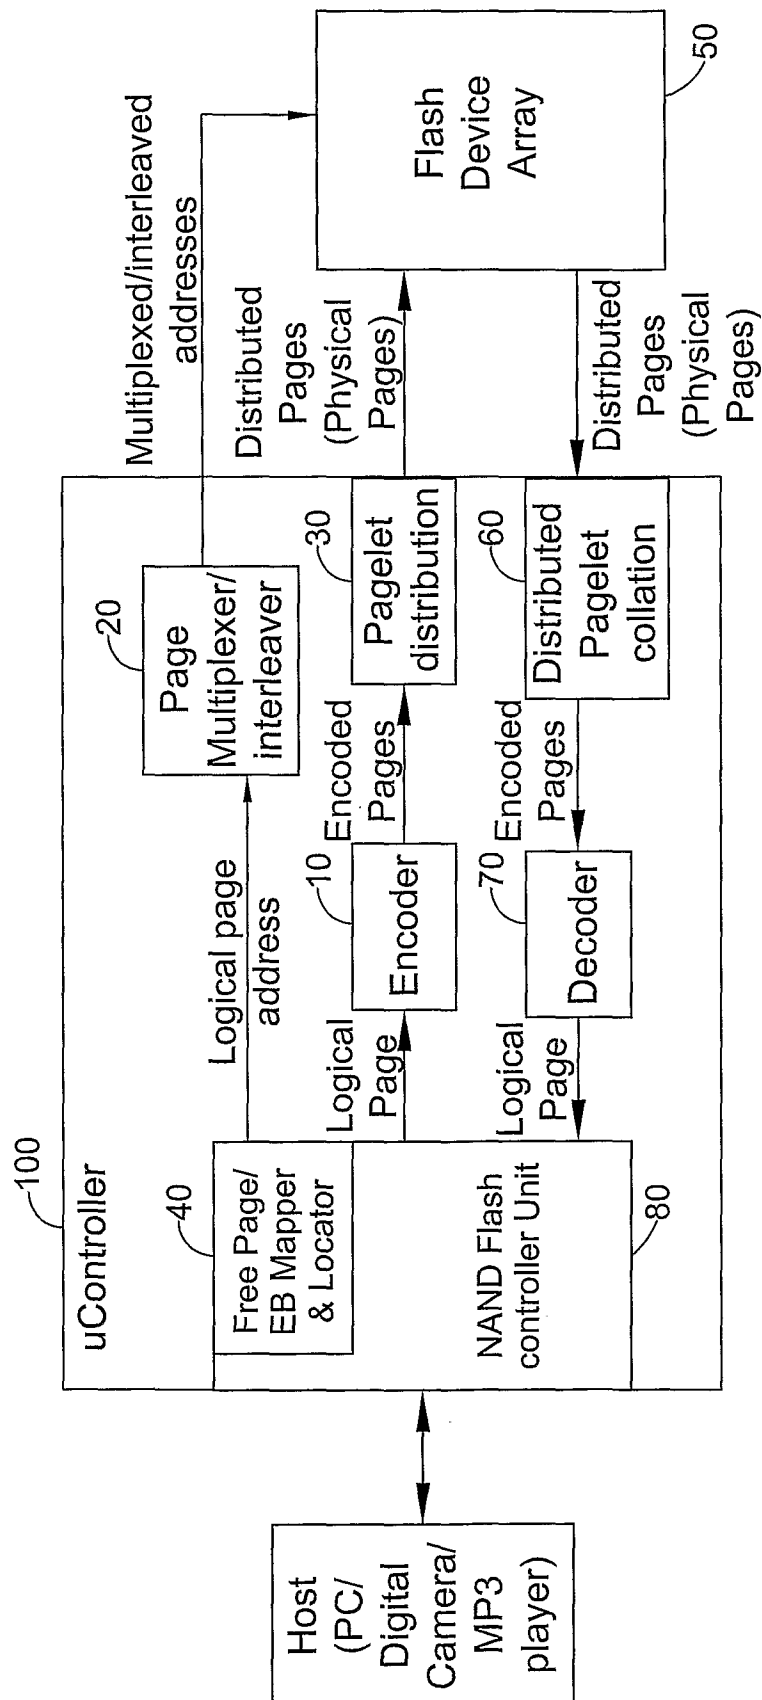


Figure 1b

3/14

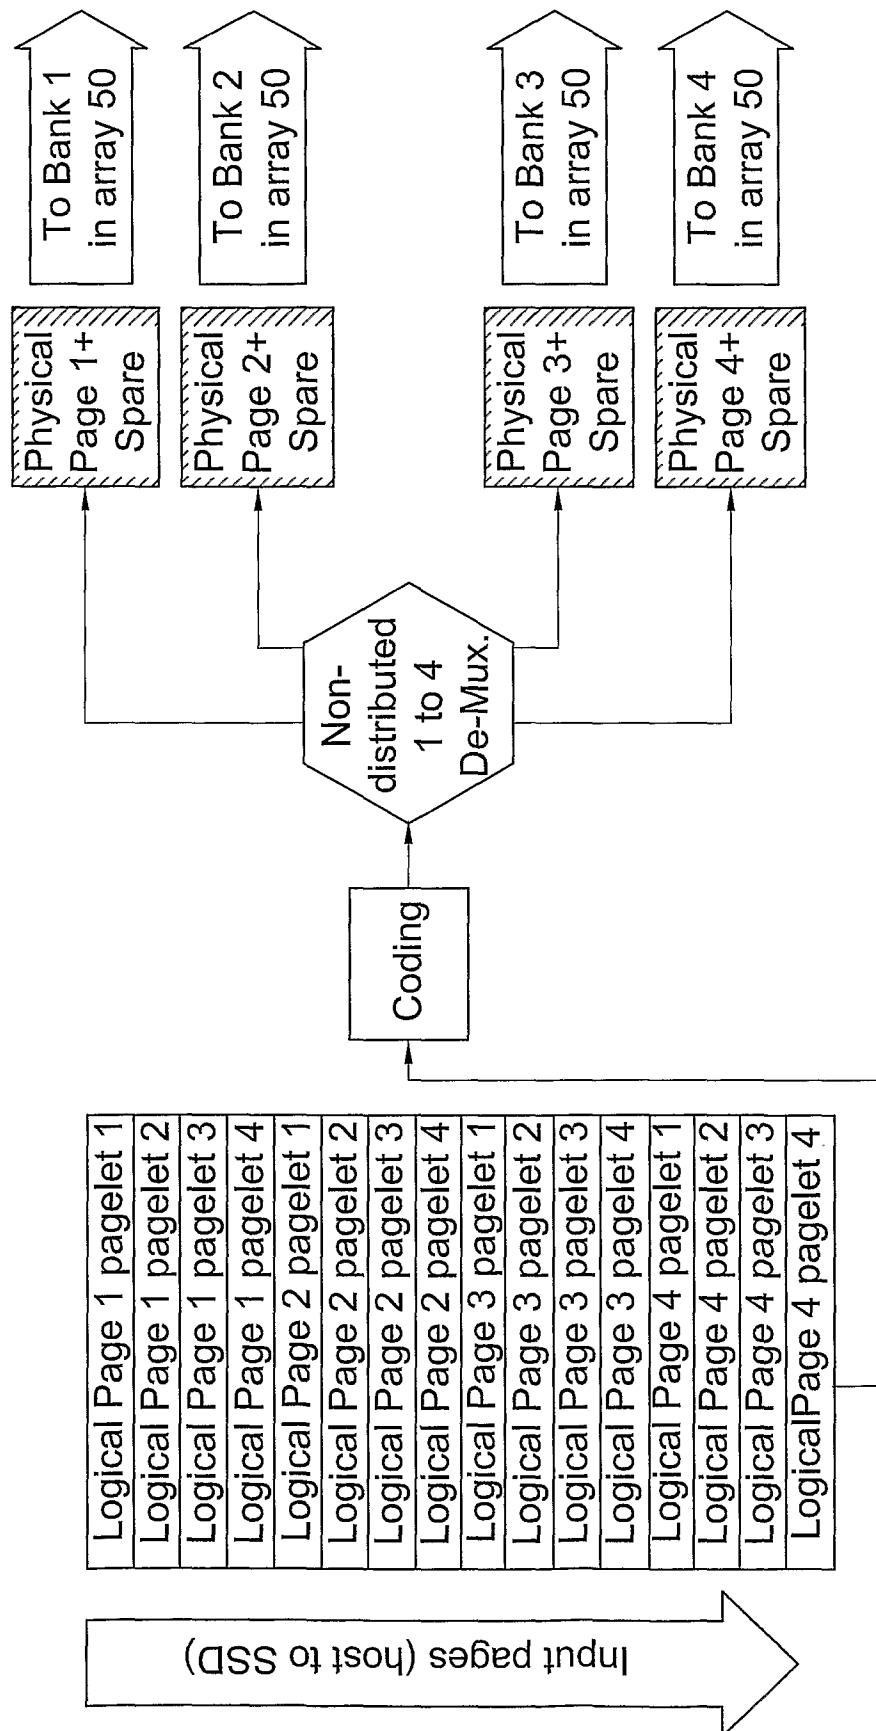


Figure 2 (PRIOR ART)

