



- (51) International Patent Classification:
G06F 12/08 (2016.01)
- (21) International Application Number:
PCT/US2020/022907
- (22) International Filing Date:
16 March 2020 (16.03.2020)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/820,144 18 March 2019 (18.03.2019) US
- (71) Applicant: **RAMBUS INC.** [US/US]; 1050 Enterprise Way, Suite 700, Sunnyvale, CA 94089 (US).
- (72) Inventors: **WARE, Frederick**; 13961 Fremont Pines Lane, Los Altos Hills, CA 94022 (US). **VOGELSANG, Thomas**;

1050 Enterprise Way, Suite 700, Sunnyvale, CA 94089 (US). **MILLER, Michael**; 1050 Enterprise Way, Suite 700, Sunnyvale, CA 94089 (US). **WILLIAMS, Collins**; 1050 Enterprise Way, Suite 700, Sunnyvale, CA 94089 (US).

(74) Agent: **OVANEZIAN, Daniel, E.** et al.; Womble Bond Dickinson (US) LLP, P.O. Box 7037, Atlanta, GA 30357-0037 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, 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,

(54) Title: SYSTEM APPLICATION OF DRAM COMPONENT WITH CACHE MODE




FIG. 13A

(57) Abstract: Disclosed is a memory system that has a memory controller and may have a memory component. The memory component may be a dynamic random access memory (DRAM). The memory controller is connectable to the memory component. The memory component has at least one data row and at least one tag row different from and associated with the at least one data row. The memory system is to implement a cache having multiple ways to hold a data group. The memory controller is operable in each of a plurality of operating modes. The operating modes include a first operating mode and a second operating mode. The first operating mode and the second operating mode have differing addressing and timing for accessing the data group. The memory controller has cache read logic that sends a cache read command, cache results logic that receives a response from the memory component, and cache fetch logic.



SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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, ST, 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 applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

SYSTEM APPLICATION OF DRAM COMPONENT WITH CACHE MODE**REFERENCE TO RELATED APPLICATION**

[001] This application is related to, and claims the benefit of, U.S. Provisional Patent Application No. 62/820,144 filed 18 March 2019, titled “SYSTEM APPLICATION OF DRAM COMPONENT WITH CACHE MODE”, the entire contents of which are hereby incorporated by reference herein.

BACKGROUND

[002] Dynamic Random Access Memory (DRAM) is in use in a wide variety of computers and computing devices as system memory. Now that DRAM scaling is slowing down, storage class memory (SCM) with lowered costs and DRAM will be useful to continue growing memory system capacity. Hybrid memory systems combining SCM with DRAM cache may be useful due to the longer latency and limited endurance of SCM compared to DRAM.

BRIEF DESCRIPTION OF THE DRAWINGS

[003] The present disclosure will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the disclosure.

[004] **FIG. 1A** illustrates an embodiment of a DRAM with data rows and tag rows, which can be operated as a multiway set associative cache, or multiway, set associative cache.

[005] **FIG. 1B** illustrates tags in a separate tag mat, as suitable for a version of the DRAM embodiment in **FIG. 1A**.

[006] **FIG. 2A** illustrates a further embodiment of a DRAM with cached data of a set distributed across multiple banks.

[007] **FIG. 2B** illustrates data rows and tag rows, as suitable for a version of the DRAM embodiment in **FIG. 2A**.

[008] **FIG. 2C** illustrates distributed tags, as suitable for a version of the DRAM embodiment in **FIG. 2A**.

[009] **FIG. 3A** illustrates circuitry and addressing for a cache hit and access to data of a set that has been cached in a way of the multiway set associative cache, in one embodiment.

[0010] **FIG. 3B** illustrates the use of set bits from tag data in column decode to access cached data pointed to by the matched tag.

[0011] **FIG. 3C** illustrates address spaces in SCM and DRAM, for accessing data of SCM that is cached in the multiway set associative cache DRAM, according to one embodiment.

[0012] **FIG. 4A** is a read timing detail illustrating a normal access of data and a modified

access of cached data, in one embodiment.

[0013] FIG. 4B illustrates reading cached data from one of the banks of the multiway set associative cache DRAM, while more of the cached data is read in parallel from other banks, using the read timing shown in FIG. 4A.

[0014] FIG. 5A is a write timing detail illustrating a normal access of data and a modified access of cached data, in one embodiment.

[0015] FIG. 5B illustrates writing cache data to one of the banks of the multiway set associative cache DRAM, while more of the cache data is written in parallel to other banks, using the write timing shown in FIG. 5A.

[0016] FIG. 6 illustrates DRAM cells and sense amplifiers in a mat, which can be used in embodiments of the DRAM.

[0017] FIG. 7 illustrates a mode register for selection of standard DRAM mode, or multiway set associative cache mode with adjustable bit length of tags and comparators.

[0018] FIG. 8A is a flow diagram of a method of operating a multiway set associative DRAM cache, which can be practiced by various embodiments described herein.

[0019] FIG. 8B is a flow diagram of selecting an operating mode, which can be practiced by various embodiments described herein.

[0020] FIG. 9 depicts an embodiment of the multiway set associative DRAM cache.

[0021] FIG. 10 depicts a further embodiment of the multiway set associative DRAM cache.

[0022] FIG. 11 depicts a DRAM, and a memory controller that can operate the DRAM as a multiway set associative DRAM cache in accordance with present embodiments.

[0023] FIG. 12 depicts a controller using a DRAM with cache mode in accordance with present embodiments.

[0024] FIG. 13A depicts system configuration example X-A, in which the data ways are held within multiple banks of multiple devices, and there is overlapped tag/data access.

[0025] FIG. 13B depicts system configuration example Y-A, in which the data ways are held within multiple banks of a single device, and there is overlapped tag/data access.

[0026] FIG. 13C depicts system configuration example Z-A, in which the data ways are held within a single bank, and there is overlapped tag/data access.

[0027] FIG. 13D depicts system configuration example Z-B, in which the data ways are held within a single bank, and there is non-overlapped tag/data access.

[0028] FIG. 14A depicts address field steering in the controller of FIG. 12 for system configuration example X-A as shown in FIG. 13A.

[0029] FIG. 14B depicts address field steering in the controller of FIG. 12 for system configuration example Y-A as shown in FIG. 13B.

- [0030] FIG. 14C depicts address field steering in the controller of FIG. 12 for system configuration example Z-A as shown in FIG. 13C.
- [0031] FIG. 14D depicts address field steering in the controller of FIG. 12 for system configuration example Z-B as shown in FIG. 13D.
- [0032] FIG. 15A depicts command format for system configuration example Z-B as shown in FIG. 13D.
- [0033] FIG. 15B depicts command format for system configuration examples Z-A, Y-A and X-A as shown in FIGS. 13C, 13B and 13A.
- [0034] FIG. 15C depicts command format for system configuration examples Z-B, Z-A, Y-A and X-A as shown in FIGS. 13D, 13C, 13B and 13A.
- [0035] FIG. 16 depicts a read timing detail for a cache miss, with fetch and evict, in a DRAM with cache mode.
- [0036] FIG. 17 is a read timing detail for a 64 byte read of a DRAM with cache mode, for a normal (direct) access, and a modified (cached) access.
- [0037] FIG. 18 is a flow diagram depicting controller sequencing for a DRAM cache, experiencing a cache miss.
- [0038] FIG. 19 is a system action diagram depicting the controller using DRAM with cache mode, performing the actions of FIG. 18.
- [0039] FIG. 20 depicts a touch flag enhancement for the eviction decision process.
- [0040] FIG. 21A depicts a 4 kB cache line distributed across banks/channels/devices in accordance with present embodiments.
- [0041] FIG. 21B depicts 4 kB cache line tags and data for 16 sets (ways) located in the same DRAM row in accordance with present embodiments.

DETAILED DESCRIPTION

[0042] In the following description, various aspects of the illustrative embodiments will be described using terms commonly employed by those skilled in the art to convey the substance of their work to others skilled in the art. For purposes of explanation, specific numbers, materials and configurations are set forth in order to provide a thorough understanding of the illustrative embodiments. However, it will be apparent to one skilled in the art that the present disclosure may be practiced with only some of the described aspects or without the specific details. In other instances, well-known features are omitted or simplified in order not to obscure the illustrative embodiments.

[0043] DRAM architecture and controller modifications described herein enable tag storage and comparison on the DRAM, and access to cached data. In some versions, the DRAM is

operable as a multiway set associative cache, and also operable as a standard DRAM, e.g., if not part of a hybrid system with storage class memory or other memory. Sets of data from storage class memory or other memory can be cached in ways in the DRAM cache and accessed when a tag comparison indicates a cache hit.

[0044] In typical systems, circuitry outside of the DRAM die performs the tag matching. Using standard DRAM parts as cache in a hybrid memory system includes transferring tag information bits from the DRAM to the chip doing the tag matching prior to being able to access the data information on the DRAM. This creates both a latency overhead and a power overhead compared to modifying the DRAM to enable tag matching on the DRAM, as described herein.

[0045] Various embodiments of the DRAM cache have some or all of the following features.

[0046] Transfers to and from storage class memory (SCM) are typically done using very large block sizes. This makes the ratio of tag to data much smaller than in conventional caches, enabling DRAM modifications for cache with little overhead. Various embodiments of a DRAM with data rows and tag rows, and data mat rows and tag mat rows, are shown and described with reference to FIGS. 1-11, and further embodiments with further arrangements of data rows, tag rows, data mat rows and tag mat rows are readily devised for various ratios of tag to data. A controller and DRAM implementing a multiway set associative cache are described in multiple configurations in FIGS. 12-21B.

