

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2015/088476 A1

(43) International Publication Date
18 June 2015 (18.06.2015)

(51) International Patent Classification:
G06F 11/00 (2006.01) *G06F 12/08* (2006.01)

(21) International Application Number:
PCT/US2013/073871

(22) International Filing Date:
9 December 2013 (09.12.2013)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P. [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).

(72) Inventors: **WARNES, Lidia**; 8000 Foothills Blvd., Roseville, California 95747 (US). **HANDGEN, Erin, A.**; 3404 E Harmony Rd., Ft. Collins, Colorado 80528-9544 (US). **WALTON, Andrew, C**; 8000 Foothills Blvd., Roseville, California 95747 (US).

(74) Agents: **HA, Miranda, J.** et al; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

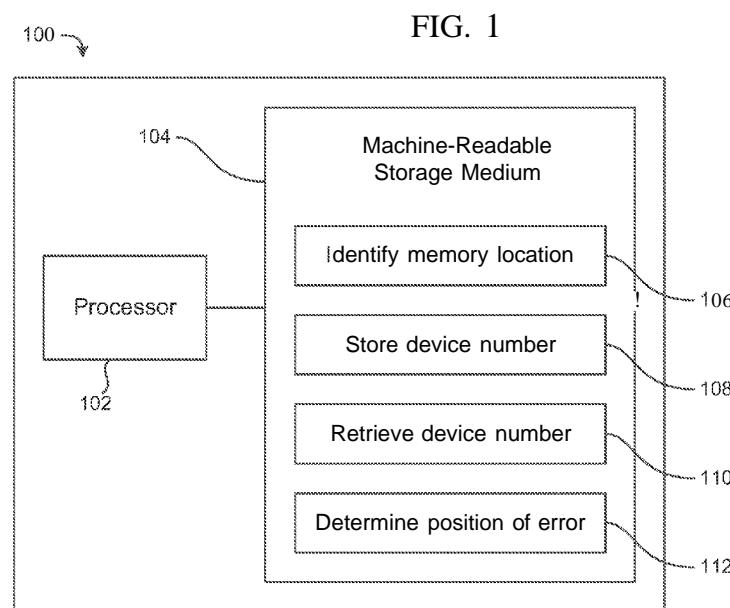
Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(l))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(H))

Published:

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

(54) Title: MEMORY ERASURE INFORMATION IN CACHE LINES



(57) Abstract: Example implementations relate to storing memory erasure information in memory devices on a memory module. In example implementations, a memory location associated with an error in a first cache line may be identified. The first cache line may include data read from the memory location, and the memory location may be in a first memory device of a plurality of memory devices on a memory module. A device number corresponding to the first memory device may be written to one of the plurality of memory devices. When the memory location is read for a second cache line, the device number corresponding to the first memory device may be retrieved. The second cache line may include the retrieved device number and data read from the memory location.

MEMORY ERASURE INFORMATION IN CACHE LINES**BACKGROUND**

[0001] A memory module may include many memory devices. Error correction logic may be used to correct errors detected in the memory devices. Memory modules may include spare memory devices so that when a memory device malfunctions, a spare memory device may be used instead of the malfunctioning memory device.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] The following detailed description references the drawings, wherein:

[0003] FIG. 1 is a block diagram of an example memory controller that includes a machine-readable storage medium encoded with instructions to store memory erasure information;

[0004] FIG. 2 is a block diagram of an example memory controller that includes a machine-readable storage medium encoded with instructions that enable memory erasure for a region of a memory device;

[0005] FIG. 3 is a block diagram of an example memory controller that includes a machine-readable storage medium encoded with instructions to correct cache line errors;

[0006] FIG. 4 is a block diagram of an example memory controller that includes a machine-readable storage medium encoded with instructions to identify and use device numbers in a cache line for memory erasure;

[0007] FIG. 5 is a block diagram of an example memory controller that includes a machine-readable storage medium encoded with instructions to track errors in a memory device;

[0008] FIG. 6 is a diagram of an example cache line having device numbers for memory erasure;

[0009] FIG. 7 is a flowchart of an example method for storing memory erasure information;

[0010] FIG. 8 is a flowchart of an example method for correcting a cache line error;

- 2 -

[001 1] FIG. 9 is a flowchart of an example method for storing device numbers to enable memory erasure on a cache line basis;

[0012] FIG. 10 is a flowchart of an example method for memory erasure for a region of a memory device; and

[001 3] FIG. 11 is a flowchart of an example method for identifying and using device numbers in a cache line for memory erasure.

DETAILED DESCRIPTION

[0014] Most errors in a memory module do not permeate through an entire memory device. However, even if an error is localized to a small portion (e.g., bit, row, column) of a memory device, the entire memory device may be replaced by a spare memory device. Thus, large error-free portions of a memory device may be unused because of a small error-prone portion of the memory device.

[001 5] In light of the above, the present disclosure provides for an increase in the granularity at which memory devices may be replaced. A memory controller may be programmed to identify portions of memory devices from which errors originate, to ignore and/or correct data that is read from the identified portions, and to process data from error-free portions of memory devices as usual (e.g., for a single memory device that has an error-prone portion and an error-free portion, data read from the error-prone portion may be ignored/corrected, while data read from the error-free portion may be processed as usual). The process of ignoring and/or correcting data read from a memory device may be referred to herein as "memory erasure". Information regarding which data to ignore/correct may be referred to herein as "memory erasure information".

[001 6] A memory controller may receive data in a cache line from a memory module. Data may be read from various memory devices of a memory module at substantially the same time, and the memory module may output a cache line that includes the data read from the various memory devices. In some implementations, data may be read from memory devices in bursts, such that multiple addresses in a memory device are read in response to each read command. Each memory device of a memory module may be assigned a device number, and the position of data in a cache line may correspond to the device number of the memory device from which the data was read.

[001 7] For example, a memory module may include 18 memory devices (e.g., dynamic random-access memories [DRAMs]) numbered 0-17, and a burst length of 8 may be used when reading the memory devices, with each address of a memory device holding 4 bits. Thus, when a read command is sent to such a memory module, the memory module may output a 576-bit cache line (calculated by multiplying 32, which is the number of bits read per memory device, by 18, which is the number of memory devices). In some implementations, the 32 most significant bits of a cache line may be the 32 bits read from memory device 0, the following 32 bits may be the 32 bits read from memory device 1, and so on, with the 32 least significant bits of the cache line being the 32 bits read from memory device 17. In some implementations, the 32 most significant bits of a cache line may be the 32 bits read from memory device 17, the following 32 bits may be the 32 bits read from memory device 16, and so on, with the 32 least significant bits of the cache line being the 32 bits read from memory device 0. It should be understood that other suitable orders of bits may be used, and that other numbering schemes for the device numbers may be used (e.g., the memory devices may be numbered from 1 to 18). Although the present disclosure discusses a memory module having 18 memory devices and a read burst length of 8, it should be understood that the concepts discussed herein are applicable to memory modules having different numbers of memory devices and different burst lengths.

[001 8] Referring now to the drawings, FIG. 1 is a block diagram of an example memory controller 100 that includes a machine-readable storage medium encoded with instructions to store memory erasure information. Memory controller 100 may be communicatively coupled to, and control access to, various memory devices (e.g., memory devices on a memory module). In FIG. 1, memory controller 100 includes processor 102 and machine-readable storage medium 104. As used herein, the terms "include", "have", and "comprise" are interchangeable and should be understood to have the same meaning.

[001 9] Processor 102 may include a central processing unit (CPU), microprocessor (e.g., semiconductor-based microprocessor), and/or other hardware device suitable for retrieval and/or execution of instructions stored in

machine-readable storage medium 104. Processor 102 may fetch, decode, and/or execute instructions 106, 108, 110, and 112 to enable storage of memory erasure information, as described below. As an alternative or in addition to retrieving and/or executing instructions, processor 102 may include an electronic circuit comprising a number of electronic components for performing the functionality of instructions 106, 108, 110, and/or 112.

[0020] Machine-readable storage medium 104 may be any suitable electronic, magnetic, optical, or other physical storage device that contains or stores executable instructions. Thus, machine-readable storage medium 104 may include, for example, a random-access memory (RAM), an Electrically Erasable Programmable Read-Only Memory (EEPROM), a storage device, an optical disc, and the like. In some implementations, machine-readable storage medium 104 may include a non-transitory storage medium, where the term "non-transitory" does not encompass transitory propagating signals. As described in detail below, machine-readable storage medium 104 may be encoded with a set of executable instructions 106, 108, 110, and 112.

[0021] Instructions 106 may identify a memory location associated with an error in a first cache line. For example, erroneous data in the first cache line may be detected or identified, and instructions 106 may identify the memory device and/or memory address from which the data was read. The first cache line may include data read from the memory location. The memory location may be in a first memory device of a plurality of memory devices on a memory module. The memory module may be an in-line memory module, such as a single in-line memory module (SIMM) or a dual in-line memory module (DIMM), or any memory module suitable for mounting memory integrated circuits (ICs).

[0022] Instructions 108 may store, in one of the plurality of memory devices on the memory module, a device number corresponding to the first memory device. The device number may be in binary form. The device number may be stored in the first memory device or a memory device other than the first memory device. In some implementations, instructions 108 may store, in another one of the plurality of memory devices, the device number corresponding to the first memory device.

Such redundant storage of the device number may allow the device number to be retrieved even if one of the memory devices storing the device number fails.

[0023] Instructions 110 may retrieve, when the memory location is read for a second cache line, the device number corresponding to the first memory device. The second cache line may include data read from the same memory addresses as those read for the first cache line, but at a later point in time. The second cache line may include the retrieved device number and data read from the memory location identified by instructions 106. In some implementations, instructions 110 may retrieve a first copy of the device number from one of the plurality of memory devices, and retrieve a second copy of the device number from a different one of the plurality of memory devices.

[0024] Instructions 112 may determine, based on the device number, a position of an error in the second cache line. As discussed above, the position of data in a cache line may correspond to the device number of the memory device from which the data was read. A device number in a cache line may indicate which group of bits in the cache line memory controller 100 should ignore or correct. For example, if a memory module has memory devices numbered 0-17, and 32 bits are read from each memory device, then if the device number in a cache line is 0, memory controller 100 may ignore or correct the 32 least significant bits in the cache line.

[0025] FIG. 2 is a block diagram of an example memory controller 200 that includes a machine-readable storage medium encoded with instructions that enable memory erasure for a region of a memory device. Memory controller 200 may be communicatively coupled to, and control access to, various memory devices (e.g., memory devices on a memory module). In FIG. 2, memory controller 200 includes processor 202 and machine-readable storage medium 204.