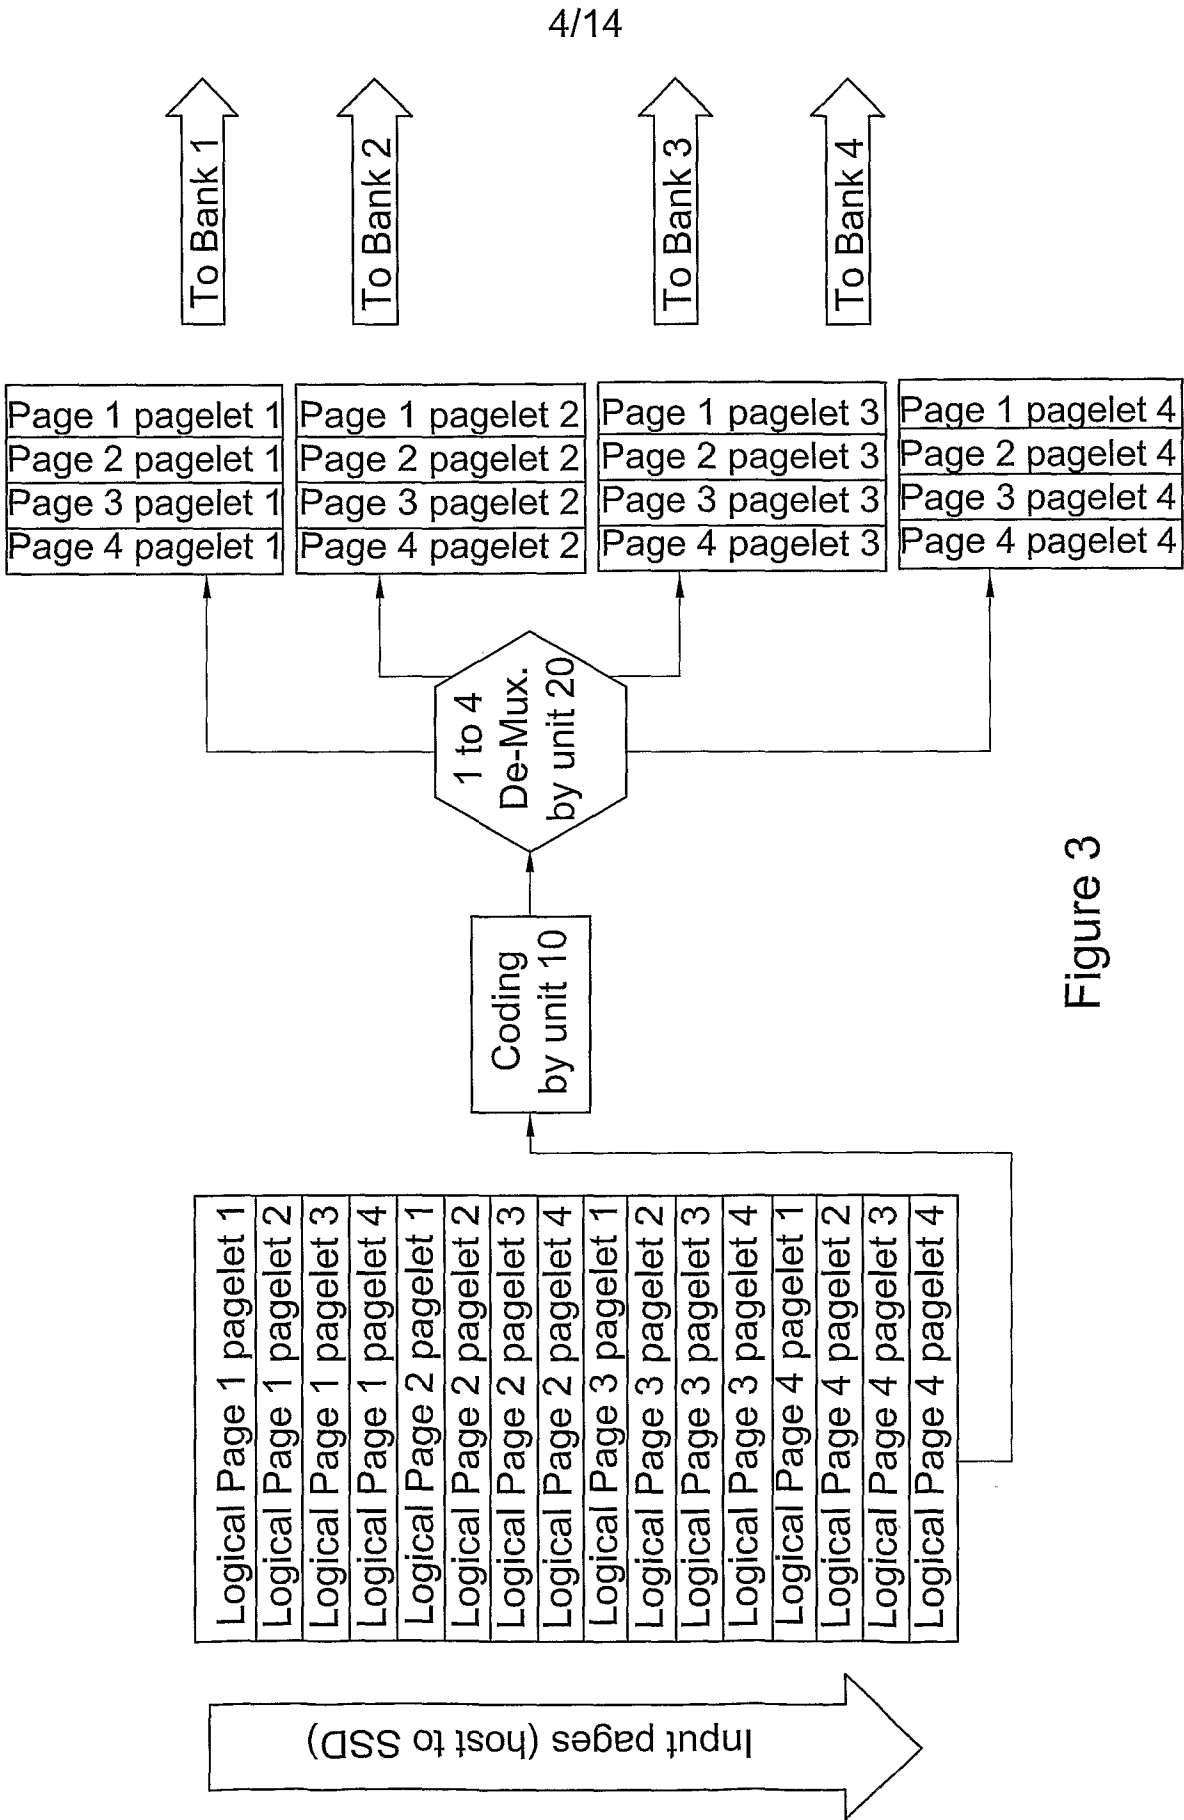


Figure 3

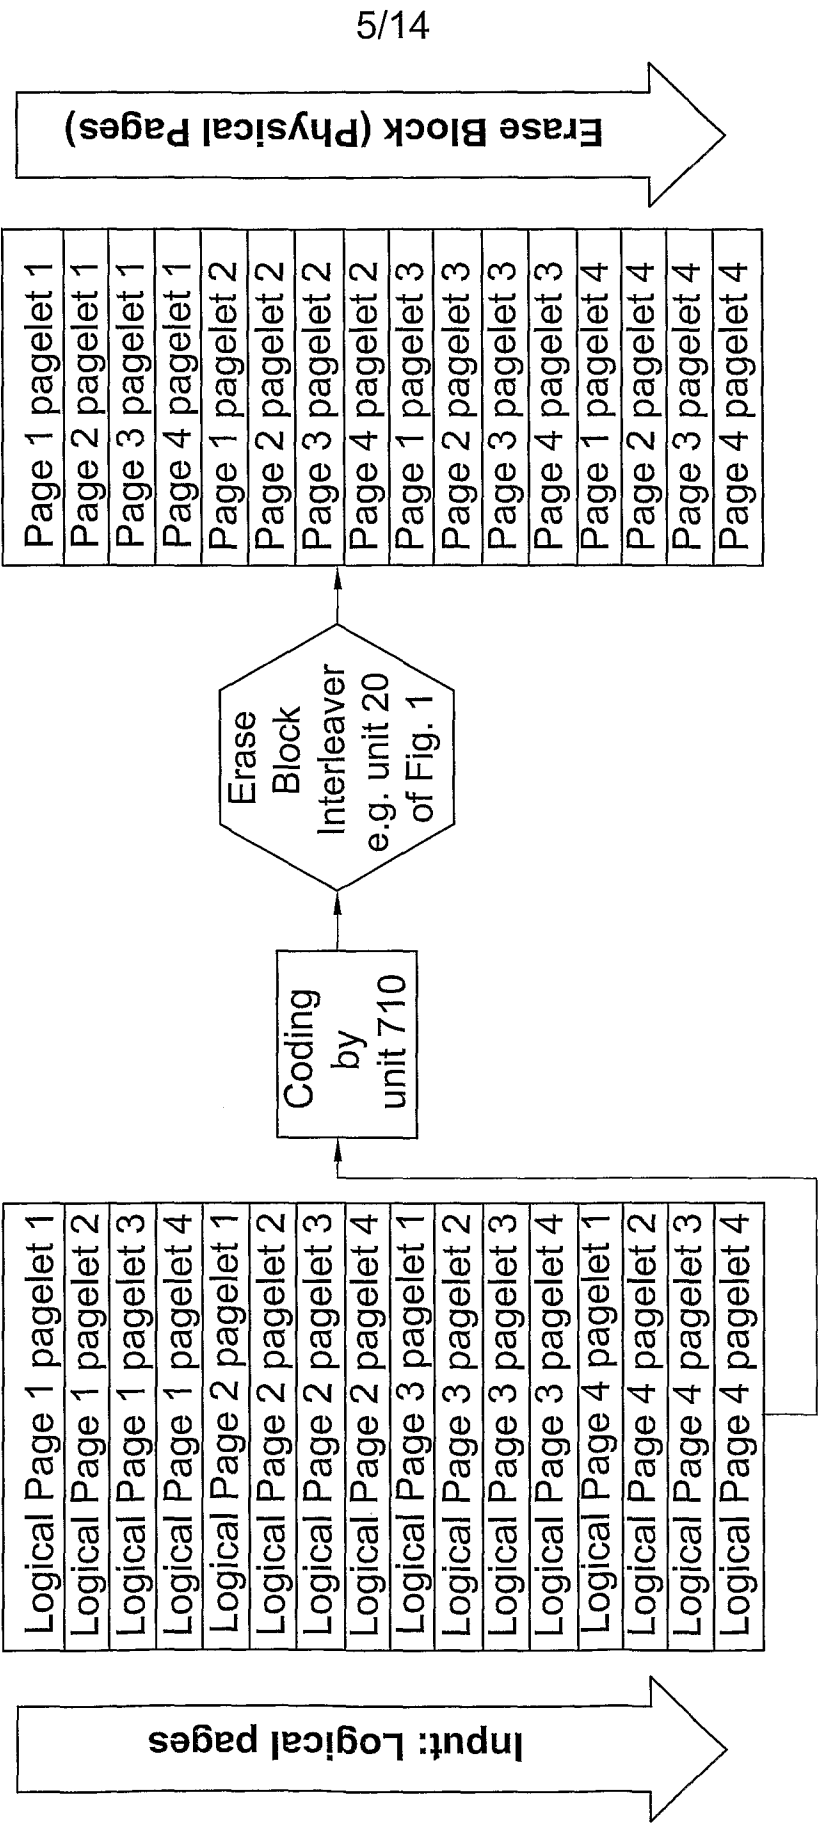