[0047] The DRAM cache is organized as a set associative cache. In a set associative cache, a data set from storage class memory could be cached in a set, and in the set in any of multiple ways, or locations, in cache memory. Multiplying the total number of sets by the total number of ways in each set gives the total number of lines in cache memory. Each way, and each set, maps to a line. A tag matches part of an address of a data set, when the data set is cached, and points to the set and the way or location in cache memory, so that the tagged, cached data set can be read from cache memory. The tag associates the data set to the set and the way or location in cache memory, thus the name for this type of cache. The amount of data that can be cached for a data set, the number of sets, and the number of ways or locations in which that amount can be cached, generally denoted N , determines the amount of cache memory for the N way set associative cache and the number of address bits used in total and for each tag.

[0048] Tag storage in the DRAM is done by adding additional rows of cells, either distributed in the existing mats or by adding mats. A mat, whether a data mat with data rows, a tag mat with tag rows, or a mat that has both data rows and tag rows, is a group of DRAM cells of a defined height (number of cells or rows) and width (number of cells or columns). Each DRAM cell has a transistor (generally, NMOS or N type metal oxide semiconductor) and a capacitor. Various embodiments of mats are shown and described with reference to FIGS. 1A-6,

and further arrangements of DRAM cells and mats are readily devised.

[0049] Associating tag rows with data rows that do not share a primary sense-amplifier avoids the need to double sense-amplifier stripes and saves area. A tag row in a mat with one sense amplifier is associated to and points to a data location elsewhere in DRAM that has another sense amplifier, in some embodiments. This supports timing overlap in the tag sense, through the sense amplifier working for the tag row, and the data sense, through the sense amplifier working for the data location associated to the tag row, as shown in timing details of FIGS. 4 and 5. Such timing overlap would require doubling the sense amplifiers if the tag row and associated data location were in the same mat.

[0050] Tags of different ways within a set are located on one row. A tag hit activates a tag match to one of the ways located on this one row, as shown and described with reference to FIG. 3A. This feature supports compare blocks, one for each tag and tag match, to be ganged together in a logical OR configuration for compact circuitry in the compare block shown in FIG. 3A.

[0051] Tag comparison is done with column accesses only. Tags for all of the ways of a specific set are on the same row, and are compared to address bits in parallel across the column accesses of the tag data for that row, as shown in FIGS. 3A and 3B.

[0052] Parallel access at the edge of a bank is used to maximize parallel tag comparison. Tags are compared at the edge of each bank, so that multiple tag comparisons can occur across multiple banks in parallel. Cached data is then accessed in multiple banks in parallel, in some embodiments, for example as shown in FIGS. 1A-2C.

[0053] Addressing a set uses device, bank and row addresses. Data from storage class memory can be cached in one of N ways or locations, in one of the sets, in the N way set associative cache. The address for set includes a device address, a bank address, and a row address, as decoded through address decoding.

[0054] Set addresses, compare values and addresses determined by tag matching are distributed in the DRAM using the internal address busses. See, for example, FIG. 9 showing cache logic and multiplexing for distributing addresses to row decoding and column decoding for banks in DRAM.

[0055] Internal data busses outside of the array and DQ pins are used for data only. See, for example, FIG. 9 showing internal data from banks in the DRAM to data logic. Address and data are not mixed or multiplexed on the data busses.

[0056] **FIG. 1A** illustrates an embodiment of a DRAM with data mat rows 104 (each with multiple data rows) and tag mat rows 106 (each with multiple tag rows for tag information bits), which can be operated as a multiway set associative cache. In this example, each bank 102 of DRAM is an array of 32 mats 108 by 64 mats 108 of DRAM, and each mat 108 is a 512 bit by

512 bit array or sub-array (see FIG. 6), for a total of 256 kilobits (kb) per mat 108. There are eight banks 102 in each of channels A and B of DRAM 0, and eight more banks 102 in each of channels C and D in DRAM 1, for a total of 32 banks 102. The DRAM can perform 64 byte read or write access to 4 kilobyte (kB) DRAM blocks (cache lines). Tags, in tag rows in 32 mats 108 forming a tag mat row 106 at the bottom (or the top, in further embodiments) of the bank 102, take up about 1.6% extra memory cell area in this example.

[0057] A DRAM bank 102 is built from an array of mat blocks (sub-array) containing memory cells and row/column peripheral logic. Data wires connect the bank 102 to the interface block. A single row of unmodified mats provides enough capacity and bandwidth (BW) for the tags.

[0058] FIG. 1B illustrates tags in a separate tag mat 112, as suitable for a version of the DRAM embodiment in FIG. 1A. A tag read bit determines whether cached data or tag data is read during a read operation. If the tag read bit is 1, data of a tag in the tag mat 112 is read. If the tag read bit is 0, cached data is read from a data mat 110.

[0059] In various embodiments, including for tags as shown in FIGS. 1A, 1B, 2A, 2B and 2C, the tags can be read using an access like a normal access (no tag match, direct access from the memory controller to the row address to be activated). In some embodiments, there needs to be an additional address space that signals tag access instead of a normal access. Internally this can be set by a tag read bit as FIG. 1B shows for tags in a separate mat 112, or distributed in mats as shown in FIGS. 2B and 2C, etc. There are multiple options to do this:

1. When DRAM is in cache mode, the bit assignment to decode a sequence of bits sent over the CA bus is modified to have an additional state that signals tag access.
2. When DRAM is in cache mode and there is a bidirectional signal like DMI used to transfer additional information, this signal can be used to switch between cache access and tag read.
3. A mode register set (MRS) switches between cache access and tag read.

[0060] FIG. 2A illustrates a further embodiment of a DRAM with cached data of a set distributed across multiple banks 202. In this example, the set is distributed across four banks 202, at 32 kB per bank 202, and each way of the set is distributed as 16 bytes per bank 202 in each of the same four banks 202. Each of these same four banks 202 has 32 tags. Other sets are distributed across other banks 202 of the DRAM, at four banks 202 per set. A tag match, in a tag in each of the four banks 202 occupied by the set, points to cached data of a way of the set, in each of the four banks 202.

[0061] FIG. 2B illustrates data rows 204 and tag rows 206, as suitable for a version of the DRAM embodiment in FIG. 2A. This, and other numbers and arrangements of data rows and tag rows in a DRAM as readily devised in keeping with the teachings herein, can be operated as

a multiway set associative cache. In this example, each bank 202 of DRAM is an array of eight mats 208 by 64 mats 208, and each mat 208 is a 512 bit by 1024 bit array or sub- array (see FIG. 6), for a total of 512 kb per mat 208. In this example, each mat 208 includes redundant rows at top and bottom, a tag row next to the upper redundant rows, or alternatively next to the lower redundant rows, and regular rows (i.e., data rows), and uses 32 rows of data per set. In various embodiments, a tag row 206 could be at the top, the bottom, or elsewhere in the mat 208.

Various arrangements for tag, redundant and regular data rows are possible. In some embodiments, tags in the lower half of a segment (or bank 202) are assigned to data in the upper half of the segment (or bank 202). Tags in the upper half of the segment (or bank 202) are assigned to data in the lower half of the segment (or bank 202). There are 16 banks 202 in each of channels A and B of DRAM 0, and 16 more banks 202 in each of channels A and B of DRAM 1. 32 tags in each of four banks 202 in DRAM 0 can cache a 4 by 32 kB set with 4 by 16 bytes of hit data.

[0062] FIG. 2C illustrates distributed tags, as suitable for a version of the DRAM embodiment in FIG. 2A. Regular wordlines (depicted in solid lines) and tag wordlines (depicted in dashed lines) are present in mats 210. Data of a regular wordline can be read when the tag read bit is 0. Data of a tag wordline can be read when the tag read bit is 1. In this embodiment, in a given mat 210, there could be multiple regular wordlines, for cached data, and a single tag wordline, for tag data.

[0063] With reference to FIGS. 1A, 1B and 2A-2C, various embodiments have some or all of the following architectural details.

[0064] Optimizing the distribution of data units corresponding to one way in the multiway cache (e.g. 64B transactions within 4kB way) can be based on tradeoffs between access speed of tag matching, random data access to single unit and access to whole 4kB blocks vs. bank utilization to allow continuous data streams. For example, the multiway set associative cache DRAM embodiment in FIG. 1A distributes cached data for the ways over all of the banks. The multiway set associative cache DRAM embodiment in FIG. 2A distributes cached data for a given way over only a few banks. One version, in FIG. 1A or variations thereof, can have simultaneous row activation for tags and data, another version, in FIG. 2A or variations thereof can wait until a comparison has a match, to activate a row for cached data access. One version, in FIG. 1A, favors random access to data, another version, in FIG. 2A, favors bank utilization and continuous streams.

[0065] Associating the way address with a column address allows full overlap of tag row activation and data row activation but requires either more banks to be accessed or multiple row accesses in the same bank when accessing a full way for fill, for example fetch and evict. Fill

involves bringing data into the cache and then getting rid of dirty data at the end.