[0026] As with processor 102 of FIG. 1, processor 202 may include a CPU, microprocessor (e.g., semiconductor-based microprocessor), and/or other hardware device suitable for retrieval and/or execution of instructions stored in machine-readable storage medium 204. Processor 202 may fetch, decode, and/or execute instructions 206, 208, 210, 212, 214, and 216 to enable memory

erasure for a region of a memory device, as described below. As an alternative or in addition to retrieving and/or executing instructions, processor 202 may include an electronic circuit comprising a number of electronic components for performing the functionality of instructions 206, 208, 210, 212, 214, and/or 216.

[0027] As with machine-readable storage medium 104 of FIG. 1, machine-readable storage medium 204 may be any suitable physical storage device that stores executable instructions. Instructions 206, 208, 210, and 212 on machine-readable storage medium 204 may be analogous to (e.g., have functions and/or components similar to) instructions 106, 108, 110, and 112 on machine-readable storage medium 104. Instructions 214 may embed tag encoding bits in a stored device number corresponding to a first memory device. Tag encoding bits may be used to indicate ownership of cache lines and maintain cache coherency when multiple processors are used to access the same memory devices. Tag encoding bits may be embedded in stored device numbers when a device number may be discernible using less than all the bits reserved for the device number.

[0028] For example, 5 bits may be used to store a device number (in binary form) in a memory module having 18 memory devices numbered 0-17. For device numbers 0-15 (00000 to 01111 in binary form), all 5 bits may be needed to recognize the proper device number. However, for device numbers 16 (10000 in binary form) and 17 (10001 in binary form), the bits that matter are the most significant bit and the least significant bit; the three bits in between may be ignored for purposes of recognizing a device number. Thus, when an error is associated with device number 16 or 17, the third-most significant bit and the fourth-most significant bit of the 5 bits may be tag encoding bits. When an error is associated with device numbers 0-15, no tag encoding bits may be embedded in the device number, and snooping may be used to maintain cache coherency.

[0029] When none of the memory locations read for a cache line are associated with an error, the most significant bit and another bit other than the least-significant bit of the 5 bits may both be set to 1. For example, the two most significant bits of the 5 bits may both be set to 1 (e.g., the 5-bit string '11xxx' may be stored when there are no errors in a cache line, where each 'x' may represent either a 0 or a 1). If the most significant bit and any other bit except for the least significant bit of the 5

bits are both 1, it may not matter what the rest of the bits are for purposes of indicating that there are no errors in the cache line. Thus, any of the three remaining bits may be tag encoding bits. For consistency with situations where a memory location read for a cache line is associated with an error, the third-most significant bit and fourth-most significant bit may be tag encoding bits (e.g., the 5-bit string '11tx' may be stored when there are no errors in a cache line, where each 't' represents a tag encoding bit and the 'x' may represent either a 0 or a 1). Although 5-bit strings for storing device numbers are mentioned herein, it should be understood that the principles and concepts of the present disclosure are applicable for longer or shorter strings for storing device numbers. The number of bits used to store device numbers may be adjusted based on how many memory devices are on a memory module and/or how many tag encoding bits are to be embedded in a device number.

[0030] If tag encoding bits are embedded in a device number retrieved by instructions 210, instructions 212 may apply a mask to the device number when using the device number to determine the position of an error in a cache line. For example, if the device number is 5 bits long and the third- and fourth-most significant bits are tag encoding bits, instructions 212 may use a mask to block out the third- and fourth-most significant bits and then determine what the device number is based on the remaining bits.

[0031] Instructions 216 may designate a region around a memory location in a first memory device as defective. For example, instructions 216 may designate a row, column, and/or bank that includes a particular memory location as defective. A region may be designated as defective when the number of errors originating from the region exceeds a specified threshold. When an address in the designated region is read, instructions 210 may retrieve a device number corresponding to the first memory device.

[0032] FIG. 3 is a block diagram of an example memory controller 300 that includes a machine-readable storage medium encoded with instructions to correct cache line errors. Memory controller 300 may be communicatively coupled to, and control access to, various memory devices (e.g., memory devices

on a memory module). In FIG. 3, memory controller 300 includes processor 302 and machine-readable storage medium 304.

[0033] As with processor 102 of FIG. 1, processor 302 may include a CPU, microprocessor (e.g., semiconductor-based microprocessor), and/or other hardware device suitable for retrieval and/or execution of instructions stored in machine-readable storage medium 304. Processor 302 may fetch, decode, and/or execute instructions 306, 308, 310, and 312 to enable correction of cache line errors, as described below. As an alternative or in addition to retrieving and/or executing instructions, processor 302 may include an electronic circuit comprising a number of electronic components for performing the functionality of instructions 306, 308, 310, and/or 312.

[0034] As with machine-readable storage medium 104 of FIG. 1, machine-readable storage medium 304 may be any suitable physical storage device that stores executable instructions. In some implementations, machine-readable storage medium 304 may include a non-transitory storage medium. As described in detail below, machine-readable storage medium 304 may be encoded with a set of executable instructions 306, 308, 310, and 312.

[0035] Instructions 306 may detect an error in a cache line. For example, instructions 306 may detect a non-matching checksum computed for a cache line. The detected error may be associated with a first memory device of a plurality of memory devices on a memory module. The memory module may be an in-line memory module, such as a SIMM or DIMM, or any memory module suitable for mounting memory ICs.

[0036] Instructions 308 may identify, in a cache line, a first device number. The first device number may correspond to a memory device associated with a detected error (e.g., data read from the memory device may contain the detected error), and may be read from one of a plurality of memory devices on a memory module. The first device number may appear in a designated position in a cache line. For example, a device number may be 5 bits, and any cache line that contains a device number may have the device number in the 5 most significant bits of the cache line. The 5 bits of a device number may appear consecutively in a cache line, or may be split up into different parts of a cache line.

[0037] In some implementations, tag encoding bits may be embedded in the first device number, as discussed above with respect to FIG. 2. In such implementations, instructions 308 may apply a mask to the first device number to block out the tag encoding bits. For example, if the device number is 5 bits long and the third- and fourth-most significant bits are tag encoding bits, instructions 308 may use a mask to block out the third- and fourth-most significant bits and then determine what device number is specified by the remaining bits.

[0038] Instructions 310 may determine, based on a first device number, a position of a detected error in a cache line. As discussed above, the position of data in a cache line may correspond to the device number of the memory device from which the data was read. Instructions 310 may identify the position, in the cache line, that corresponds to the first device number. For example, if 32 bits are read from each memory device of a memory module and the bits read from memory device 0 are the 32 least significant bits of a cache line, then if the cache line includes the device number 00000, instructions 310 may determine that a detected error is in the last 32 bits of the cache line.

[0039] Instructions 312 may correct a detected error in a cache line. In some implementations, instructions 312 may use non-erroneous data in a cache line to reconstruct data for a portion of the cache line having the detected error. The portion of the cache line having the detected error may include data read from a memory device corresponding to a device number identified in the cache line. In some implementations, instructions 312 may ignore the data in a certain position of a cache line (e.g., the portion of the cache line that includes data read from the memory device corresponding to a device number in the cache line) and use data in other positions in the cache line to determine what data should be in the ignored position.

[0040] FIG. 4 is a block diagram of an example memory controller 400 that includes a machine-readable storage medium encoded with instructions to identify and use device numbers in a cache line for memory erasure. Memory controller 400 may be communicatively coupled to, and control access to, various memory devices (e.g., memory devices on a memory module). In FIG. 4, memory

controller 400 includes processor 402 and machine-readable storage medium 404.

[0041] As with processor 302 of FIG. 3, processor 402 may include a CPU, microprocessor (e.g., semiconductor-based microprocessor), and/or other hardware device suitable for retrieval and/or execution of instructions stored in machine-readable storage medium 404. Processor 402 may fetch, decode, and/or execute instructions 406, 408, 410, 412, 414, 416, and 418 to enable memory erasure for a region of a memory device, as described below. As an alternative or in addition to retrieving and/or executing instructions, processor 402 may include an electronic circuit comprising a number of electronic components for performing the functionality of instructions 406, 408, 410, 412, 414, 416, and/or 418.

[0042] As with machine-readable storage medium 304 of FIG. 3, machine-readable storage medium 404 may be any suitable physical storage device that stores executable instructions. Instructions 406, 408, 410, and 412 on machine-readable storage medium 404 may be analogous to instructions 306, 308, 310, and 312 on machine-readable storage medium 304. Instructions 408 may identify a first device number in a cache line, and the first device number may be read from one of a plurality of memory devices on a memory module. Instructions 414 may identify a second device number in the cache line. The second device number may be read from another one of the plurality of memory devices. The second device number may correspond to a memory device associated with an error detected by instructions 406, and may be a copy of the first device number. In some implementations, tag encoding bits may be embedded in the second device number, and instructions 414 may use a mask when identifying the second device number, in a manner analogous to that discussed above with respect to FIG. 3.

[0043] The second device number may appear in a designated position in a cache line. For example, a device number may be 5 bits, and any cache line that contains two device numbers may have the second device number in the 5 least significant bits of the cache line. The 5 bits of the second device number may appear consecutively in a cache line, or may be split up into different parts of a cache line.

[0044] Instructions 416 may compare the first device number identified by instructions 408 and the second device number identified by instructions 414. For example, instructions 416 may input the first device number and the second device number into a comparator. If the first device number and the second device number match, the device number may be used to determine the position of an error in a cache line detected by instructions 406, and the error may be corrected. If the first device number and the second device number do not match, an attempt may be made to use the first device number to determine the position of an error in a cache line and to correct the error.

[0045] Instructions 418 may determine, based on the second device number, the position of a detected error in the cache line. Instructions 418 may be executed if the first device number and the second device number do not match, and if the detected error is not successfully corrected when the first device number is used to determine the position of the detected error. As discussed above, the position of data in a cache line may correspond to the device number of the memory device from which the data was read. Instructions 418 may identify the position, in the cache line, that corresponds to the second device number. For example, if 32 bits are read from each memory device of a memory module and the bits read from memory device 0 are the 32 least significant bits of a cache line, then if the second device number is 00000, instructions 418 may determine that a detected error is in the last 32 bits of the cache line. Instructions 412 may correct the error in the determined position of the cache line.

[0046] FIG. 5 is a block diagram of an example memory controller 500 that includes a machine-readable storage medium encoded with instructions to track errors in a memory device. Memory controller 500 may be communicatively coupled to, and control access to, various memory devices (e.g., memory devices on a memory module). In FIG. 5, memory controller 500 includes processor 502 and machine-readable storage medium 504.