Figure 4

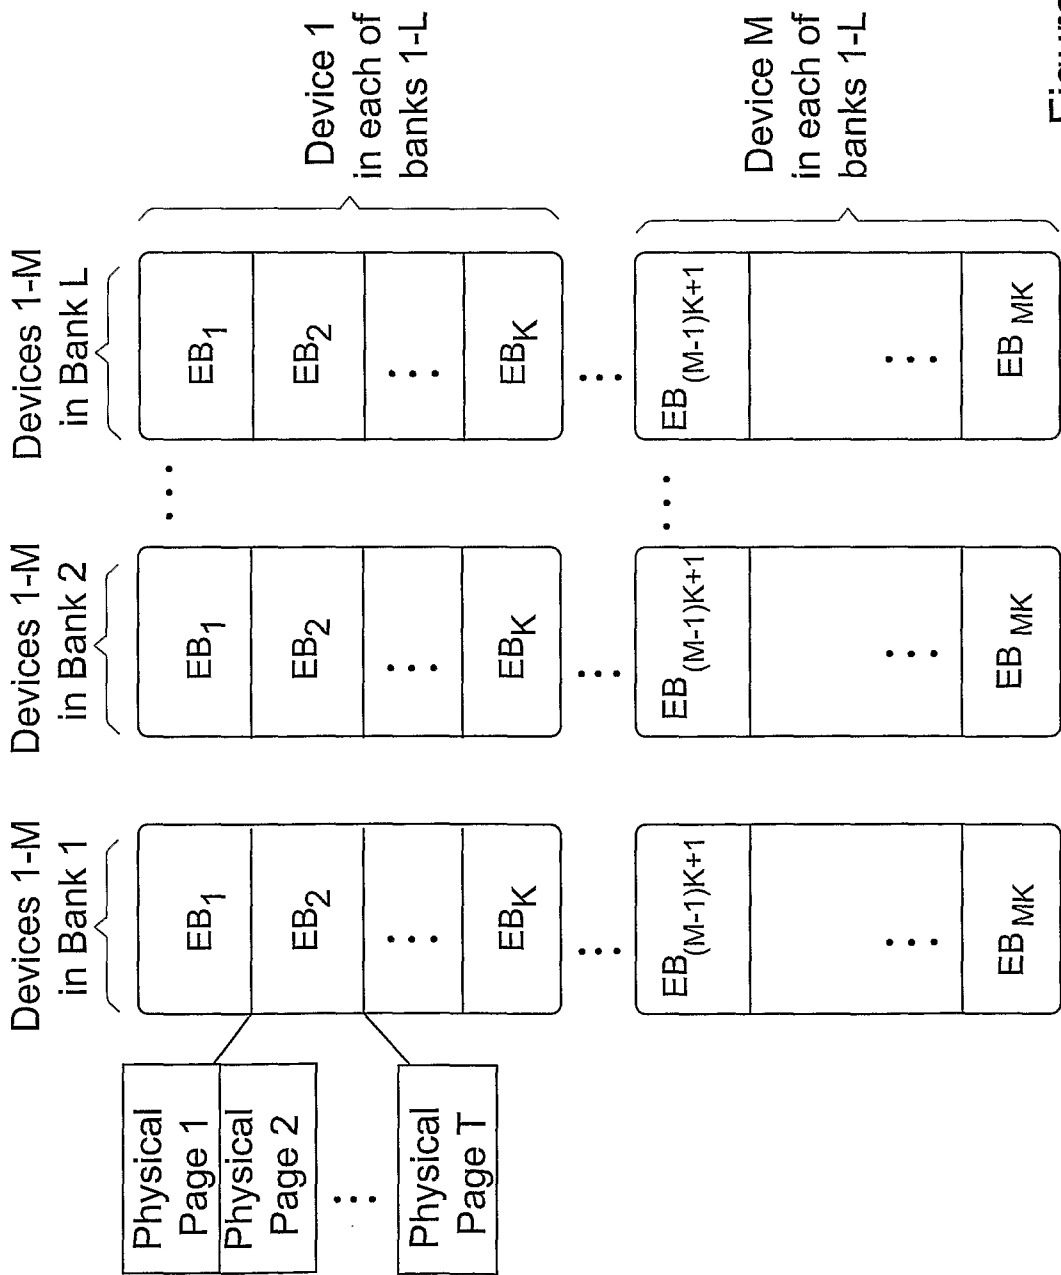


Figure 5 (PRIOR ART)