[0066] Tags are replicated as needed so that tag matching can be done locally to where the match information is needed. Multiple banks each have copies of tags for the ways in which data could be cached in that bank, so that whichever bank(s) actually have the tag match can then produce the data. See, for example, FIG. 2A showing tag hits in four banks of the multiway set associative cache DRAM. Each bank has a tag match for a way, and the cached data for that way in all four banks is read in parallel. Each bank has data for all of the ways of a set, but does not have the complete data for each way. In FIG. 1A, tags are replicated 32 times (each of 32 banks has a full, duplicate set of tags) and all of the cached data for a way is available at once, through parallel access to data of a way of a set, in all 32 banks. In FIG. 2A, tags are replicated four times (each of four banks has a duplicate set of a group of tags, each other group of four banks has another duplicate set of another group of tags, etc.), and these four banks are accessed multiple times in order to obtain all of the cached data for the way.

[0067] Compare logic is located close to the array edge. For example, the comparators and the compare block (see FIG. 3) are located along an edge of an array of mats of DRAM cells in a bank. Such location shortens signal paths, keeping parasitic signal line loading low, e.g., reducing resistance and capacitance, possibly also reducing incidences of signal reflections.

[0068] Feedback to a controller (e.g. hit or miss) is done using available feedback pins (alert (ALERT_n), data mask inversion (DMI), etc.) For example, feedback to the controller can be done using available pins that have a DRAM to controller direction and are not needed in the case of using the DRAM in its cache mode. Examples are using the DMI pins in case of modifying a LPDDR4 DRAM as cache DRAM (see FIG. 11) or using the ALERT_n pin in case of modifying a DDR4 DRAM. A further example is using the error detection code (EDC) pin in the case of modifying a GDDR5 or GDDR6 DRAM.

[0069] Writing dirty bit and writing data can start at the same time. The dirty bit is written along with tag data, and the data to be cached is written to the location (or way) in cache memory associated with the tag. If the same bank in the DRAM is used for tag data and cached data, two column cycles may be required pushing out closing of the tag row by one column cycle.

[0070] Redundancy: tag rows could get priority replacement or be always doubled up.

[0071] DRAM is dual use and can be set in one mode where it works as cache DRAM (e.g., multiway set associative cache mode) or in another mode as standard DRAM (e.g., standard DRAM mode). An embodiment of a mode register, to select standard DRAM mode versus cache DRAM mode, is shown and described with reference to FIG. 7. Mode selection controls address multiplexing and data multiplexing, which differ between these two modes.

[0072] Implementation with changes to periphery logic only, from a standard DRAM design layout, is possible if one redundant row per mat is assigned to be a tag row. Standard DRAM design layout has multiple redundant rows, and one of the redundant rows could be designated a tag row, without changing the design layout of the arrangement of mats or rows. Periphery logic can then implement changes to address decoding and data multiplexing, and the addition of comparators and compare blocks for the N way set associative cache operation of the tag rows and remainder of DRAM.

[0073] The most recent tag match information can be stored on the DRAM to allow subsequent accesses to the same way without another comparison, in one embodiment. If the cache DRAM is used according to a closed page policy, both tag and data row (way) will be closed after one access. The cache DRAM could also be used according to an open page policy where the data row (way) is kept open after an initial cache hit has occurred. In that case a consecutive column command can be used to access additional columns of the open row. An explicit precharge command will close the data row and the next access will initiate a tag comparison again.

[0074] The size of a tag and the bits allocated to address and metadata can be made configurable. An embodiment of a mode register, to select bit length for tags and comparators, corresponding to a configurable number of ways for the multiway set associative cache, is shown and described with reference to FIG. 7.

[0075] One embodiment is a DRAM device that supports cache operations, where cache line tags are held in dedicated row storage elements adjacent to associated data row storage elements on separate physical wordlines, that are accessed with two associated operations. The first operation accesses the data of a tag in a tag match and accesses the data rows storage elements. The second operation takes the data of the tag, uses this data to point to columns for the data, and accesses the cached data. These two operations are overlapped in timing shown and described with reference to FIGS. 4A and 5A. The data and cache accesses are thus overlapping, for one of the two associated operations. See, for example, the design in FIG. 1A, where tag and data rows are accessed concurrently, and the tag match activates column decode for the cached data access.

[0076] The data and cache accesses are non-overlapping, for the other one of the two associated operations. See, for example, the design in FIG. 2A, which causes a tag match to occur before another row is opened for the cached data access.

[0077] One embodiment of the multiway set associative cache has features as follows. 2GB cache size. This is the amount of DRAM available for caching data from storage class memory. Other sizes of cache are readily devised for further embodiments.

[0078] Two 8Gb LPDDR4 DRAM die. This is the number of dies available for the multiway

set associative cache, in one embodiment. Other numbers of dies are readily used for further embodiments.

[0079] 32 ways. This is the number of ways or addressed locations in each set in cache memory in which data from storage class memory can be cached in the multiway (now 32 way for this embodiment) set associative cache. Other numbers of ways are readily used for further embodiments.

[0080] 4kB per way. This is the amount of data that can be cached in each way, or line, in the 32 way or N way set associative cache. Other amounts of data for each way are readily used for further embodiments.

[0081] 64B transactions. This is the amount of data that can be read from cache memory, or written to each cache memory, in each (read or write) transaction. Other amounts of data for each transaction are readily used for further embodiments.

[0082] Operation compatible with low powered double data rate version 4 (LPDDR4) signaling and timing: one burst with 16b burst length on two 16b channels (32B per channel). In one embodiment, standard timing and signaling for LPDDR4 DRAM is used for both standard DRAM operation and N way set associative cache operation, when enabled. In another embodiment, the DRAM operates as N way set associative cache using standard timing and signaling for LPDDR4 DRAM, but is not selectable for standard DRAM operation without the N way set associative cache operation. Further embodiments could use standard timing and signaling for other types of DRAM as readily devised.

[0083] Addressing, including address decoding, and address line multiplexing is readily devised for the following configuration in one embodiment, and other variations in further embodiments.

[0084] A 2GB cache memory is 2^{31} bytes = (divided into) 2^5 ways, with 2^{12} bytes of cache memory per way and 2^{14} sets. Thus, a 14b set address identifies the set, and the size of the set is 128kB. Within a 4kB way, there are $64 = 2^6$ possible 64B transactions. A 6b offset address identifies the transaction in case of a cache hit.

[0085] With a 1TB SCM size, there is a 40b address for locations in SCM. The address to match in a tag is $40b - 14b$ (number of sets) $- 12b$ (size of way) = 14b. Adding two status bits makes a tag $2B = 16b$. 2B of tag per 4kB of data (in each way in this example 32 way set associative cache) is very little overhead for tag storage. Even when tags are stored at multiple location in the DRAM to facilitate parallel access, the overhead will be well below 1%.

[0086] Addressing, including address decoding, and address line multiplexing is readily devised for the following configuration in one embodiment, and variations in further embodiments. Addressing of 1b for device, 3b for bank, and 10b for the block of rows belonging

to one set identifies a region of 32 1kB rows per sub-bank. Four sub-banks are simultaneously used in this embodiment.

[0087] Depending on how and where data is cached, and design decisions for address multiplexing, bank select, row decode and row enable, and column decode and column enable, there are various possible timings for access to cached data. It is possible to combine multiple embodiments for timing in one device. In one embodiment, there is no overlap between tag compare and access to the cached data in the DRAM. The address of a set of data from main memory (e.g. SCM) which may be cached is separated into bit fields, and one field, a tag field of the address, is presented to the comparators for the tag match. Tags are read and compared to the tag field. A tag match generates part of the address for the cached data in the DRAM, and another part of the address for the cached DRAM comes from the address of the set of data from main memory. Only after all of the address bits are available is the access performed in the DRAM.

[0088] In other embodiments, there is full overlap or partial overlap of the access to the tags (i.e., the tag rows) for the tag match and the access for the cached data in the DRAM. Some of the address bits for the set of data in the main memory are used for bank select and row enable in the data rows of the DRAM, in parallel with or starting soon after the access to the tag rows. A tag match generates part of the address for the cached data in the DRAM, and this is used for the column decode to access the cached data.

[0089] Various embodiments with various numbers of tag address bits for compare, numbers of ways in a set, numbers of flag bits, arrangements of address bits for column and row decodes, designs and functions for sense amplifiers, etc., are presented herein. It should be appreciated that these are examples, and that further embodiments are readily devised in light thereof.

[0090] **FIG. 3A** illustrates circuitry and addressing for a cache hit and access to a data of a set that has been cached in a way of the multiway set associative cache, in one embodiment. The tag mats are identical to the normal data mats of the bank. For example in the embodiment shown, each mat 108 is a 512 bit by 512 bit array or sub-array (see FIG. 6), for a total of 256 kilobits (kb) per mat. The tag array allows a 16kb row to be accessed from the bank, and a 256b column from the row. The 256b column contains sixteen 16b tags 106, each of which could be a possible match for the physical address AP[39:27] field of the access. Tags are accessed using nine address bits AR[14:6] for row decode and six bits AR[5:0] for column decode 306. Tags are read in parallel and compared in parallel, with each tag access providing 16 bits, 13 of which are presented to a comparator 304 in a compare block 302 for comparison to the AP[39:27] address field, and the other three of which are flag bits (e.g., dirty, valid, parity). A further embodiment could use 14 bits for address comparison, leaving two flag bits. Other numbers of tag bits,

address comparison bits and flag bits are readily envisioned.