[0047] As with processor 302 of FIG. 3, processor 502 may include a CPU, microprocessor (e.g., semiconductor-based microprocessor), and/or other hardware device suitable for retrieval and/or execution of instructions stored in

machine-readable storage medium 504. Processor 502 may fetch, decode, and/or execute instructions 506, 508, 510, 512, 514, and 516 to enable tracking of errors in a memory device, as described below. As an alternative or in addition to retrieving and/or executing instructions, processor 502 may include an electronic circuit comprising a number of electronic components for performing the functionality of instructions 506, 508, 510, 512, 514, and/or 516.

[0048] As with machine-readable storage medium 304 of FIG. 3, machine-readable storage medium 504 may be any suitable physical storage device that stores executable instructions. Instructions 506, 508, 510, and 512 on machine-readable storage medium 504 may be analogous to instructions 306, 308, 310, and 312 on machine-readable storage medium 304. Instructions 514 may keep track of how many errors are associated with a region of a first memory device on a memory module. The memory module may be an in-line memory module, such as a SIMM or DIMM, or any memory module suitable for mounting memory ICs. The region of the first memory device may be a row or column of the first memory device. In some implementations, instructions 514 may keep track of how many errors are associated with a memory bank that includes part of the first memory device.

[0049] In some implementations, instructions 514 may maintain an error counter that is incremented whenever an error associated with the region (e.g., an error in data read from an address in the region) is detected. If the number of errors associated with the region exceeds a specified threshold value, the region may be designated as defective. Data read from an address in a region designated as defective may be deemed unreliable.

[0050] Instructions 516 may retrieve, when an address in the region is read, a device number corresponding to the first memory device. Instructions 516 may be executed if the region is designated as defective. The device number may be retrieved from one of a plurality of memory devices on the memory module. In some implementations, instructions 516 may retrieve a copy of the device number from another one of the plurality of memory devices, such that the device number appears twice in a cache line. If one of the memory devices from which

the device number is retrieved fails, an accurate copy of the device number may still appear in a cache line and be used for memory erasure.

[0051] In some implementations, a patrol scrubber may be used to program memory devices in the memory module such that whenever an address in a defective region is read, the device number corresponding to the memory device containing the defective region is also retrieved. Thus, both the data read from the defective region and the corresponding device number may appear in the same cache line. Instructions 510 may use the device number to determine where in the cache line the data read from the defective region is so that the data can be ignored or corrected.

[0052] FIG. 6 is a diagram of an example cache line 600 having device numbers for memory erasure. Cache line 600 may be output by a memory module having 18 DRAMs. A burst length of 8 may be used when reading the DRAMs, as illustrated by the 8 rows in FIG. 6. Each DRAM address may hold 4 bits. Thus, cache line 600 may have 576 bits (calculated by multiplying 32, which is the number of bits read per DRAM, by 18, which is the number of DRAMs).

[0053] The left-most column in FIG. 6 represents data read for cache line 600 from a first DRAM of the memory module. The second column from the left in FIG. 6 represents data read for cache line 600 from a second DRAM of the memory module. The third column from the left in FIG. 6 represents data read for cache line 600 from a third DRAM of the memory module. The right-most column in FIG. 6, which is wider than the other three columns, represents data read for cache line 600 from the remaining 15 DRAMs of the memory module.

[0054] The first DRAM may store parity bits. Cache line 600 includes 32 parity bits, which may enable detection of 100% of errors in cache line 600. The second and third DRAMs may store device numbers (represented by "DRAM #" in FIG. 6), cyclic redundancy check (CRC) bits, and other data. Cache line 600 includes 10 device number bits. The 10 bits may be two copies of the same 5-bit device number. In some implementations, the device numbers may include embedded tag encoding bits (e.g., 2 tag encoding bits per device number, 4 tag encoding bits total in cache line 600), as discussed above with respect to FIG. 2. Cache line 600 includes 22 CRC bits, which may enable 99.999976% of errors in

cache line 600 to be corrected. The remaining 15 DRAMs may store data other than parity bits, device numbers, and CRC bits.

[0055] Methods related to memory erasure are discussed with respect to FIGS. 7-1 1. FIG. 7 is a flowchart of an example method 700 for storing memory erasure information. Although execution of method 700 is described below with reference to processor 102 of FIG. 1, it should be understood that execution of method 700 may be performed by other suitable devices, such as processors 202 and 302 of FIGS. 2 and 3, respectively. Method 700 may be implemented in the form of executable instructions stored on a machine-readable storage medium and/or in the form of electronic circuitry.

[0056] Method 700 may start in block 702, where processor 102 may identify a memory location, in a first memory device of a plurality of memory devices on a memory module, associated with an error in a first cache line. For example, processor 102 may detect or identify erroneous data in the first cache line, and identify the memory device and/or memory address from which the data was read. The memory module may be an in-line memory module, such as a SIMM or DIMM, or any memory module suitable for mounting memory ICs. The first cache line may include data read from the memory location.

[0057] Next, in block 704, processor 102 may write a device number to one of the plurality of memory devices. The device number may correspond to the first memory device, and may be in binary form. The device number may be stored in the first memory device or a memory device other than the first memory device. In some implementations, 5 bits may be used to store the device number. In some implementations, tag encoding bits may be embedded in the device number. The number of bits used to store the device number may be adjusted based on how many memory devices are on the memory module and/or how many tag encoding bits are to be embedded in a device number.

[0058] Finally, in block 706, processor 102 may retrieve the device number written in block 704 when the memory location identified in block 702 is read for a second cache line. The second cache line may include the retrieved device number and data read from the memory location. The second cache line may include data read from the same memory addresses as those read for the first cache line, but

at a later point in time. The device number retrieved in block 706 may be used to correct an error in the second cache line, as discussed above with respect to FIG. 3.

[0059] FIG. 8 is a flowchart of an example method 800 for correcting a cache line error. Although execution of method 800 is described below with reference to processor 302 of FIG. 3, it should be understood that execution of method 800 may be performed by other suitable devices, such as processors 402 and 502 of FIGS. 4 and 5, respectively. In some implementations, some blocks of method 800 may be performed in parallel with and/or after method 700. Method 800 may be implemented in the form of executable instructions stored on a machine-readable storage medium and/or in the form of electronic circuitry.

[0060] Method 800 may start in block 802, where processor 302 may detect an error in a cache line that includes data read from a memory location in a first memory device corresponding to a stored device number. For example, processor 302 may detect a non-matching checksum computed for the cache line. The first memory device may be one of a plurality of memory devices on a memory module, and the detected error may be associated with the first memory device (e.g., data read from the first memory device for the cache line may contain errors). The memory module may be an in-line memory module, such as a SIMM or DIMM, or any memory module suitable for mounting memory ICs. Along with data read from the first memory device, the cache line may also include the device number corresponding to the first memory device.

[0061] Next, in block 804, processor 302 may determine, based on the device number, a position, in the cache line, of data read from the memory location. As discussed above, the position of data in a cache line may correspond to the device number of the memory device from which the data was read. Processor 302 may identify the position, in the cache line, that corresponds to the device number in the cache line. For example, if 32 bits are read from each memory device of the memory module and the bits read from memory device 0 are the 32 least significant bits of the cache line, then if the cache line includes the device number 00000, processor 302 may determine that the error detected in block 802 is in the last 32 bits of the cache line.

[0062] In some implementations, tag encoding bits may be embedded in the device number, as discussed above with respect to FIG. 2. In such implementations, processor 302 may apply a mask to the device number to block out the tag encoding bits. For example, if the device number is 5 bits long and the third- and fourth-most significant bits are tag encoding bits, processor 302 may use a mask to block out the third- and fourth-most significant bits and then determine what device number is specified by the remaining bits.

[0063] Finally, in block 806, processor 302 may correct the data in the determined position. In some implementations, processor 302 may use non-erroneous data in the cache line to reconstruct data for a portion of the cache line having the detected error. In some implementations, processor 302 may ignore the data in the determined position of the cache line (e.g., the portion of the cache line that includes data read from the memory device corresponding to the device number in the cache line) and use data in other positions in the cache line to determine what data should be in the determined position.

[0064] FIG. 9 is a flowchart of an example method 900 for storing device numbers to enable memory erasure on a cache line basis. Although execution of method 900 is described below with reference to processor 302 of FIG. 3, it should be understood that execution of method 900 may be performed by other suitable devices, such as processors 402 and 502 of FIGS. 4 and 5, respectively. In some implementations, some blocks of method 900 may be performed in parallel with and/or after method 700. Method 900 may be implemented in the form of executable instructions stored on a machine-readable storage medium and/or in the form of electronic circuitry.

[0065] Method 900 may start in block 902, where processor 302 may identify a memory location, in a first memory device of a plurality of memory devices on a memory module, associated with an error in a cache line. For example, processor 302 may detect or identify erroneous data in the cache line, and identify the memory address from which the data was read. The cache line may include data read from the memory location. The memory module may be an in-line memory module, such as a SIMM or DIMM, or any memory module suitable for mounting memory ICs.

[0066] Next, in block 904, processor 302 may determine whether there is a device number in the cache line. For example, processor 302 may determine whether a device number appears in a designated position in the cache line, as discussed above with respect to FIG. 3. If processor 302 determines that there is a device number in the cache line, method 900 may proceed to block 906, in which processor 302 may determine, using the device number, the position of the error in the cache line and correct the error, as discussed above with respect to FIG. 8.

[0067] If, in block 904, it is determined that there is not a device number in the cache line (e.g., if there is not a valid device number in the designated position in the cache line), method 900 may proceed to block 908, in which processor 302 may scrub the identified memory location in the first memory device (e.g., write correct data to the identified memory location). Next, in block 910, processor 302 may read data from the scrubbed memory location.

[0068] In block 912, processor 302 may determine whether the data read from the memory location is erroneous. For example, processor 302 may run a checksum algorithm on the data. If the data is not erroneous, method 900 may proceed to block 914, and the device number corresponding to the first memory device may not be stored.

[0069] If, in block 912, it is determined that the data read from the memory location is erroneous, method 900 may proceed to block 916, in which processor 302 may write, to one of the plurality of memory devices, a device number corresponding to the first memory device. The device number may be in binary form (e.g., 5 bits). The device number may be written to the first memory device or to a memory device other than the first memory device.