7/14

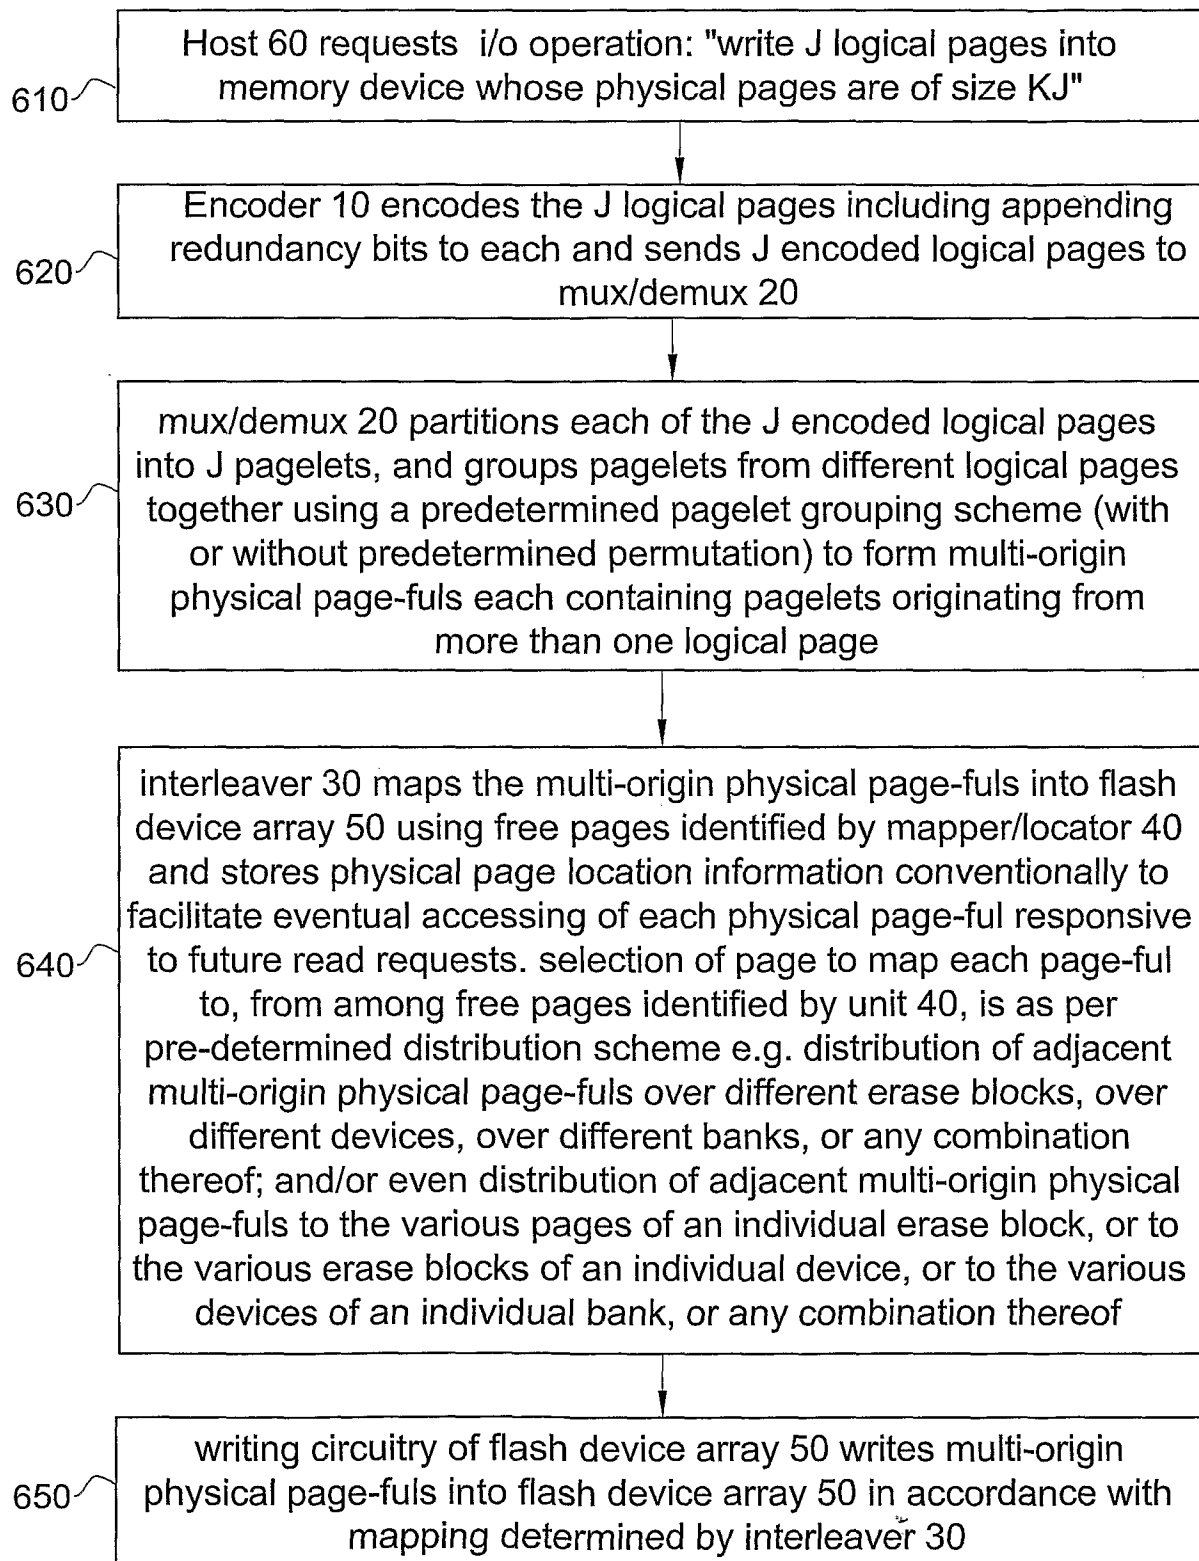


Figure 6a

8/14

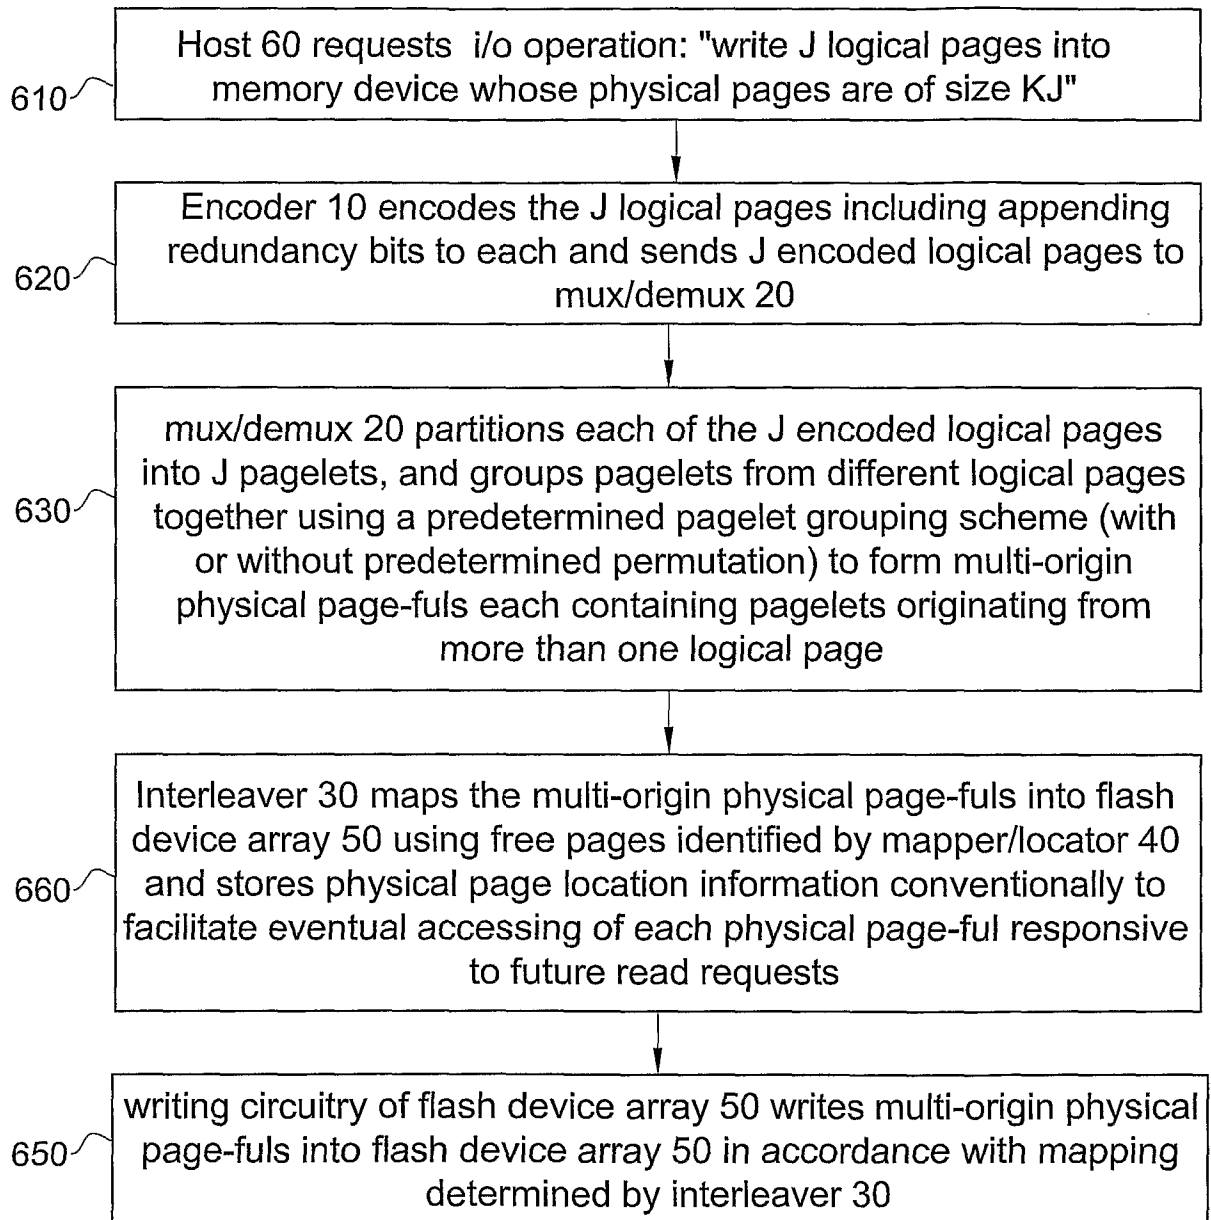


Figure 6b

9/14

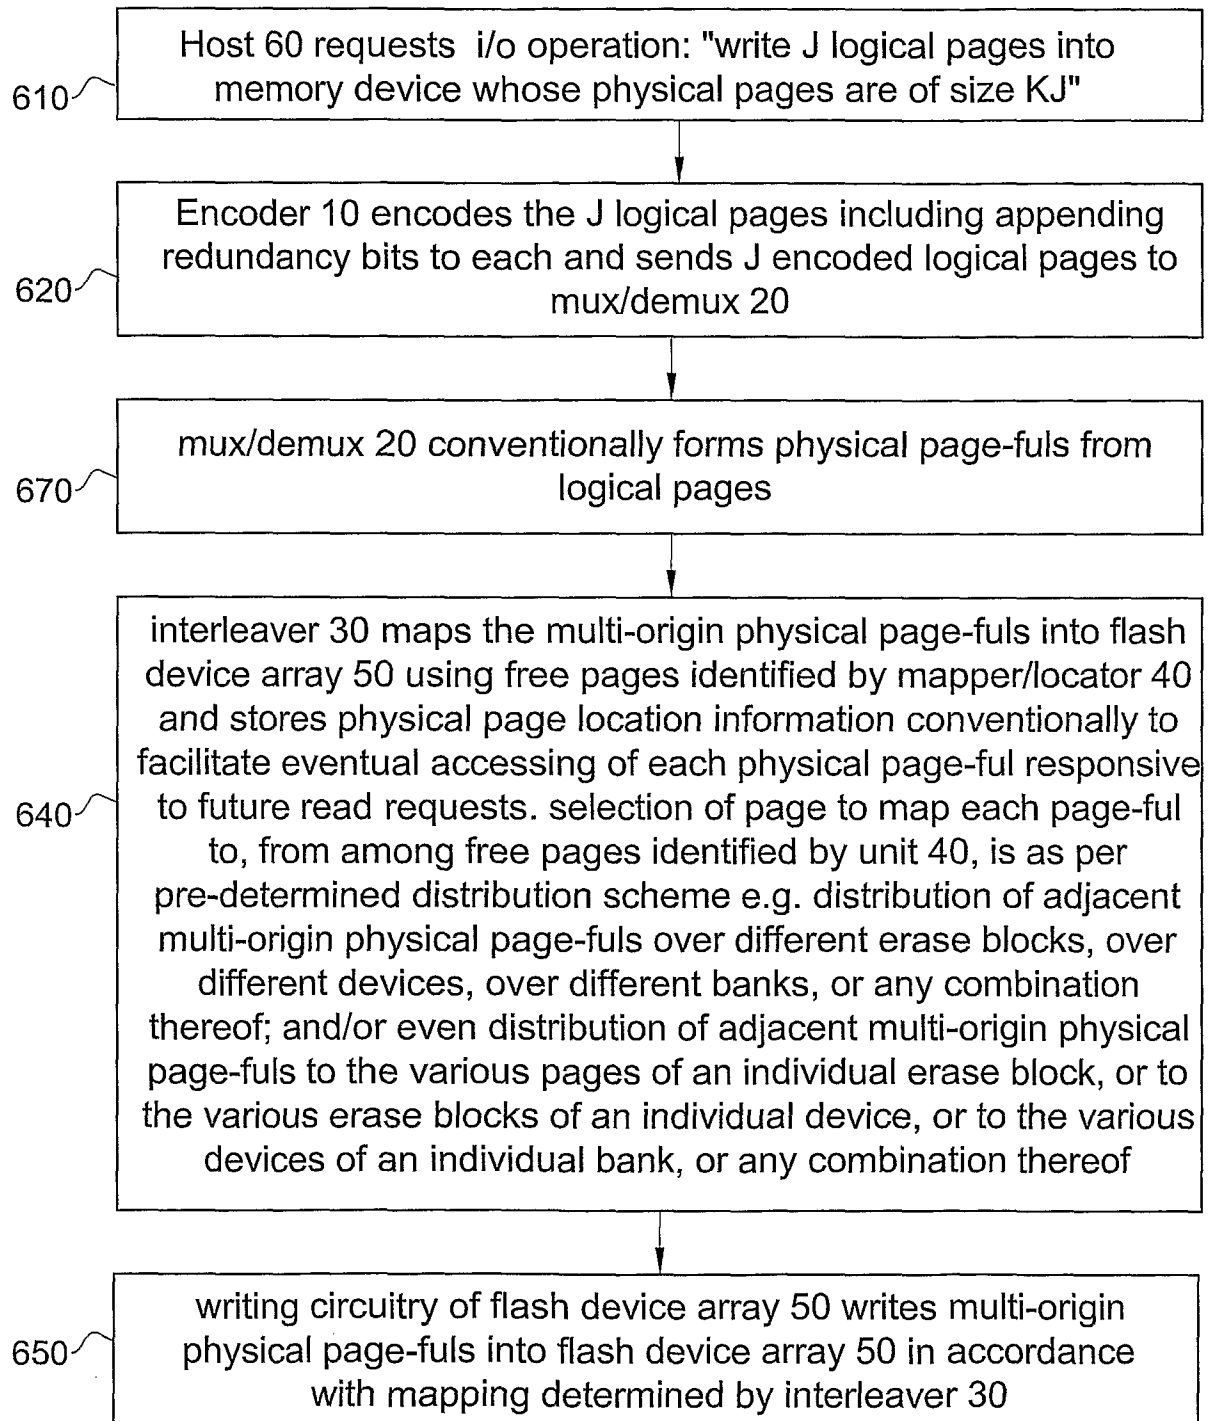