[0091] FIG. 3B illustrates the use of set bits from tag data in column decode to access cached data pointed to by the matched tag. A matching tag 106 is a HIT, and the set number (from 0 thru 15) Set[3:0] is passed by the compare blocks 302 to the column access path for the bank, column decode 306. In the circuitry shown, each compare block 302 is hardwired with a set number N, which is enabled onto the set number bits Set[3:0] by a hit in that compare block 302. The matching data is then accessed. Comparison and data access can happen in multiple banks 102 in parallel, in various embodiments (e.g., see FIGS. 1A, 2A, 9 and 10).

[0092] FIG. 3C illustrates address spaces in SCM and DRAM, for accessing data of SCM that is cached in the multiway set associative cache DRAM, according to one embodiment. In one embodiment, the SCM physical address is 40 bits $A_P[39:0]$, to address 2^{40} bytes (1 TB) of SCM. To access cached data, the 13 most significant bits $A_P[39:27]$ are used as a tag field, which generates Set[3:0], used for the column address bits $A_C[5:2]$ of the cached data. The block address, 15 bits of the SCM physical address, $A_P[26:12]$ is used for the row address $A_R[14:0]$ of the cached data. The five most significant bits of the sub block address of the SCM physical address, $A_P[11:6]$ provide two bits for the device address $A_D[1:0]$ and three bits for the bank address $A_B[2:0]$ of the cached data. The next two bits of the sub block address of the SCM physical address provide two more bits for the column address of the cached data, $A_C[1:0]$. Column decode 306 uses address bits $A_C[5:2]$ and $A_C[1:0]$. All sets are available in the DRAM cache after tag compare.

[0093] FIG. 4A is a read timing detail illustrating a normal access 402 of data and a modified access 404 of cached data, in one embodiment. A copy of a 4 kB block in DRAM (cache memory) is distributed across 32 banks 102 for better performance. In the embodiment shown, the 4 kB block is composed of mats, and each mat 108 is a 512 bit by 512 bit array or sub-array (see FIG. 6), for a total of 256 kilobits (kb) per mat. Further arrangements of banks, blocks and mats are readily devised. Column decode 306 is as described above with reference to FIG. 3A. Tags in tag mat rows 106 are as described above with reference to FIG. 1A. Variations for further embodiments are readily devised.

[0094] FIG. 4A shows a clock signal CK, grouped with command/address CA as clock, command and addresses (controller to DRAM), data bus DQ as data (DRAM used in normal mode), and bank control Bank as internal activities of the DRAM (DRAM used in normal mode), for a normal (direct) access 402 for reading data. Below these are tag activity Tag, grouped with bank control Bank as internal activities of the DRAM (DRAM used in cache mode), and data bus DQ as data (DRAM used in cache mode), for a modified (cached) access 404 for reading cached data. According to the bank control, for the normal (direct) access 402,

there is first an activate and sense of the bank, followed by a read and transfer to the interface Xface. Xface in the timing diagrams denotes the time it takes inside the DRAM for data from coming out of the array to showing up on the external DQ (read), e.g., FIG. 4A, or from coming in on the external DQ to being at the array (write), e.g., FIG. 5A.

[0095] The timing for a normal (direct) access 402 read delivers two columns of data in 41ns, with the activate and sense of the bank 102 taking 36 clock cycles or 18 ns, read taking eight clock cycles or 4 ns, and transfer to the interface (xface) taking eight clock cycles or 4 ns. FIGS. 4A and 5A show CA[5:0] as an example of the transfer of command and address from the controller to the DRAM following an LPDDR4 protocol. If DRAMs of another protocol are modified to be usable as cache DRAM, these functions could be done using explicit address pins and separate command pins like RAS, CAS and WE.

[0096] According to the tag activity, for the modified (cached) access 404, there is first an activate and sense of the tag, followed by a read of the tag data and compare, in which the determination of a tag hit or miss is made. CA input packets are the same for the normal (direct) access 402 as for the modified (cached) access 404. Overlapped with this activity, according to the bank control, there is the activate and sense of the bank, with a tag hit providing the set bits Set [3:0], followed by a data read of the cached data, and the transfer to the interface Xface. The timing for a modified (cached) read access 404 is 51ns, with activate and sense of the tag taking the same 36 clock cycles or 18 ns that a read of data in a normal access 402 takes, followed by tag compare taking 12 clock cycles or 6 ns. Timing is overlapped for the cached data access, with activate and sense of the bank 102 taking the same 36 clock cycles or 18 ns but overlapped with the activate and sense of the tag, followed by the same eight clock cycles or 4 ns for the read, and the same eight clock cycles or 4 ns for the transfer to the interface (Xface). The extra 10ns (for the overall time for the access to read cached data) includes 4ns for a column access and 6ns for a 16b compare.

[0097] **FIG. 4B** illustrates reading cached data from one of the banks of the multiway set associative cache DRAM, while more of the cached data is read in parallel from other banks, using the read timing shown in **FIG. 4A**. A copy of a 4 kB block of data from SCM is distributed across 32 banks in the DRAM (cache) memory, for better performance, in one embodiment (see, e.g., FIG. 1A). 256 bits of tag data are read out in parallel from the bank, using selected address bits for mat (and tag row) selection and column decode 306. A tag match, for one of the tags 106, produces indication of a hit (i.e., a cache hit, tag match or tag hit, see FIG. 3A) and set bits Set[3:0], along with dirty, valid and parity bits in one embodiment. The set bits are combined with specified address bits for column decode 306, and 256 bits of cache data are read out from the bank. Similar activity occurs across 32 banks of DRAM, resulting in the

reading, in parallel, of the 4 kB block of cached data. Variations, with different numbers of bits and different addressing and timing schemes are readily developed for further embodiments of the multiway set associative DRAM cache, including the embodiment shown in FIG. 2A.

[0098] FIG. 5A is a write timing detail illustrating a normal access 502 of data and a modified access 504 of cached data, in one embodiment. A copy of a 4 kB block in DRAM (cache memory) is distributed across 32 banks 102 for better performance. In the embodiment shown, the 4 kB block is composed of mats, and each mat 108 is a 512 bit by 512 bit array or sub-array (see FIG. 6), for a total of 256 kilobits (kb) per mat. Further arrangements of banks, blocks and mats are readily devised. Further arrangements of banks, blocks and mats are readily devised. Column decode 306 is as described above with reference to FIG. 3A. Tags in tag mat rows 106 are as described above with reference to FIG. 1A. Variations for further embodiments are readily devised.

[0099] FIG. 5A shows a clock signal CK, grouped with command/address CA as clock, command and addresses (controller to DRAM), data bus DQ as data (DRAM used in normal mode), and bank control Bank as internal activities of the DRAM (DRAM used in normal mode), for a normal (direct) access 502 for writing data. Below these are tag activity Tag, grouped with bank control Bank as internal activities of the DRAM (DRAM used in cache mode), and data bus DQ as data (DRAM used in cache mode), for a modified (cached) access 504 for writing cached data. According to the bank control, for the normal (direct) access 502 there is first an activate and sense of the bank, followed by a transfer to the interface Xface, a write and a column access. The timing for a normal (direct) write access 502 stores two columns of data in 34ns, with the activate and sense of the bank 102 taking 36 clock cycles or 18 ns, a waiting period, a transfer to the interface (Xface) taking eight clock cycles or 4 ns, write WR of eight clock cycles or 4 ns, and column COL of eight clock cycles or 4 ns.

[00100] According to the tag activity, for the modified (cached) access 504, there is first an activate and sense of the tag, followed by a read of the tag data and compare, in which the determination of a tag hit or miss is made. CA input packets are the same for the normal (direct) access 502 as for the modified (cached) access 504. Overlapped with this activity, according to the bank control, there is the activate and sense of the bank, with a tag hit providing the set bits Set [3:0], followed by a transfer to the interface (Xface) and two write cycles to write the cached data. Xface in the timing diagram denotes the time it takes inside the DRAM for data from coming out of the array to showing up on the external DQ (read), e.g., FIG. 4A, or from coming in on the external DQ to being at the array (write), e.g., FIG. 5A. The second WR in “Bank” for the modified (cached) access 504 is for the tag write WR. There is a difference between the line above labeled “Tag” and the line labeled “Bank”. The line “tag” shows what happens in the tag

logic, i.e. for WR the tag needs to be read, compared and a new tag with dirty bits etc. set needs to be prepared. The line “bank” shows what’s going on in the array. FIG. 4A and FIG. 5A both read the tag, but the line “Bank” of FIG. 4A and FIG. 5A does not show the tag read (it would be underneath the “Activate/Sense Bank” block as the tag is read while the data row is activated and sensed).

[00101] The timing for a modified (cached) write access 504 is 44ns, with activate and sense of the tag taking the same 36 clock cycles or 18 ns that a read of data in a normal access 502 takes, followed by tag read taking eight clock cycles or 4 ns and tag compare taking 12 clock cycles or 6 ns. This is followed immediately by the write wr, taking eight clock cycles or 4 ns. Tags are read (rd) and (if hit) written back (wr) with the dirty bit set and parity adjusted. Hit or miss status is returned to the controller. If there is a miss, the 64 byte write and tag write dirty operations are canceled. The extra 10ns (for the overall time for the access to write cached data) includes 4ns for a column access and 6ns for a 16b compare. The 16b tag is written back with its “dirty” flag set – this might increase the tRC of the tag access in some embodiments.