[0070] In block 918, processor 302 may embed tag encoding bits in the device number. For example, the device number may be stored in a 5-bit string, and two of the five bits may be tag encoding bits, as discussed above with respect to FIG. 2. It should be understood that block 918 may not be performed in implementations where snooping is used to maintain cache coherency.

[0071] In block 920, processor 302 may write, to another one of the plurality of memory devices, a copy of the device number corresponding to the first memory device. The copy of the device number may include embedded tag encoding bits.

Such redundant storage of the device number may allow the device number to be retrieved even if one of the memory devices to which the device number is written fails.

[0072] FIG. 10 is a flowchart of an example method 1000 for memory erasure for a region of a memory device. Although execution of method 1000 is described below with reference to processor 502 of FIG. 5, it should be understood that execution of method 1000 may be performed by other suitable devices, such as processor 202 of FIG. 2. In some implementations, some blocks of method 1000 may be performed in parallel with and/or after method 700. Method 1000 may be implemented in the form of executable instructions stored on a machine-readable storage medium and/or in the form of electronic circuitry.

[0073] Method 1000 may start in block 1002, where processor 502 may keep track of how many errors are associated with a region around a memory location of a first memory device. The region may be a row, column, and/or memory bank that includes the memory location. In some implementations, processor 502 may maintain an error counter that is incremented whenever an error associated with the region (e.g., an error in data read from an address in the region) is detected.

[0074] In block 1004, processor 502 may determine whether the number of errors in the region exceeds a threshold number of errors. For example, processor 502 may input the value of the error counter and the threshold number into a comparator. If the number of errors in the region does not exceed the threshold number, method 1000 may loop back to block 1002.

[0075] If, in block 1004, it is determined that the number of errors in the region exceeds the threshold number of errors, method 1000 may proceed to block 1006, in which processor 502 may designate the region around the memory location as defective. In some implementations, method 1000 may proceed to block 1006 if the number of errors in the region equals the threshold number of errors. Data read from an address in a region designated as defective may be deemed unreliable.

[0076] Finally, in block 1008, processor 502 may retrieve, when an address in the designated region is read, the device number corresponding to the first memory device. Processor 502 may retrieve the device number from one of a plurality of

memory devices on a memory module. In some implementations, processor 502 may retrieve a copy of the device number from another one of the plurality of memory devices, such that the device number appears twice in a cache line. If one of the memory devices from which the device number is retrieved fails, an accurate copy of the device number may still appear in a cache line and be used for memory erasure.

[0077] FIG. 11 is a flowchart of an example method 1100 for identifying and using device numbers in a cache line for memory erasure. Although execution of method 1100 is described below with reference to processor 402 of FIG. 4, it should be understood that execution of method 1100 may be performed by other suitable devices, such as processors 102 and 302 of FIGS. 1 and 3, respectively. In some implementations, some blocks of method 1100 may be performed in parallel with and/or after method 800. Method 1100 may be implemented in the form of executable instructions stored on a machine-readable storage medium and/or in the form of electronic circuitry.

[0078] Method 1100 may start in block 1102, where processor 402 may detect an error in a cache line. For example, processor 402 may detect a non-matching checksum computed for the cache line. The detected error may be associated with a memory location in a memory device on a memory module. The memory module may have a plurality of memory devices, and may be an in-line memory module, such as a SIMM or DIMM, or any memory module suitable for mounting memory ICs.

[0079] Next, in block 1104, processor 402 may identify a first device number in the cache line. The first device number may correspond to a memory device associated with the detected error (e.g., data read from the memory device may contain the detected error), and may be read from one of a plurality of memory devices on a memory module. The first device number may appear in a designated position in the cache line, as discussed above with respect to FIG. 3. In some implementations, tag encoding bits may be embedded in the first device number, and processor 402 may use a mask when identifying the first device number, as discussed above with respect to FIG. 3.

- 20 -

[0080] In block 1106, processor 402 may identify a second device number in the cache line. The second device number may be read from another one of the plurality of memory devices. The second device number may correspond to a memory device associated with the detected error, and may be a copy of the first device number. The second device number may appear in a designated position in the cache line, as discussed above with respect to FIG. 4. In some implementations, tag encoding bits may be embedded in the second device number, and processor 402 may use a mask when identifying the second device number, in a manner analogous to that discussed above with respect to FIG. 3.

[0081] Next, in block 1108, processor 402 may determine whether the first device number matches the second device number. For example, processor 402 may input the first device number and the second device number into a comparator. If the first device number and the second device number match, method 1100 may proceed to block 1110, in which processor 402 may determine, using the device number, the position of the error in the cache line and correct the error, as discussed above with respect to FIG. 8. In block 1112, processor 402 may continue normal operations of memory controller 400.

[0082] If, in block 1108, it is determined that the first device number does not match the second device number, method 1100 may proceed to block 1114, in which processor 402 may use the first device number to determine the position of the error in the cache line and attempt to correct the error. As discussed above, the position of data in a cache line may correspond to the device number of the memory device from which the data was read. In block 1114, processor 402 may determine which bits in the cache line are in the position corresponding to the first device number, and may attempt to correct those bits (e.g., reconstruct the data in the position using data in other positions in the cache line).

[0083] Next, in block 1116, processor 402 may determine whether the correction attempted in block 1114 was successful. For example, processor 402 may run a checksum algorithm on the corrected bits and/or the entire cache line with the corrected bits. If it is determined that the correction was successful, method 1100 may proceed to block 1112.

- 21 -

[0084] If, in block 1116, it is determined that the attempted correction was not successful, method 1100 may proceed to block 1118, in which processor 402 may use the second device number to determine the position of the error in the cache line and attempt to correct the error. For example, processor 402 may determine which bits in the cache line are in the position corresponding to the second device number, and may attempt to correct those bits (e.g., reconstruct the data in the position using data in other positions in the cache line). Next, in block 1120, processor 402 may determine whether the correction attempted in block 1118 was successful. For example, processor 402 may run a checksum algorithm on the corrected bits and/or the entire cache line with the corrected bits. If it is determined that the correction was successful, method 1100 may proceed to block 1112.

[0085] If, in block 1120, it is determined that the correction attempted in block 1118 was not successful, method 1100 may proceed to block 1122, in which a system crash may occur due to the uncorrected error. The defective memory device associated with the error may need to be replaced for memory controller 400 to operate normally.

[0086] The foregoing disclosure describes inclusion of memory erasure information for a cache line in the cache line. Such memory erasure information in a cache line increases the granularity at which memory devices can be spared, for example up to the contribution of a single memory device to the cache line.

We claim:

1. A method for storing memory erasure information, the method comprising:

identifying a memory location associated with an error in a first cache line, wherein the first cache line comprises data read from the memory location, and wherein the memory location is in a first memory device of a plurality of memory devices on a memory module;

writing, to one of the plurality of memory devices, a device number corresponding to the first memory device; and

retrieving, when the memory location is read for a second cache line, the device number corresponding to the first memory device, wherein the second cache line comprises the retrieved device number and data read from the memory location.

2. The method of claim 1, further comprising:

determining, based on the retrieved device number, a position, in the second cache line, of the data read from the identified memory location; and
correcting the data in the determined position.

3. The method of claim 1, further comprising writing, to another one of the plurality of memory devices, a copy of the device number corresponding to the first memory device.

4. The method of claim 1, further comprising embedding tag encoding bits in the device number.

5. The method of claim 1, further comprising:

scrubbing the identified memory location in the first memory device;
and

reading data from the scrubbed memory location, wherein the device number corresponding to the first memory device is written to the one of

- 23 -

the plurality of memory devices if the data read from the scrubbed memory location is erroneous.

6. The method of claim 1, further comprising:

designating a region around the memory location as defective; and retrieving, when an address in the designated region is read, the device number corresponding to the first memory device.

7. A machine-readable storage medium encoded with instructions executable by a processor of a memory controller for storing memory erasure information, the machine-readable storage medium comprising:

instructions to identify a memory location associated with an error in a first cache line, wherein the first cache line comprises data read from the memory location, and wherein the memory location is in a first memory device of a plurality of memory devices on a memory module;

instructions to store, in one of the plurality of memory devices, a device number corresponding to the first memory device;

instructions to retrieve, when the memory location is read for a second cache line, the device number corresponding to the first memory device, wherein the second cache line comprises the retrieved device number and data read from the memory location; and

instructions to determine, based on the device number, a position of an error in the second cache line.

8. The machine-readable storage medium of claim 7, further comprising instructions to store, in another one of the plurality of memory devices, the device number corresponding to the first memory device.

9. The machine-readable storage medium of claim 7, further comprising instructions to embed tag encoding bits in the device number.

10. The machine-readable storage medium of claim 7, further comprising:

instructions to designate a region around the memory location as defective; and

instructions to retrieve, when an address in the designated region is read, the device number corresponding to the first memory device.

11. A machine-readable storage medium encoded with instructions executable by a processor of a memory controller for correcting cache line errors, the machine-readable storage medium comprising:

instructions to detect an error in a cache line, wherein the detected error is associated with a first memory device of a plurality of memory devices on a memory module;

instructions to identify, in the cache line, a first device number, wherein the first device number is read from one of the plurality of memory devices;

instructions to determine, based on the first device number, a position of the detected error in the cache line; and

instructions to correct the detected error.

12. The machine-readable storage medium of claim 11, further comprising instructions to identify, in the cache line, a second device number, wherein the second device number is read from another one of the plurality of memory devices.

13. The machine-readable storage medium of claim 12, further comprising:

instructions to compare the first device number and the second device number;

instructions to determine, if the first device number and the second device number do not match, and if the detected error is not successfully corrected when the first device number is used to determine the position of the detected error, the position of the detected error in the cache line, based on the second device number.