Figure 6c

10/14

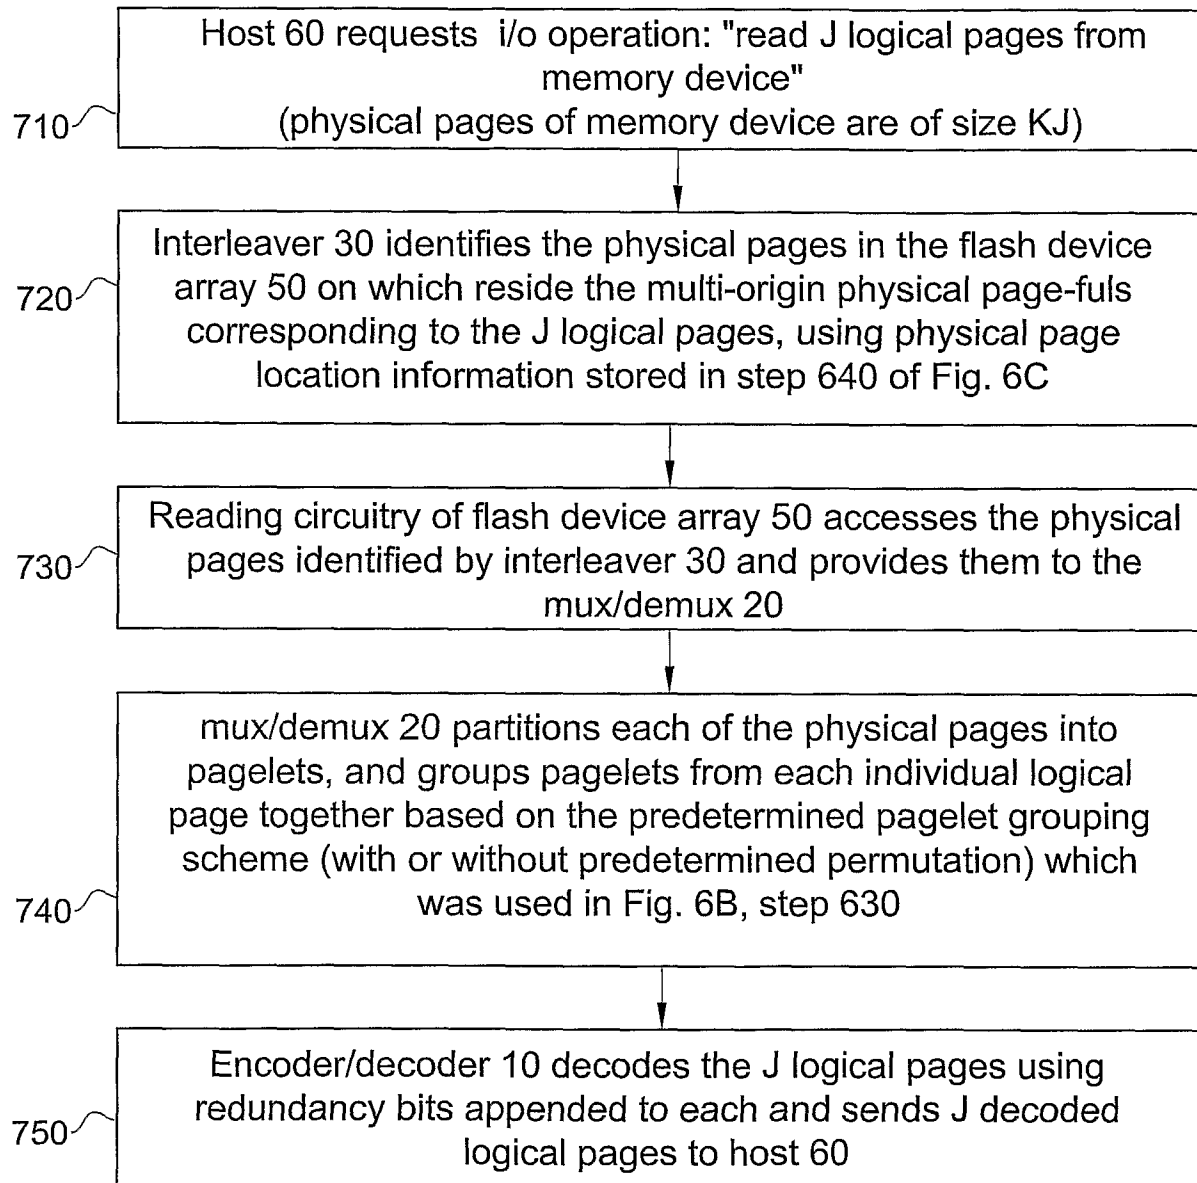


Figure 7

11/14

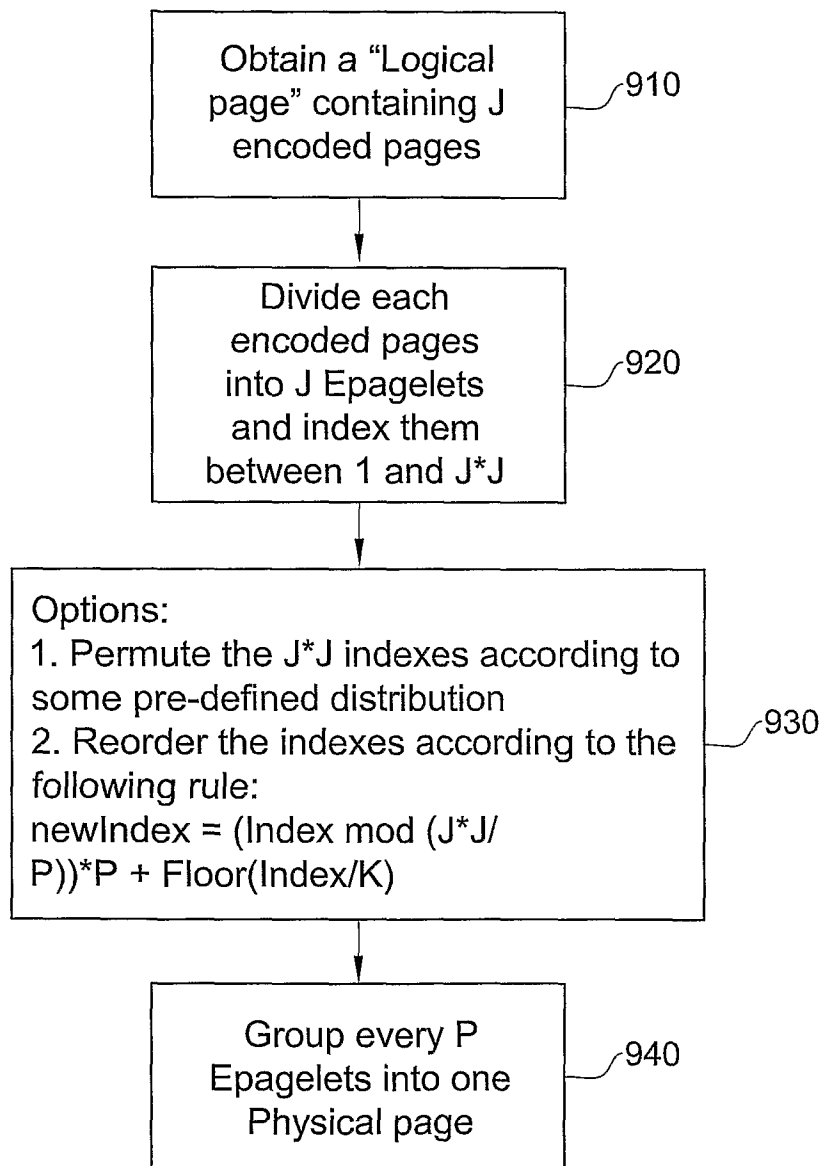


Figure 8

12/14

Virtual address	Block#	Page Usage Vector
0	10	1111101111..
256	5	1101101111..
512	1	1111111111..
768	2	0001111111..

Figure 9

Erase Block	Free
0	No
1	No
2	No
3	Yes

Figure 10

Virtual address	Block#	Page Usage Vector	Short/long Vector
0	10	1111101111..	1111001111..
256	5	1101101111..	1100001111..
512	1	1111111111..	1111111111..
768	2	0001111111..	0000111111..