[00102] FIG. 5B illustrates writing cache data to one of the banks of the multiway set associative cache DRAM, while more of the cache data is written in parallel to other banks, using the write timing shown in FIG. 5A. A copy of a 4 kB block of data from SCM is distributed across 32 banks in the DRAM (cache) memory, for better performance, in one embodiment (see, e.g., FIG. 1A). 256 bits of tag data are read out in parallel from the bank, using selected address bits for mat (and tag row) selection and column decode 306. A tag match, for one of the tags 106, produces indication of a hit (i.e., a cache hit, tag match or tag hit, see FIG. 3A) and set bits Set[3:0], along with dirty, valid and parity bits in one embodiment. The set bits are combined with specified address bits for column decode 306, and 256 bits of cache data are written into the bank. If appropriate to a hit, the newly prepared tag with dirty bit and adjusted parity, etc. is also written back to the bank. Similar activity occurs across 32 banks of DRAM, resulting in the writing, in parallel, of the 4 kB block of cache data. Variations, with different numbers of bits and different addressing and timing schemes are readily developed for further embodiments of the multiway set associative DRAM cache, including the embodiment shown in FIG. 2A.

[00103] FIG. 6 illustrates DRAM cells 604 and sense amplifiers 608 in a mat 602, which can be used in embodiments of the DRAM. DRAM cells 604 are arranged in rows 612 and columns 610. Each DRAM cell 604 has a MOSFET and a capacitor, which can store a charge. The DRAM cell 604 in a given row 612 is activated by the word line 614 for that row 612, and reads out onto the bit line 616 for that column 610 of DRAM cells 604. All of the bit lines 616 are pre-charged by the bit line precharge 606, and selected DRAM cells 604 are read out through the

sense amplifiers 608. Sense amplifiers 608 may have further circuitry. Generally, a sense amplifier will have three functions, precharge, sense and connection to array data lines, and various sense amplifier designs are readily implemented for various embodiments. Of note is that in this design, each mat 602 has sense amplifiers 608 specific to the DRAM cells 604 of the mat 602. It is thus possible to arrange a tag mat row 106 with one set of mats 602 and data rows 104 associated to tags, i.e., as tag rows, with other sets of mats 602, so that each tag and tag mat row 106 is coupled to a sense amplifier (i.e., sense amplifiers 608 for the set of mats 602 making up the tag mat row 106) and the associated data rows are coupled to a different sense amplifier (i.e., sense amplifiers 608 for those sets of mats 602 making up the data mat rows 104), and access to tags can be overlapped with access to cached data pointed to by the tags.

[00104] In various embodiments, there are different ways of avoiding sense-amplifier conflicts. FIGS. 1A and 1B may have a tag mat and a data mat adjacent to each other, so they may need to have a double up sense-amplifier stripe only at that border, nowhere else in the bank. FIG. 2 and other possible embodiments do not need to double up the usual interleaved sense-amplifier since they are architected in a way that tags are never adjacent to the related data. Having adjacent mats that are operated at the same time in a 6F2 open bitline architecture requires the overhead of dummy arrays to provide complement bitlines and has therefore quite a bit of overhead.

[00105] FIG. 7 illustrates a mode register 702 for selection of standard DRAM mode, or multiway set associative cache mode with adjustable bit length of tags and comparators. One or more mode registers 702 control an address multiplexer (mux) 704, a data multiplexer (mux) 706, and the comparators 304. Address bits for row decode and row enable, and column decode and column enable, in tag access and cached data access, and selection of number and position of data bits for tag compares (see FIGS. 3A-5B) are thus controlled according to mode selection, through the address multiplexer 704 and/or the data multiplexer 706. In various embodiments, there could be a specified range of bit length for tags and comparators 304, corresponding to a range of numbers of ways for the multiway set associative cache. The mode register 702 could select numbers of bits for tag comparison or numbers of ways in a set. For example, a higher performance solution could be converted from 32 way to 16 way with twice the number of sets. The same DRAM could be used as a standard DRAM, compatible with a specified DRAM timing and signaling and the comparators disabled, or used as a multiway set associative cache with the comparators enabled, or even switched back and forth between the two modes. In a further embodiment, some or all of these features could be enabled or selected using fuse logic, instead of a mode register 702.

[00106] FIG. 8A is a flow diagram of a method of operating a multiway set associative

DRAM cache, which can be practiced by various embodiments described herein. For example, hardware, software executing on a processor, firmware, or combination thereof can perform these various actions, using timing and circuits as shown in FIGS. 3A, 4A or 5A, or variation thereof to perform an access of cached data in a DRAM as shown in FIGS. 1A-7 or variation thereof. In an action 802, tag rows of the DRAM are read. Tag rows are read in parallel, and could be located as depicted in and FIGS. 1A and 2A, or in variations thereof. An example of the read timing for the action 802 is shown in FIG. 4A, for the modified (cached) access 404, in which the tag activity shows activate/sense tag, followed by read (rd).

[00107] In an action 804, address bits of a set that may be cached (e.g., from main memory or SCM) in DRAM are presented to the comparators. Selection of address bits could be controlled as shown in FIG. 7 or FIG. 8B, or variation thereof. An example of the read timing for the action 804 is shown in FIG. 4A, for the modified (cached) access 404, in which the tag activity shows read (rd), followed by compare.

[00108] In an action 806, the address bits of the set and the data bits (i.e., tag information bits) from the tag rows of the DRAM are compared. Comparators as shown in FIGS. 3A-5B perform the comparisons in parallel across banks of the DRAM. An example of the read timing for the action 806 is shown in FIG. 4A, for the modified (cached) access 404, in which the tag activity shows compare, with a tag hit providing the set bits Set [3:0].

[00109] In a decision action 808, it is determined if there is a cache hit. Cache hit is determined from results of the comparisons in the action 806. An example of this timing is shown in FIG. 4A, for the modified (cached) access 404, in which the compare is performed to determine whether there is a cache hit or a cache miss. If the answer is no, there is no cache hit, flow proceeds to the action 810, to report cache miss. If the answer is yes, there is a cache hit, flow proceeds to the action 812.

[00110] In the action 812, address bits are generated for the multiway set associative cache hit. FIGS. 3A-5B show examples of generation of set bits Set [3:0] showing which way of the multiway set associative cache has the set of cached data.

[00111] In an action 814, data is read from data rows using address bits based on the cache hit. FIGS. 3A-5B show examples with column decode for reading cache data. An example of this timing is shown in FIGS. 4, for the modified (cached) access 404, in which the DQ bits are read onto the DQ bus after a cache hit.

[00112] **FIG. 8B** is a flow diagram of selecting an operating mode, which can be practiced by various embodiments described herein. In an action 820, the mode register is written to, to select standard DRAM mode or multiway set associative cache mode, with selected bit length of tags and comparators.

[00113] In an action 822, address multiplexing, data multiplexing, the bit length of tags and bit length of comparators are controlled, based on contents of the mode register. FIG. 7 shows an example of a mode register controlling these aspects.

[00114] FIG. 9 depicts an embodiment of the multiway set associative DRAM cache. Four banks 902 are shown, each with a row decoder 904. Each bank 902 has a column decode 906 and secondary sense amplifier (SA) 908. Each bank 902 also has cache logic 914. This example shows one multiplexer 920 per two banks 902, but could have one multiplexer 920 per bank or per any number of banks 902 in further embodiments. Internal command and address lines, from command/address logic/re-driver 916 connect to logic at the edge of the array (for bank control, decoder and row decoder), either from cache logic or from global command and address logic and re-driver 916. Internal data lines 910 connect secondary sense amplifiers to global data logic and re-driver 918, not to cache logic 914.

[00115] FIG. 10 depicts a further embodiment of the multiway set associative DRAM cache. Sense amplifiers 1008 are shared between neighboring mats 1004. Load mats 1002 are necessary at the edge of the bank. Activating a wordline (WL) blocks assessing wordlines sharing a sense amplifier 1008, e.g., WL b blocks both WL a and WL c.

[00116] Options for an embodiment of the multiway set associative DRAM cache (see left side of FIG. 10) include the following. Option A) Mat 1 and mat n store tags, mat 1 stores tags for data in mats 2 to $n/2$, and mat n stores tags for data in mats $n/2 + 1$ to $n-1$. Option B) Only mat 1 stores tags, mat 2 is a load mat shared between mat 1 and mat 3, and data are stored in mats 3 to n. Option C) Tags are distributed as extra row in all tags, tags for data in mats 1 to $n/2$ are in mats $n/2 + 1$ to n, and tags for data in mats $n/2 + 1$ to n are in mats 1 to $n/2$.

[00117] Options for an embodiment of the multiway set associative DRAM cache (see right side FIG. 10) include the following. Option D) Mat t 1006 stores tags (i.e., tag mat), mat t 1006 has different number of wordlines than mats 1 to n, mat t 1006 needs its own load mats with the same number of wordlines to have balanced load (cutting load bit lines and sharing load mat may create a process issue), and data are stored in mats 1 to n.

[00118] FIG. 11 depicts a DRAM 1102, and a memory controller 1104 that can operate the DRAM 1102 as a multiway set associative DRAM cache in accordance with present embodiments. In some versions, the DRAM 1102 is operated as a multiway set associative DRAM, and in other versions the DRAM 1102 can be operated selectably as either a standard DRAM or a multiway set associative DRAM (see, e.g., description of mode register 702 in FIG. 7). Signal lines connecting corresponding pins on both the DRAM 1102 and memory controller 1104 include CK, CKE, CS, CA, DQ, DQS, and DMI (see DRAM standards below). In this embodiment, the DMI signal and pin are repurposed for communicating cache information, and

labeled DMI/CACHE_INFO.