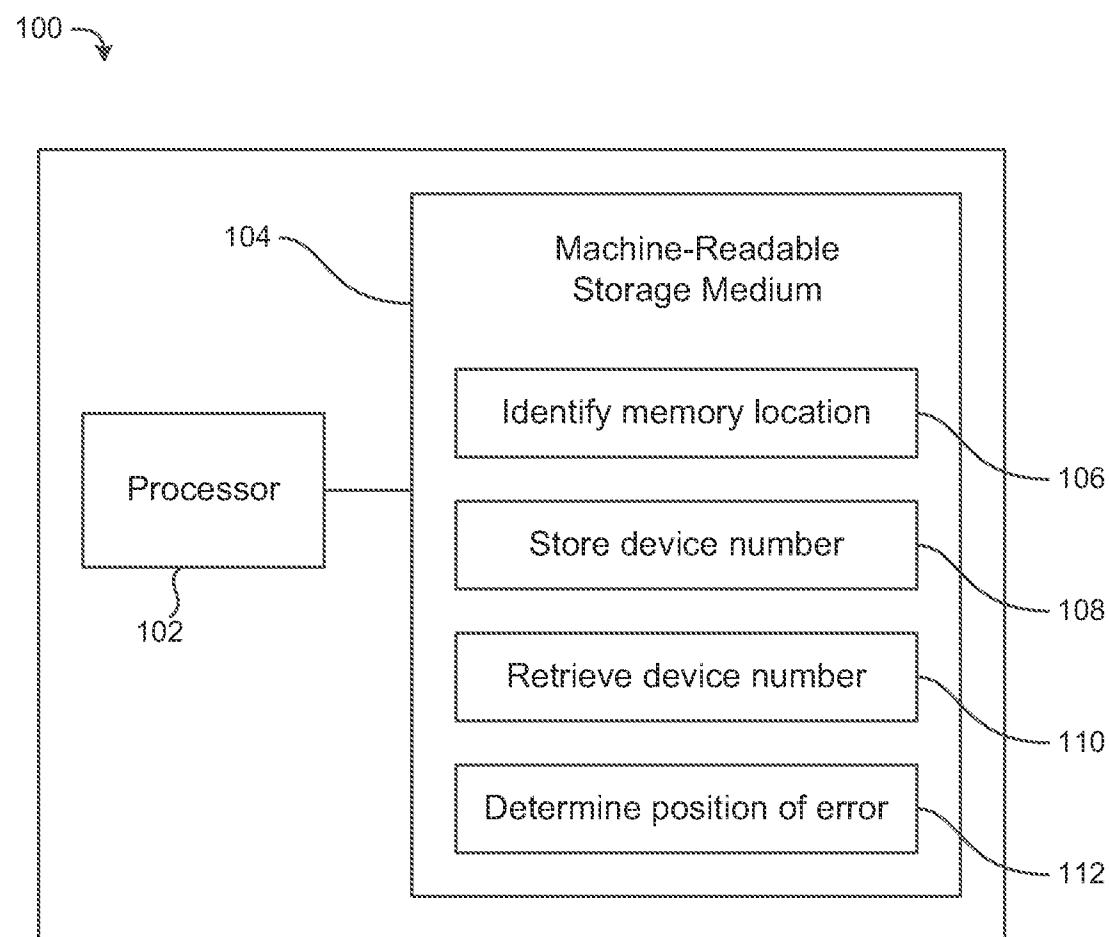
- 25 -

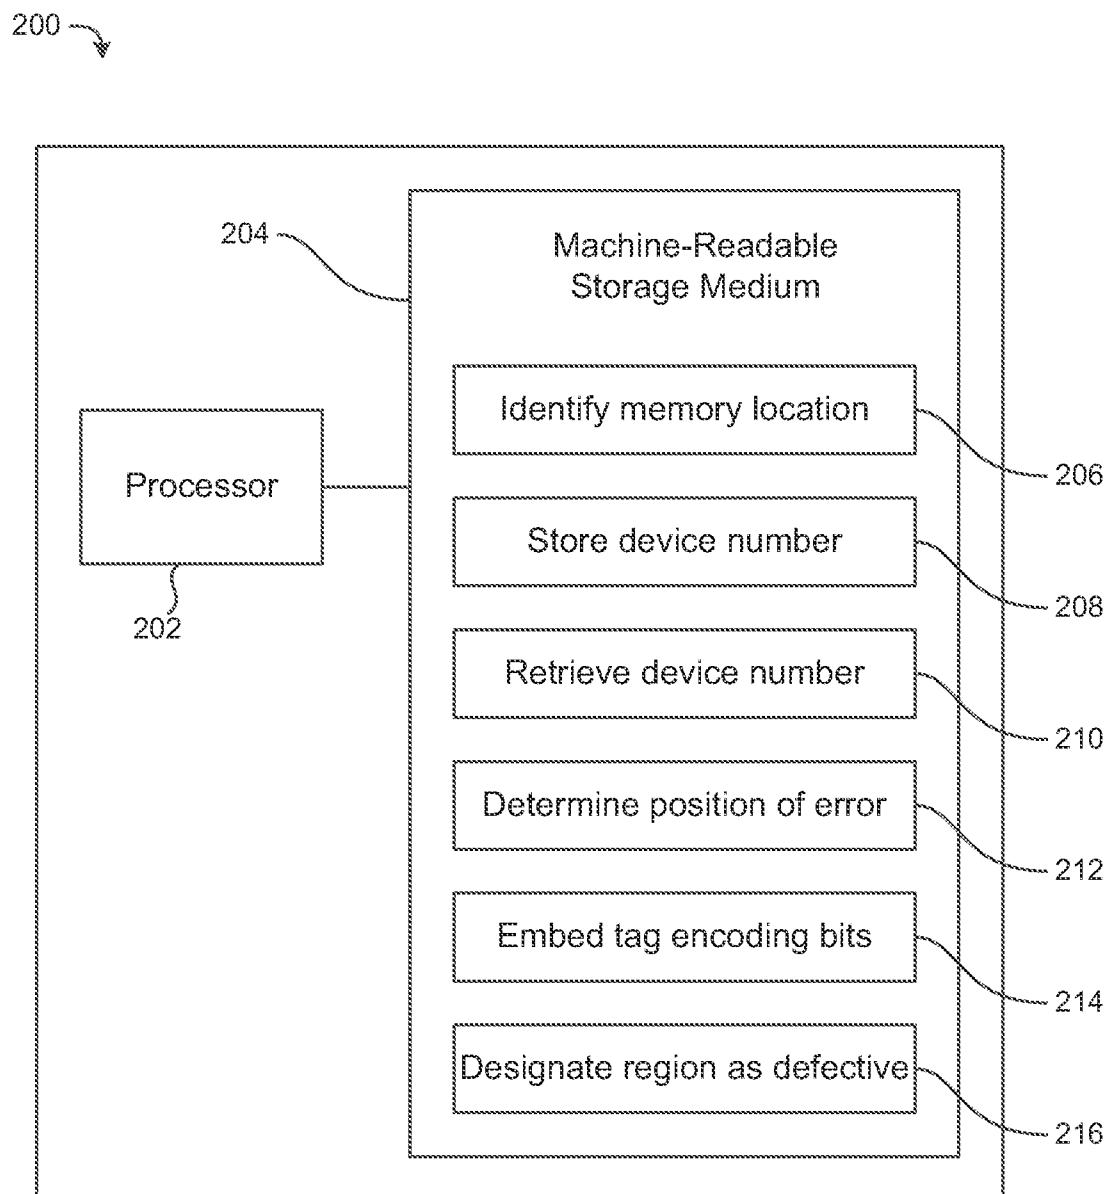
14. The machine-readable storage medium of claim 11, wherein tag encoding bits are embedded in the first device number.

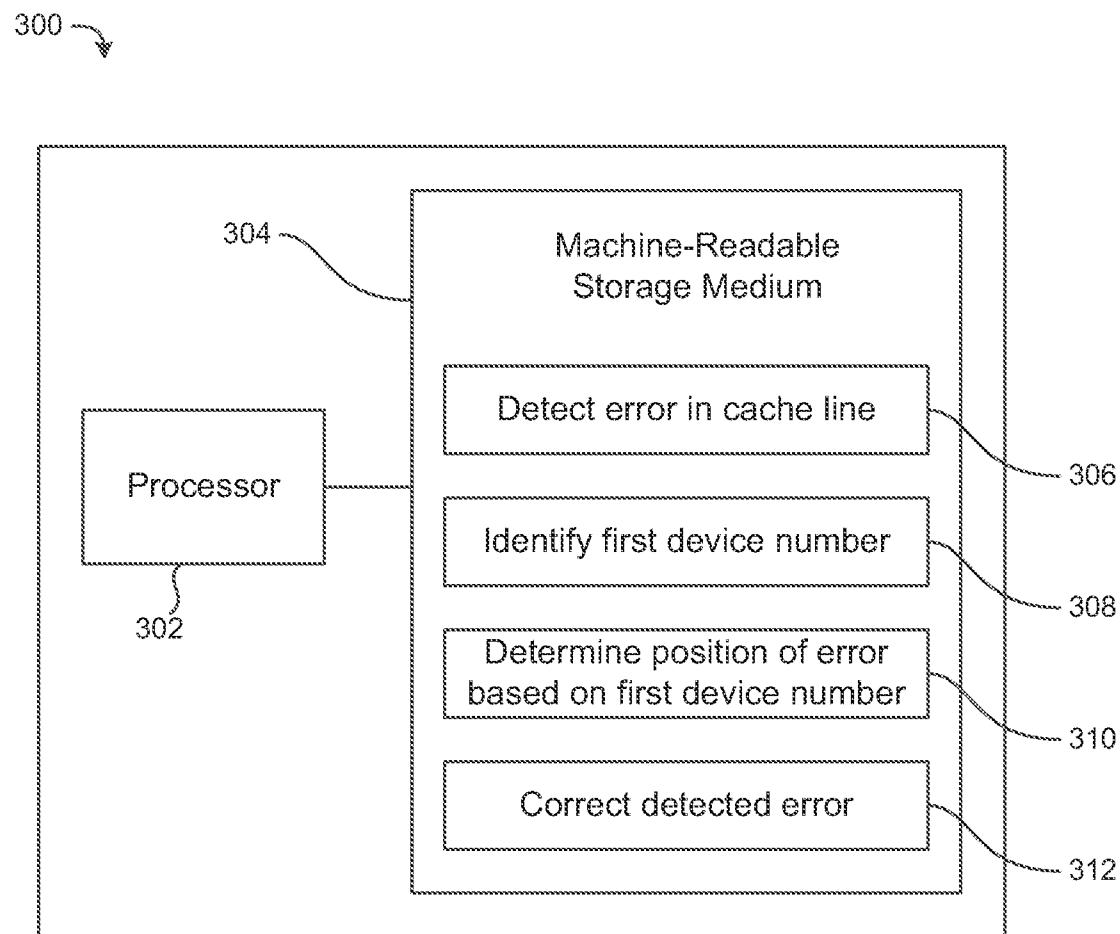
15. The machine-readable storage medium of claim 11, further comprising:

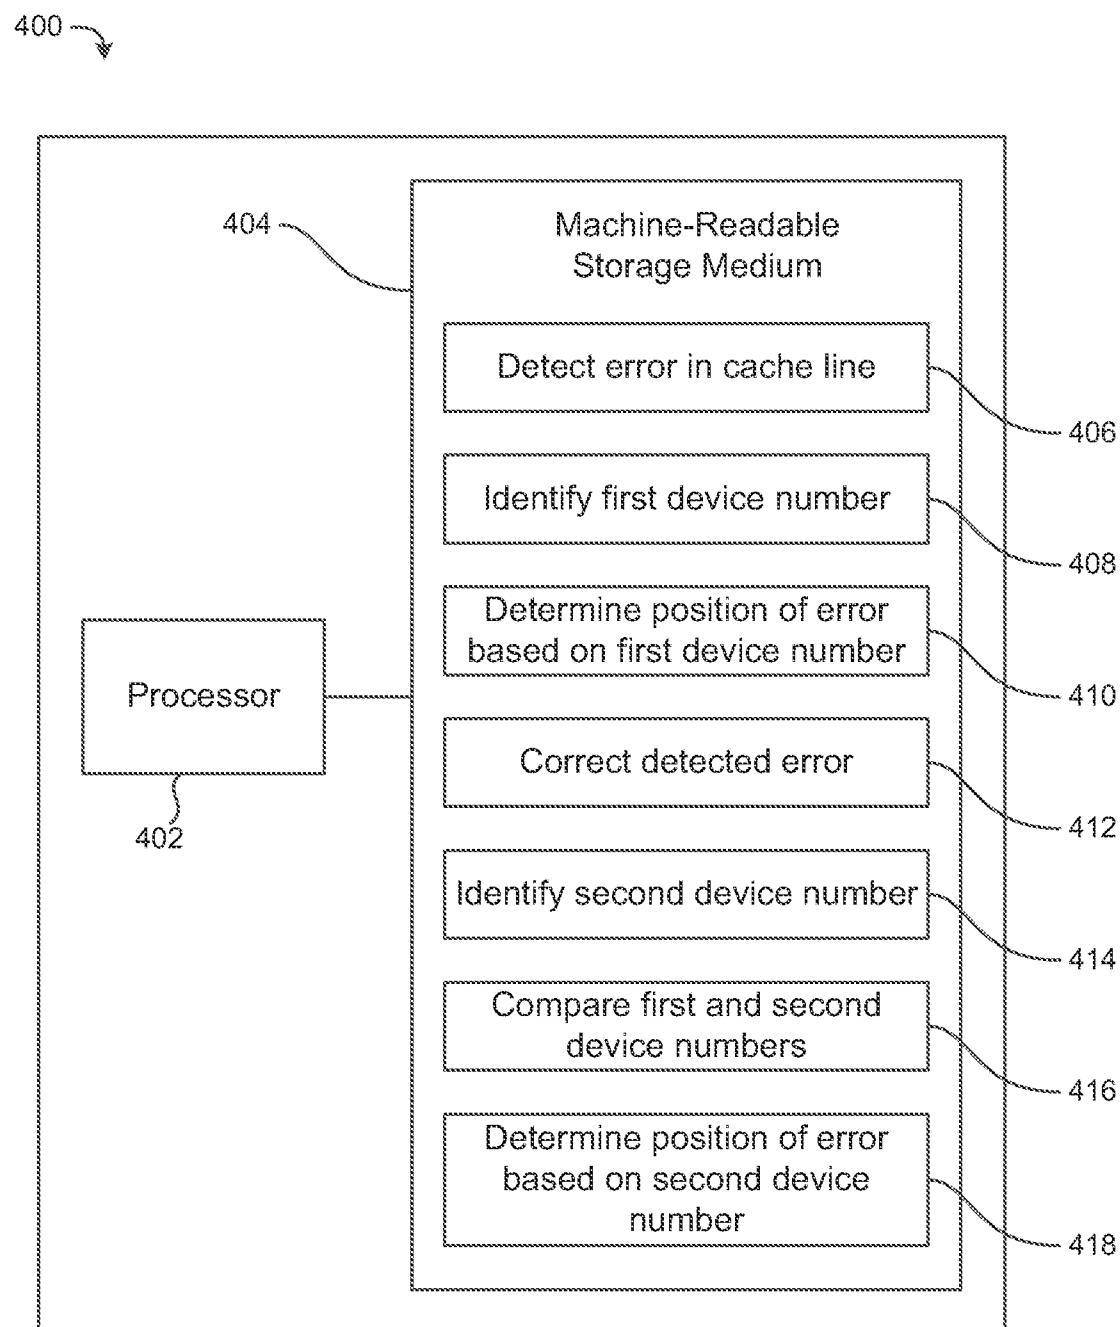
instructions to keep track of how many errors are associated with a region of the first memory device; and

instructions to retrieve, when an address in the region is read, the device number corresponding to the first memory device.

**FIG. 1**

**FIG. 2**

**FIG. 3**

**FIG. 4**

500

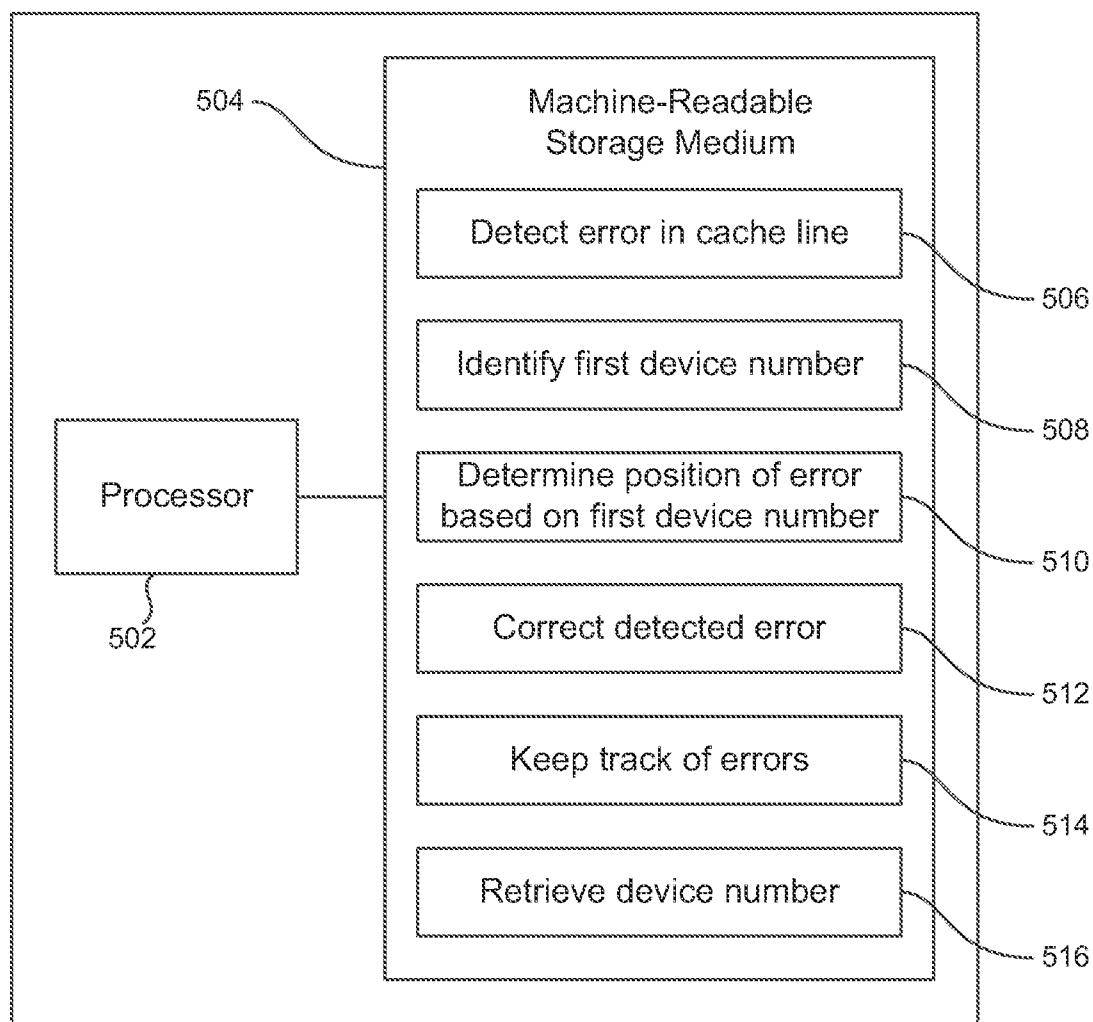


FIG. 5

600 ↘

Parity (4 bits)	DRAM # (4 bits)	DRAM # (4 bits)	Data (60 bits)	
Parity (4 bits)			Data (68 bits)	
Parity (4 bits)	DRAM # (1 bit) CRC (3 bits)	DRAM # (1 bit) CRC (3 bits)	Data (60 bits)	
Parity (4 bits)			Data (68 bits)	
Parity (4 bits)	CRC (4 bits)	CRC (4 bits)	Data (60 bits)	
Parity (4 bits)			Data (68 bits)	
Parity (4 bits)	CRC (4 bits)	CRC (4 bits)	Data (60 bits)	
Parity (4 bits)			Data (68 bits)	