Figure 11

13/14

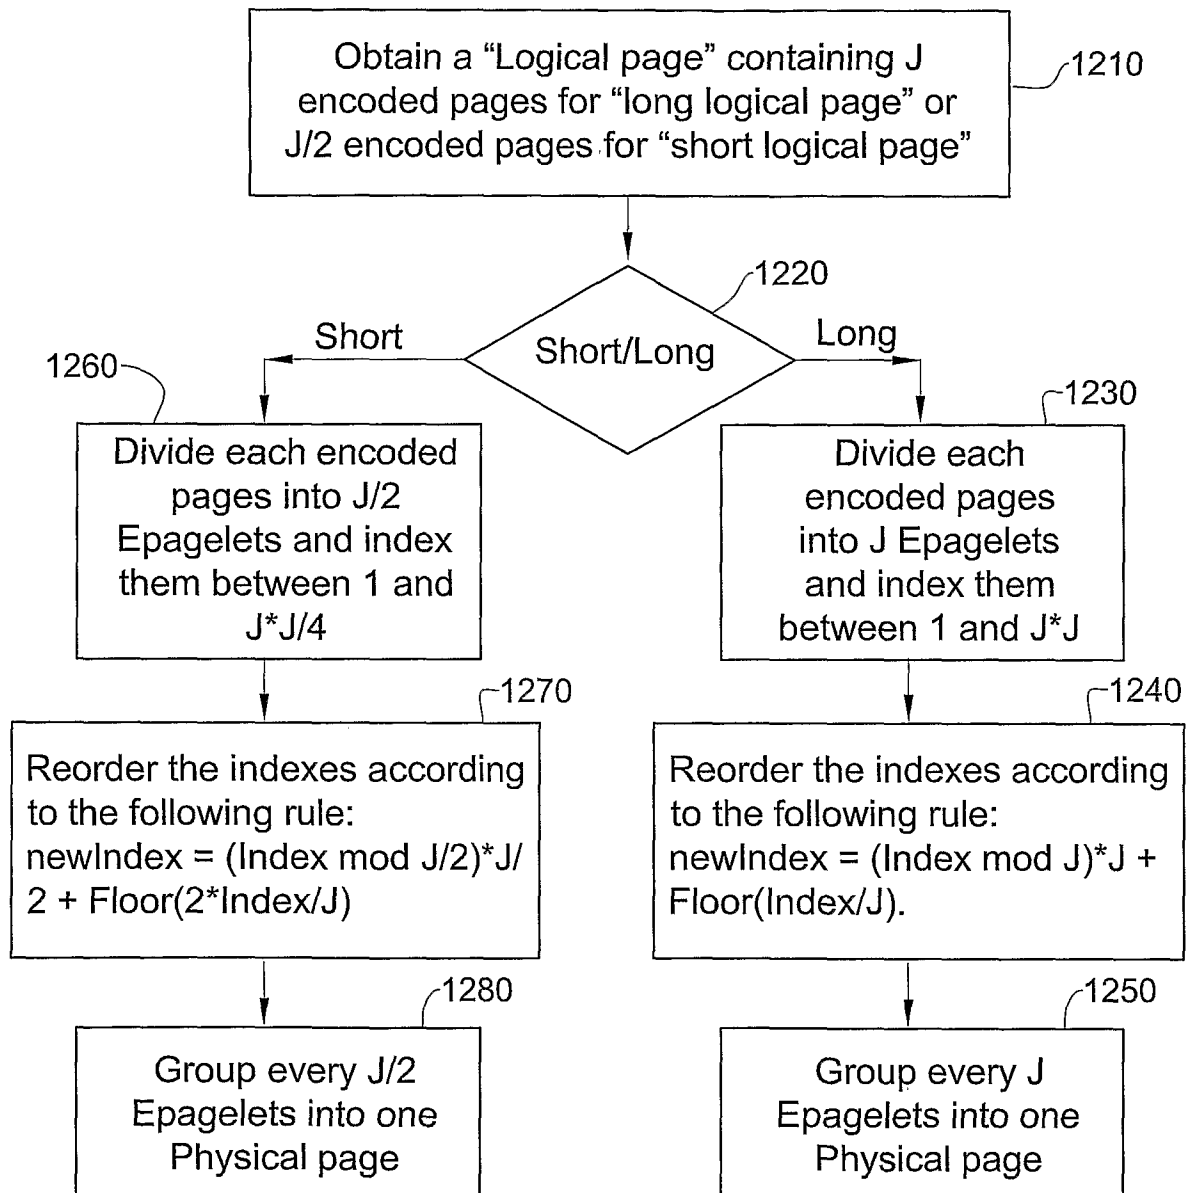


Figure 12

14/14

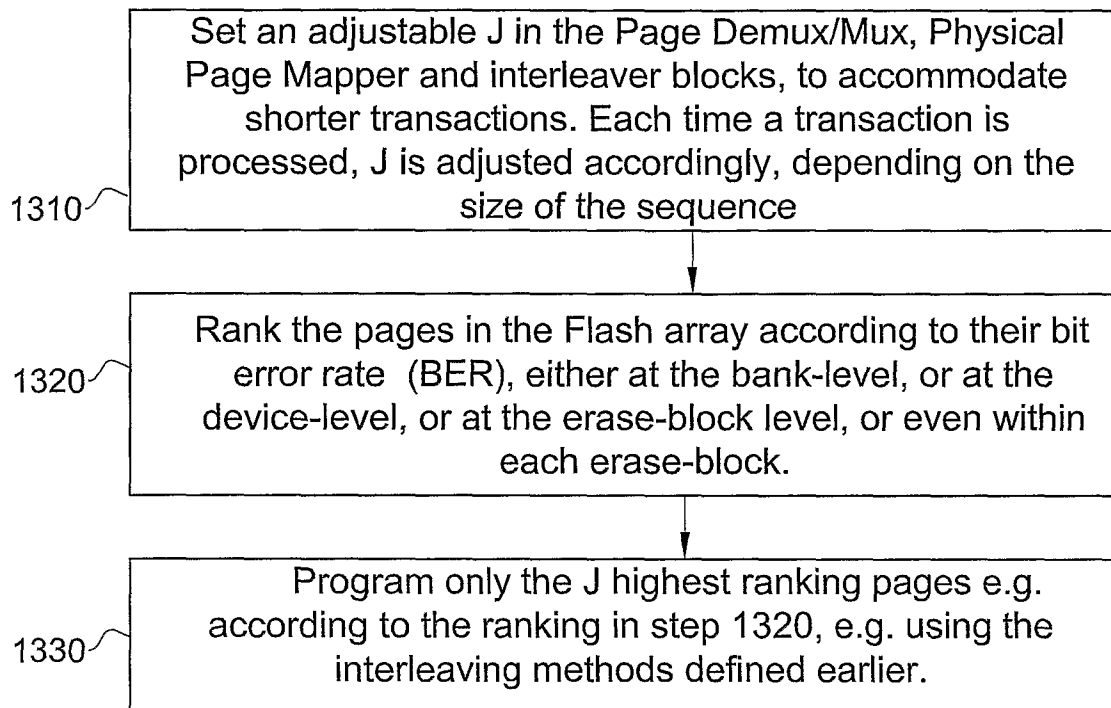


Figure 13