[00119] One example DRAM standard that is suitable for embodiments of the DRAM 1102 is the JEDEC standard number 209-4B, for LPDDR4 (low-power double data rate version 4)

DRAM, summarized below. Additional information is available under the published standard.

Clock, input, symbol CK_t_A, CK_c_A, CK_t_B, CK_c_B. Differential clock inputs, with each channel (A and B) having its own clock pair.

Clock Enable, input, symbol CKE_A, CKE_B. Each channel (A and B) has its own clock enable signal.

Chip Select, input, symbol CS_A, CS_B. Each channel (A and B) has its own chip select signal.

Command/Address Inputs, input, symbol CA[5:0]_A, CB[5:0]_B. Each channel (A and B) has its own command/address signals.

Command/address on-die-termination control, input, symbol ODT_CA_A, ODT_CA_B. Turns on or off the on-die-termination for the CA pins.

Data Input/Output, I/O, symbol DQ[15:0]_A, DQ[15:0]_B. Bidirectional data bus. Each channel (A and B) has its own bus.

Data Strobe, I/O, symbol DQS[1:0]_t_A, DQS[1:0]_c_A, DQS[1:0]_t_B, DQS[1:0]_c_B.

Bidirectional differential output clock signals used to strobe data during READ or WRITE. Each channel (A and B) has its own DQS strobes.

Data Mask Inversion, I/O, symbol DMI[1:0]_A, DMI[1:0]_B. Bidirectional signal indicates when data on the data bus is inverted or in normal state, or provides write data masking information to the DRAM, depending on mode register setting for data inversion or data mask.

Each channel (A and B) has its own DMI signals.

Calibration Reference, reference, symbol ZQ. Calibrates output drive strength and termination resistance.

Power Supplies, supply, symbol VDDQ, VDD1, VDD2.

Ground Reference, GND, symbol VSS, VSSQ.

Reset, input, symbol RESET_n. Active low reset signal.

[00120] Another example DRAM standard that is suitable for embodiments of the DRAM 1102 is the JEDEC standard number 79-4, for DDR4 (double data rate version 4) DRAM, summarized below. Similarly or identically named pins or signals in LPDDR4 and DDR4 standards perform similar or identical functions. Additional information is available under the published standard.

Clock, input, symbol CK_t, CK_c. Differential clock inputs.

Clock Enable, input, symbol CKE, (CKE1). Activates and deactivates internal clock signals, device input buffers and output drivers.

Chip Select, input, symbol CS_n, (CS1_n).

Chip ID, input, symbol C0, C1, C2. Selects each slice of a stacked component.

On die termination, input, symbol ODT, (ODT1). Applies to selected pins for selected configurations.

Activation Command Input, input, symbol ACT_n). Defines Activation command being entered along with CS_n.

Command Inputs, input, symbol RAS_n/A16, CAS_n/A15, WE_n/A14. Multifunction pins define command being entered.

Input Data Mask and Data Bus Inversion, I/O, symbol DM_n/DBI_n/TDQS_t, (DMU_n/DBIU_n), (DML_n/DBIL_n). Input mask signal, or indicator of true or inverted data.

Bank Group Inputs, input, symbol BG0 - BG1. Defines to which bank group Active, Read, Write or Precharge command is applied.

Bank Address Inputs, input, symbol BA0 – BA1. Defines to which bank Active, Read, Write or Precharge command is applied.

Address Inputs, input, symbol A0 - A17. Provide row address for Activate commands, and column address for Read/Write commands.

Auto-precharge, input, symbol A10/AP. Address A10 is sampled during Read Write commands to determine Auto-precharge.

Burst Chop, input, symbol A12/BC_n. Address A 12 is sampled during Read/Write commands to determine Burst Chop.

Active Low Asynchronous Reset, input, symbol RESET_n.

Data Input/Output, I/O, symbol DQ. Bidirectional data bus.

Data Strobe, I/O, symbol DQS_t, DQS_c, DQSU_t, DQSU_c, DQSL_t, DQSL_c. Output read data, input write data.

Termination Data Strobe, output, symbol TDQS_t, TDQS_c. Enables termination resistance function, data mask function or data mask inversion, depending on mode register.

Command and Address Parity Input, input, symbol PAR. Supports Even Parity.

Alert, I/O, symbol ALERT_n. Multifunction alert for CRC error, command and address parity error, connectivity test.

Connectivity Test Mode Enable, input, symbol TEN. Enables connectivity test mode operation.

No Connect, symbol NC.

DQ power supply, supply, symbol VDDQ.

DQ Ground, supply, symbol VSSQ.

Power Supply, supply, symbol VDD.

Ground, supply, symbol VSS.

DRAM activating power supply, supply, symbol VPP.

Reference voltage for CA, supply, symbol VREFCA.

Reference pin for ZQ calibration, supply, symbol ZQ.

[00121] FIG. 12 depicts a controller 1302 using a DRAM 1304 with cache mode in accordance with present embodiments. Presenting a solution to potential limitations of cache memory systems with tag SRAM on controllers, the components in FIG. 12 add tag capacity to the DRAM in a special sub-array (map) group, in various embodiments described with reference to Figs. 1A-11. In cache mode, the DRAM accesses the tag, and selects the correct data (or indicates miss). Tag and data spaces are interleaved for optimal performance.

[00122] In various embodiments, the controller 1302 could be a CPU, a separate system on chip (SoC) distinct from the CPU and memory, or on a module so long as the controller 1302 is in the path of the SCM. Buses that could connect the controller to the DRAM and the controller to the SCM include DDRX, SATA, PCIe and others. The system could work with one DRAM 1304 with one channel, one DRAM 1304 with two channels, two DRAMs 1304 each with one channel, etc., rather than two DRAMs each with two channels as shown in the example in Fig. 12. Also, there could be buffers in the path of the direct connection bus to expand capacity.

[00123] In one embodiment, described below as configuration example X-A, a copy of a 4 kB group 1202 in SCM (physical) memory stack 1312, is stored in in DRAM (cache) memory distributed across 32 banks in DRAM for best performance. DRAM 1304 is shown in two DRAM components (which could be ICs, multichip IC packages, logical units, etc.) each with two channels 1306, coupled to the controller 1302. Other configurations also described below can be implemented in the system shown in Fig. 12. The controller 1302 has a read queue 1310 and address field steering 1308, which steers address bits in various formats or schemes, according to configurations. Some embodiments have a single configuration, and some embodiments have multiple configurations. In some embodiments, the physical address fields are placed in CA packets for the DRAM. A control register field selects the cache configuration used. Address field steering 1308 uses logic, multiplexing or other circuitry, or software or firmware executing on a processor, or combinations thereof as readily devised in accordance with the teachings herein.

[00124] With the system in cache mode, cache read logic of the controller 1302 sends a cache read command (see Figs. 15A-17) across a memory bus to the DRAM 1304. Cache results logic receives a response from the DRAM 1304 SCM 1312 with the cached data, or an indication of a miss resulting from comparing the tag compare field with the tag field stored in the DRAM 1304. In the case of a miss, cache fetch logic sends a cache fetch command across another memory bus to SCM 1312 (shown as SCM stack). The cache fetch command causes the SCM

1312 to send the data associated with the tag field to the controller 1302. Cache fetch logic writes the data that it receives from the SCM 1312, in response to the cache fetch command, into the DRAM 1304 as cached data. In some embodiments, cache read logic retries the cache read command to the DRAM 1304.

[00125] Below is a table comparing system configurations of a DRAM with cache mode. For the various configurations X-A, Y-A, Z-A, Z-B, “A”= overlapped tag/data access, and “B”= non-overlapped tag/data access. By overlapped, it is meant that the data access overlaps the tag access, in other words the data access is begun before the tag access is complete. By non-overlapped, it is meant that the data access does not overlap the tag access, in other words the data access is begun after the tag access is complete. Non-overlapped tag/data access takes longer than overlapped tag/data access. “Z”= data “ways” held within a single bank, “Y”= data “ways” held within multiple banks of a single device, and “X”= data “ways” held within multiple banks of multiple devices. Configuration X-A, with data ways held within multiple banks of multiple devices and overlapped tag/data access is depicted in Fig. 13A. Configuration Y-A, with data ways held within multiple banks of multiple devices and overlapped tag/data access is depicted in Fig. 13B. Configuration Z-A, with data ways held within a single bank and overlapped tag/data access is depicted in Fig. 13C. Configuration Z-B, with data ways held within a single bank and non-overlapped tag/data access is depicted in 13D. Tags are held in associated tag banks with the same XYZ distribution option as the data in the databanks. That is, in configuration X, the data ways and the associated tag banks are held within multiple banks of multiple devices. In configuration Y, the data ways and the associated tag banks are held within multiple banks of a single device. In configuration Z, the data ways and the associated tag banks are held within a single bank.