1 DRAM 1 DRAM 1 DRAM

15 DRAM

FIG. 6

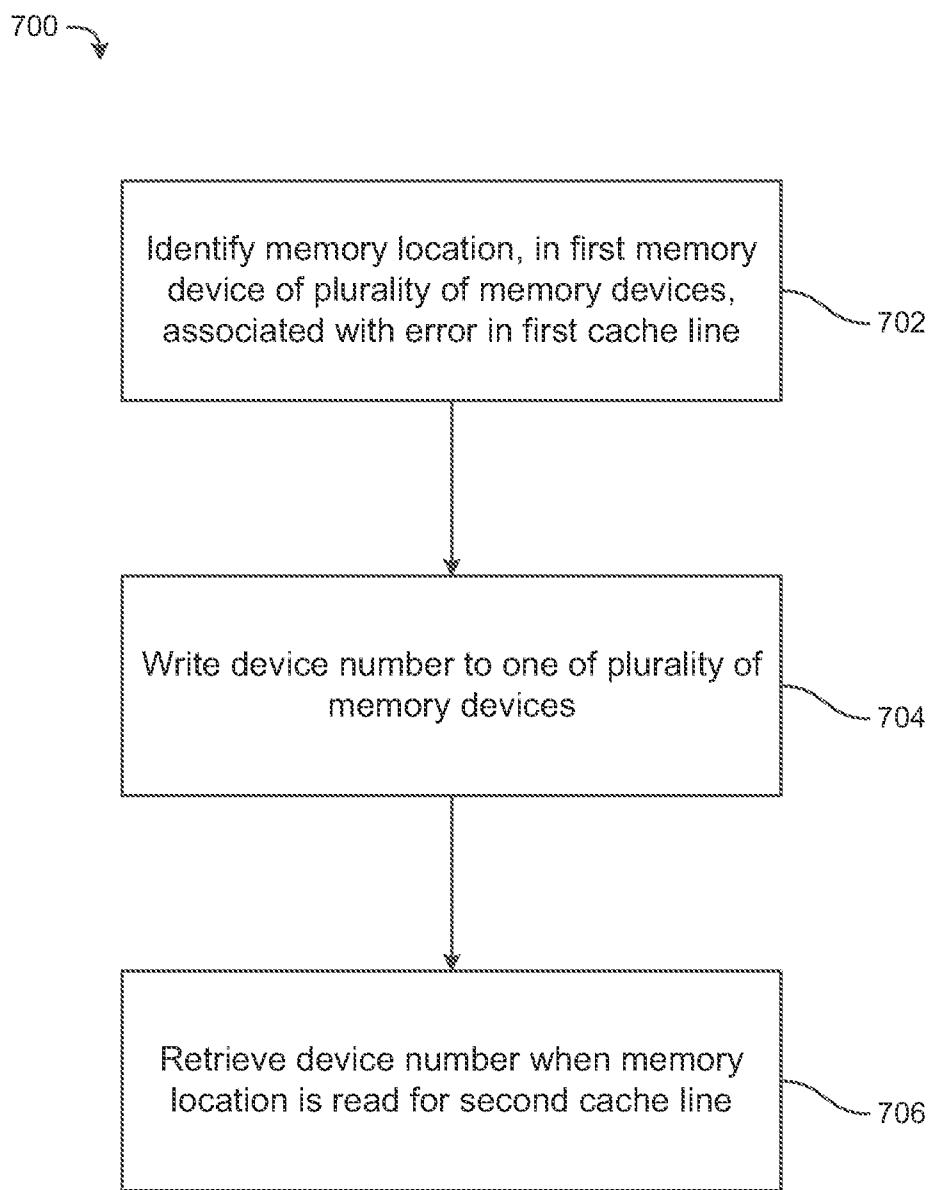
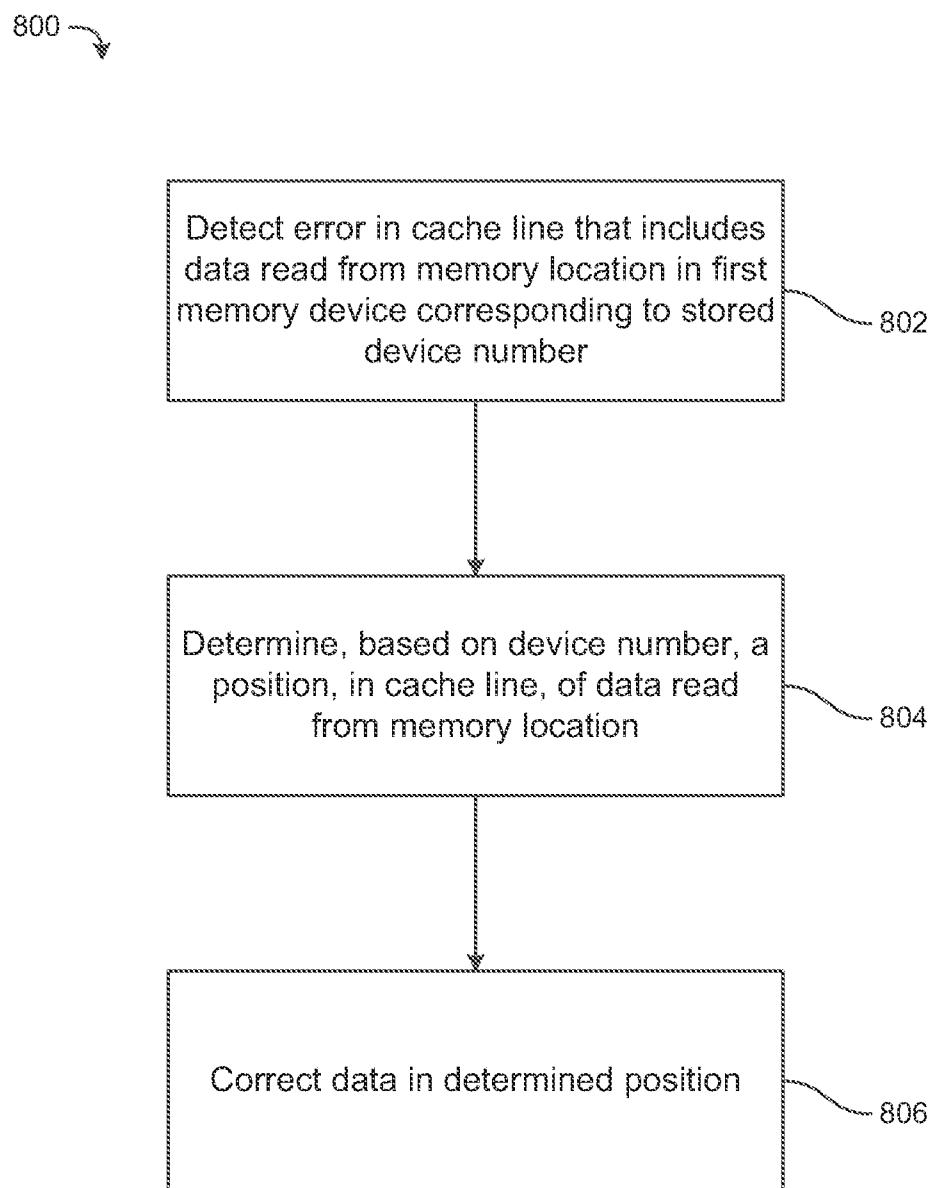