configuration ²	X-A	Y-A	Z-A	Z-B
metrics				
tag/data row access	overlapped (A)	overlapped (A)	overlapped (A)	non-overlapped (B)
4 kB group distribution	32 banks (X)	8 banks (Y)	1 bank (Z)	1 bank (Z)
tag overhead (16 bits per tag)	1.6%	0.40%	0.05%	0.05%

hit access time ¹	tCA+trCD+tmATCH+ tCOL+tdQ~2ns+ 18ns+10ns+8ns+ 8ns~46ns	tCA+trCD+tmATCH+ tCOL+tdQ~2ns+ 18ns+10ns+8ns+ 8ns~46ns	tCA+trCD+tmATCH+ tCOL+tdQ~2ns+ 18ns+10ns+8ns+ 8ns~46ns	tCA+trCD+tmATCH+ tCOL+tdQ~2ns+ 18ns+10ns+18ns+ 8ns+8ns~64ns
energy/bit for 64 byte hit access	~1x	~1x	~1x	~1x
bandwidth to access 4 kB group (after miss)	~3.2x (tFAW limited)	~0.8x (tFAW limited)	~0.4x (trC limited)	~1x
energy/bit to access 4 kB group (after miss)	~0.85x (1/2 E _{ROW})	~0.85x (1/2 E _{ROW})	~0.85x (1/2 E _{ROW})	~0.7x (1/32 E _{ROW})
¹ legacy (non-cache) access time tCA+trCD+tCOL+tdQ~ 2ns+18ns+ 8ns+8ns~36ns	² configuration assumptions 8 GB/DRAM 32 bytes/column 64 column/row (2 kB row) 16 K rows/bank 8 banks/channel		2 channels per DRAM trC ~50ns tFAW~40ns 4 Gb/second DQ bit rate 16 ways (sets)/DRAM cache line ³ 4 kB/DRAM cache line (group)	
³ note-16 ways is chosen so that the 16x16b tags for a group address in the cache fits in one column access; if more ways are needed, then more than one column access is needed to access and check all the tags				

[00126] The above comparison and following details are for one particular embodiment, but in any given implementation specifications like access time may differ. A legacy (non-cache) access time is about 36 ns. The configuration assumptions include two channels per DRAM, 8 GB per DRAM, 32 bytes per column, 64 columns per row (2 kB rows), 16 K rows per bank, eight banks per channel, trC about 50 ns, tFAW about 40 ns, 4 Gb per second DQ bit rate, 16 ways (sets) per DRAM cache line, and 4 kB per DRAM cache line (group). All of the overlapped tag/data access configurations X-A, Y-A, Z-A have hit access time about 46 ns, which is shorter than the hit access time 64 ns for the non-overlapped tag/data access of configuration Z-B. But, the non-overlapped tag/data access of configuration Z-B has the lowest energy per bit cost to

access a 4 kB group after a miss, and has a tie with configuration Z-A for the lowest tag overhead at .05%. Other characteristics and trade-offs for the various configurations are readily explored in the table.

[00127] FIG. 13A depicts system configuration example X-A, in which the data ways are held within multiple banks 102 of multiple devices, and there is overlapped tag/data access. A 4 kB group 1202 in SCM (physical) memory is copied and distributed across 32 banks 102 in this example. The 16 bit tag for the distributed 4 kB group 1202 is duplicated in the 32 tag memories 1502, one per bank 102. 16 bit tags for other ways of 4 kB groups (at the same cache group address) are held in the same column of one row in the bank 102. Transport blocks for other ways of 4 kB groups (at the same cache group address) are held in the same row as the 64 byte transport block, which is held in two 32 byte columns of the row in the bank 102. Configuration X-A may benefit DRAM cache systems using wide data buses across multiple channels.

[00128] FIG. 13B depicts system configuration example Y-A, in which the data ways are held within multiple banks 102 of a single device, and there is overlapped tag/data access. A 4 kB group 1202 in SCM (physical) memory is copied and distributed across four rows in each of eight banks 102 of one channel (channel D) in this example. The 16 bit tag for the distributed 4 kB group 1202 is duplicated in the eight tag memories 1502 (one in each of the eight banks of one channel, channel D) in this example. 16 bit tags for other ways of 4 kB groups (at the same cache group address) are held in the same column of one row in the bank 102. Transport blocks for other ways of 4 kB groups 1202 (at the same cache group address) are held in the same four rows in the bank 102 as the 64 byte transport block, which is held in two 32 byte columns of each of the four rows in the bank 102. Configuration Y-A may be beneficial for multiple sequential data accesses at successive addresses in a single channel. There may be less average power consumption for multiple sequential data accesses at successive addresses in a single channel than with other configurations.

[00129] FIG. 13C depicts system configuration example Z-A, in which the data ways are held within a single bank 102, and there is overlapped tag/data access. A 4 kB group 1202 in SCM (physical) memory is copied and distributed across 32 rows of one bank 102 of one channel in this example. A 16-bit tag 106 for a 64 byte transport block is in associated tag memory 1502 in the bank 102. There are no duplicate tags in other banks 102. 16 bit tags for other ways of 4 kB groups 1202 (at the same cache group address) are held in the same column of one row in the tag memory 1502 in the bank 102. Transport blocks for other ways of 4 kB groups 1202 (at the same cache group address) are held in the same 32 rows as the 64 byte transport block, which is held in two 32 byte columns of each of 32 rows in the bank 102. Configuration Z-A does not require duplicate tags and is thus less consumptive of rows of memory for the tags.

Configurations X, Y, and Z correspond to different schemes for addressing rows, columns, banks and channels in the DRAM, and may affect access times and energy consumption.

[00130] **FIG. 13D** depicts system configuration example Z-B, in which the data ways are held within a single bank 102, and there is non-overlapped tag/data access. A 4 kB group 1202 in SCM (physical) memory is copied and distributed across 32 rows of one bank 102 of one channel in this example. A 16-bit tag 106 for a 64 byte transport block is in associated tag memory 1502 in the bank 102. There are no duplicate tags in other banks 102. 16 bit tags for other ways of 4 kB groups 1202 (at the same cache group address) are held in the same column of one row in the tag memory 1502 in the bank 102. Transport blocks for other ways of 4 kB groups (at the same cache group address) are held in the same column of one row in the tag memory 1502 in the bank 102. Transport blocks for other ways of 4 kB groups 1202 (at the same cache group address) are held in different rows in the bank 102. The 64 byte transport block is held in two 32 byte columns in two rows in the bank 102. Configuration Z-B does not require duplicate tags and is thus less consumptive of rows of memory for the tags.

Configuration B has longer access time than configuration A, but may have lower energy consumption.

[00131] **FIG. 14A** depicts address field steering 1308 in the controller 1302 of **FIG. 12** for system configuration example X-A as shown in **FIG. 13A**. The DRAM byte address 1604, for 2 GB of DRAM, is derived from the controller (physical) address 1602 for 1 TB memory (e.g., SCM, for which various groups 1202 are cached in the DRAM), through address field steering 1308 on the controller 1302 and tag compare on the DRAM. From the tag field, $A_P[39:27]$, of the controller (physical) address 1602, the DRAM produces the set bits $Set[3:0]$ through tag compare, which are used for four column address bits $A_C[5:2]$ of the DRAM byte address 1604. All 15 bits of the group address, $A_P[26:12]$, of the controller (physical) address 1602, are used for the row address bits $A_R[14:0]$ of the DRAM byte address 1604. The controller 1302 derives the channel address bits $A_D[1:0]$, and the bank address bits $A_B[2:0]$, of the DRAM byte address 1604 from the upper two bits and next three bits of the sub-group address of the controller (physical) address 1602. The next lower two bits of the sub-group address of the controller (physical) address 1602 provide the remaining two column address bits $A_C[1:0]$ of the DRAM byte address 1608. And the lowest five bits of the sub- group address of the controller (physical) address 1602 provide the five bits of the sub-column address, for a 64 byte read/write, to complete the DRAM byte address 1604.

[00132] **FIG. 14B** depicts address field steering 1308 in the controller 1302 of **FIG. 12** for system configuration example Y-A as shown in **FIG. 13B**. The DRAM byte address 1606, for 2 GB of DRAM, is derived from the controller (physical) address 1602 for 1 TB memory (e.g.,

SCM, for which various groups 1202 are cached in the DRAM), through address field steering 1308 on the controller 1302 and tag compare on the DRAM. From the tag field, $A_P[39:27]$, of the controller (physical) address 1602, the DRAM produces the set bits $Set[3:0]$ through tag compare, which are used for four column address bits $A_C[5:2]$ of the DRAM byte address 1606. From the upper two bits of the group address, $A_P[26:25]$, of the controller (physical) address 1602, the controller 1302 provides the channel address bits $A_D[1:0]$. From the lower or remaining 13 bits of the group address, $A_P[24:12]$, of the controller (physical) address 1602, the controller 1302 provides the upper 13 bits of the row address, $A_R[14:2]$ for the DRAM byte address 1606. The upper three bits of the sub-group address, $A_P[11:9]$, of the controller (physical) address 1602 are steered as the bank address bits $A_B[2:0]$ of the DRAM byte address 1606. The next lower two bits of the sub-group address, $A_P[8:7]$, of the controller (physical) address 1602, provide the lowest two bits of the row address, $A_R[1:0]$, of the DRAM byte address 1606. The next lower two bits of the sub-group address, $A_P[6:5]$, of the controller (physical) address 1602 are used as the lowest two bits of the column address, $A_C[1:0]$, in the DRAM byte address 1606. And the lowest five bits of the sub-group address of the controller (physical) address 1602 provide the five bits of the sub-column address, for a 64 byte read/write, to complete the DRAM byte address 1606.