FIG. 7

**FIG. 8**

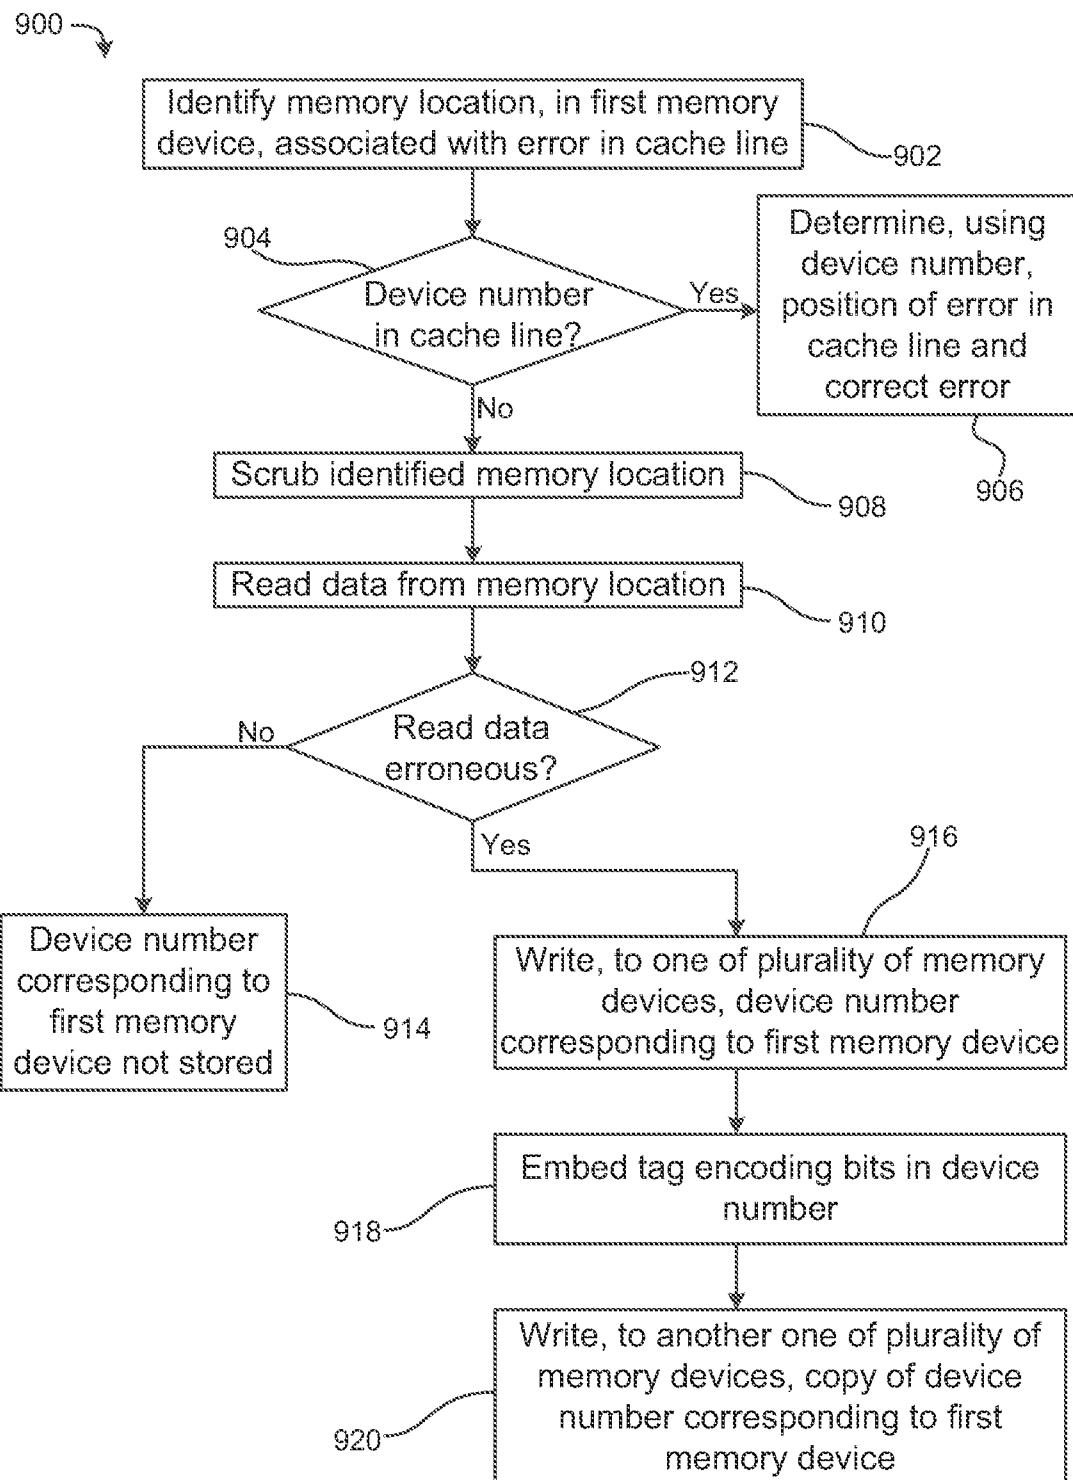


FIG. 9

1000

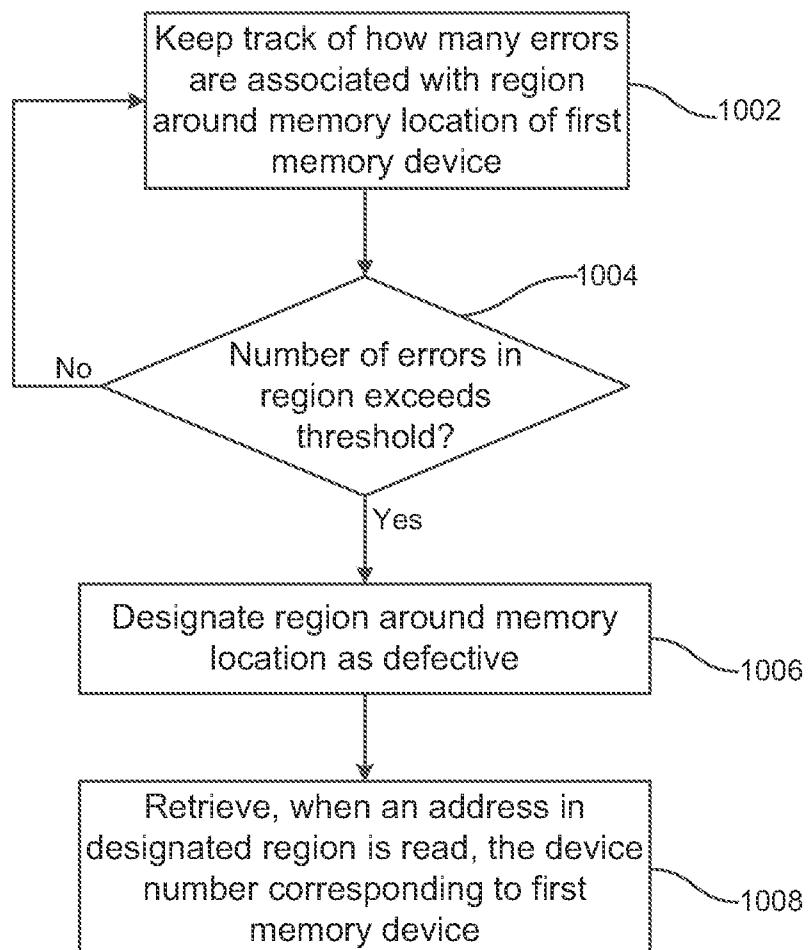


FIG. 10

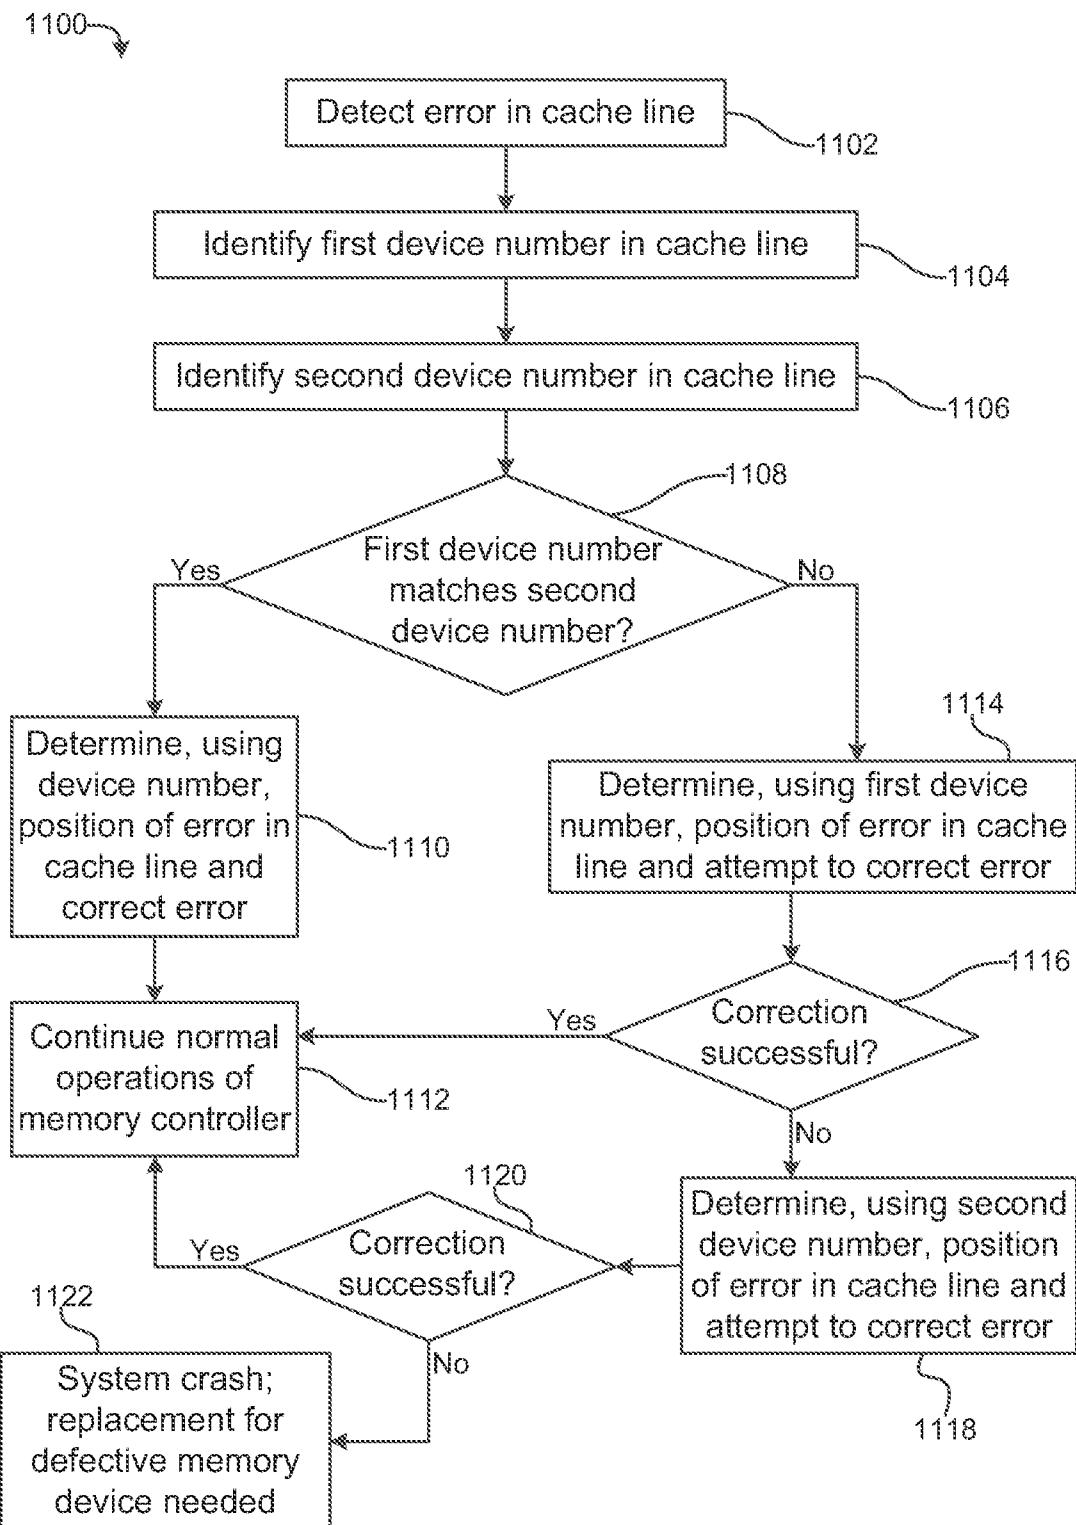


FIG. 11

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/073871

A. CLASSIFICATION OF SUBJECT MATTER

G06F 11/00(2006.01)i, G06F 12/08(2006.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 11/00; G06F 11/07; G06F 3/00; G06F 11/10; G06F 11/22; G06F 13/00; G06F 12/08

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: memory, error, identify, device, number, location, cache, line

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2007-0050688 A1 (LARRY JAY THAYER) 01 March 2007 See paragraphs 11-12, 26; and figure 1.	1-15
A	US 2004-0025094 A1 (SHAWN KENNETH WALKER et al.) 05 February 2004 See paragraphs 4, 19, 24; and figure 2.	1-15
A	US 2010-0293305 A1 (JONGGYU PARK et al.) 18 November 2010 See paragraphs 11-12; and figure 4A.	1-15
A	EP 1000395 B1 (INTERGRAPH HARDWARE TECHNOLOGIES COMPANY) 01 December 2004 See paragraphs 2, 4, 12; and figure 1.	1-15
A	US 2010-0268984 A1 (GUY L. GUTHRIE et al.) 21 October 2010 See paragraphs 7-8; and figure 1.	1-15

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family
--	--

Date of the actual completion of the international search
05 September 2014 (05.09.2014)

Date of mailing of the international search report

05 September 2014 (05.09.2014)Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea
Facsimile No. +82-42-472-7140

Authorized officer

YU, JAE CHON

Telephone No. +82-42-481-8647



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/073871

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
us 2007-0050688 A1	01/03/2007	GB 0616145 D0 GB 2429805 A GB 2429805 B US 2008-0065933 A1 US 7307902 B2 US 7599235 B2	20/09/2006 07/03/2007 22/12/2010 13/03/2008 11/12/2007 06/10/2009
us 2004-0025094 A1	05/02/2004	JP 2001-249847 A US 6591393 B1 US 6874116 B2	14/09/2001 08/07/2003 29/03/2005
us 2010-0293305 A1	18/11/2010	KR 10-2010-0122709 A US 8359412 B2	23/11/2010 22/01/2013
EP 1000395 B1	01/12/2004	DE 69827949 D1 DE 69827949 T2 EP 1000395 A1 US 6158025 A WO 1999-005599 A1	05/01/2005 27/10/2005 17/05/2000 05/12/2000 04/02/1999
us 2010-0268984 A1	21/10/2010	US 8291259 B2	16/10/2012