[00133] FIG. 14C depicts address field steering 1308 in the controller 1302 of FIG. 12 for system configuration example Z-A as shown in FIG. 13C. Solid line arrows in FIG. 14C depict the tag compare process on the DRAM. The tag field, which is part of the physical address, is compared to the stored tags in the DRAM, and if there is a match, one of the 16 sets is selected (the $Set[3:0]$ field) - the Set field forms part of the DRAM address that is used to select the data. The DRAM byte address 1608, for 2 GB of DRAM, is derived from the controller (physical) address 1602 for 1 TB of memory (e.g., SCM, for which various groups 1202 are cached in the DRAM), through address field steering 1308 on the controller 1302 and tag compare on the DRAM. From the tag field, $A_P[39:27]$, of the controller (physical) address 1602, the DRAM produces the set bits $Set[3:0]$ through tag compare, which are used for four column address bits $A_C[5:2]$ of the DRAM byte address 1608. That is, the set bits, from tag compare, are used for four column address bits to fetch the cached data. From the group address, $A_P[26:12]$, of the controller (physical) address 1602, the controller 1302 derives the channel address bits $A_D[1:0]$, the bank address bits $A_B[2:0]$, and the upper 10 row address bits $A_R[14:5]$ of the DRAM byte address 1608. Five more row address bits $A_R[4:0]$ of DRAM byte address 1608 are derived from the upper five bits of the sub-group address $A_P[11:0]$ of the controller (physical) address 1602. The middle two bits of the sub-group address of the controller (physical) address 1602 provide two more column address bits $A_C[1:0]$ of the DRAM byte address 1608. The lower five bits of

the sub-group address of the controller (physical) address 1602 provide the five bits sub-column address of the DRAM byte address 1608.

[00134] FIG. 14D depicts address field steering 1308 in the controller 1302 of FIG. 12 for system configuration example Z-B as shown in FIG. 13D. The DRAM byte address 1610, for 2 GB of DRAM, is derived from the controller (physical) address 1602 for 1 TB of memory (e.g., SCM, for which various groups 1202 are cached in the DRAM), through address field steering 1308 on the controller 1302 and tag compare on the DRAM. From the tag field, $A_P[39:27]$, of the controller (physical) address 1602, the DRAM produces the set bits $Set[3:0]$ through tag compare, which are used for the upper four row address bits $A_R[14:11]$ of the DRAM byte address 1608. From the group address, $A_P[26:12]$, of the controller (physical) address 1602, the controller 1302 derives the channel address bits $A_D[1:0]$, the bank address bits $A_B[2:0]$, and the next lower 10 row address bits, $A_R[10:1]$, of the DRAM byte address 1608. The lowest row address bit, $A_R[0]$, of DRAM byte address 1608 is from the uppermost bit of the sub-group address, $A_P[11]$, of the controller (physical) address 1602. The next lower six bits of the sub-group address, $A_P[10:5]$, of the controller (physical) address 1602 provide the six column address bits $A_C[5:0]$ of the DRAM byte address 1608. The lower five bits of the sub-group address of the controller (physical) address 1602 provide the five bits sub-column address of the DRAM byte address 1608.

[00135] FIG. 15A depicts command format for system configuration example Z-B as shown in FIG. 13D. The redundancy in FIG. 15A and FIG. 15B occurs because there are two column access per row access - other DRAM components or subsystems might set this ratio to one column access per row access, and there would not be a redundant field in the command sequence. An ACT command is followed by RD command 0 and RD command 1. In this example, the ACT command indicates it has the least significant tag field 1702 bits, $A_P[30:27]$. The first read command, RD command 0, has the most significant tag field 1704 bits, $A_P[39:31]$. The tag field is not needed in subsequent RD commands in the same row for the same group, e.g., RD command 1. This is because the read command, which has two parts per protocol, has already specified in RD command 0 the tag bits needed to complete the tag field when combined with the tag field 1702 from the ACT command, so the tag field in RD command 1 is redundant.

[00136] FIG. 15B depicts command format for system configuration examples Z-A, Y-A and X-A as shown in FIGS. 13C, 13B and 13A. An ACT command is followed by a RD command 0 and RD command 1. In this example, the ACT command has only the row address. The first read command, RD command 0, has the most significant tag field 1706 bits $A_P[39:31]$ and least significant tag field 1708 bits $A_P[30:27]$. The tag field is not needed in subsequent RD commands in the same row for the same group, e.g., RD command 1. This is because the read

command, which has two parts per protocol, has already specified in RD command 0 all of the tag bits needed to complete the tag field, so the tag field in RD command 1 is redundant.

[00137] FIG. 15C depicts command format for system configuration examples Z-B, Z-A, Y-A and X-A as shown in FIGS. 13D, 13C, 13B and 13A. An ACT command is followed by a RD command 0 and RD command 1. In this example, the first read command, RD command 0, indicates it has the most significant tag field 1710 bits $A_P[39:33]$. The next read command, RD command 1, has the least significant tag field 1712 bits $A_P[32:27]$. If only a 32 byte access is needed, then the same column can be read twice. Unlike the command format examples in FIGS. 15A and 15B, the RD command 1 does need to provide the remaining tag bits to complete the tag field and so does not have a redundant tag field. With reference to FIGS. 15A-15C, the concept of a "shared" tag field can refer to the fact that the tag field (from the physical address) is split across a row command and a column command (in FIG. 15A), or is specified by one column command and used later by a second column command (in FIG. 15B), or is split across a two successive column commands (in FIG. 15C) - any of these alternatives allow the number of CA pins to be minimized.

[00138] FIG. 16 depicts a read timing detail for a cache miss, with fetch and evict, in a DRAM with cache mode. The controller queue receives the next read (RD) transaction (T) 1720, at (1). This is shown as a read address on the address bus (ADR). The ACT/RD access 1722 to DRAM is performed, at (2). The access is canceled by the tag MISS 1724, at (3). This is signaled with the status (ST) signals from the DRAM to the controller. Transaction (T) status is retry, which the system accomplishes by sending status to the execution unit via ST2 signals, at (4). The transaction (T) is moved to the retry buffer at (5) until the fetch operation has returned the read (RD) data. The system initiates a concurrent fetch process at (6), to get the group with the missing read (RD) data from external memory (XMEM). The concurrent fetch process is shown as fetch new group 1726. At (7), the memory controller selects a group to remove from cache memory (to make room for the new group), and initiates a concurrent evict process. This is shown as evict old group 1728. At (8), the system continues with concurrent processing of transactions. The fetch access to the 4 kB fetch buffer is finished at (9), and the 4 kB buffer has been moved to DRAM. The fetch completion is signaled to the memory controller in the background field of the status packet at (10), and the original read can be retried.

[00139] FIG. 17 is a read timing detail for a 64 byte read of a DRAM with cache mode, for a normal (direct) access 1802, and a modified (cached) access 1804. At a high level (and referring to the table at [00125], the differences between the four configurations affect hit access latency, tag overhead cost, and the bandwidth and energy/bit to access a 4KB group after a miss (evict/fetch) - these interact in a complicated way, but they cannot be simultaneously optimized -

that is why four configurations are needed to illustrate the optimal point for each metric. The CA input packets are in the same position for the modified (cached) access 1804 as for the normal (direct) access 1802. The ACT command, followed by RD command 0 and RD command 1, follow the sequences and have the bits described above with reference to FIGS. 15A-15C, as appropriate to system configuration. Configurations X-A, Y-A and Z-A (see table Comparison of system configuration examples X-A, Y-A, Z-A, Z-B, and FIGS. 13A-13C, 14A-14C and 15B) have overlapped tag and data access, with timing in FIG. 17 shown for configurations XYZ-A as modified (cached access) 1804. Configuration Z-B (see table Comparison of system configuration examples X-A, Y-A, Z-A, Z-B, and FIGS. 13D, 14D and 15A) have non-overlapped tag and data access, with timing in FIG. 17 shown for configuration Z-B as modified (cached) access 1804. This timing can be compared to the faster data read in the normal (direct) access 1802, which is for reading the DRAM directly without cache access.

[00140] With reference to FIGS. 16A-16C and 17, in one embodiment there are three commands, the ACT followed by two successive column commands, for two 32 byte read or write data accesses that combine to a 64 byte access. The two successive column commands are a tag column access operation and a subsequent data column access operation. The same tag is used for both column accesses. That is, the tag compare field of the tag column access operation is also used for the tag compare field of the data column access operation.

[00141] FIG. 18 is a flow diagram depicting controller sequencing for a DRAM cache, experiencing a cache miss. With forward reference to FIG. 19, various system components participate in the actions described in FIG. 18. Starting with an action 1902, [1] get next transaction “T”, the transaction could be a read, e.g., 64 byte read (see left side of FIG. 19), or a write, e.g., 64 byte write (see right side of FIG. 18).

[00142] For the 64 byte read, flow proceeds to the action 1904, [2r] transaction T to read the DRAM 1304, and determine whether there is a hit or a miss. If a hit, flow proceeds to the action 1906, return read data. If a miss [3], flow proceeds to the actions 1908.

[00143] For the 64 byte write, flow proceeds to the action 1930, [2w] transaction T to write to the DRAM, and determine whether there is a hit or a miss. If a hit, flow proceeds to the action 1932, to write 64 byte data to the DRAM 1304 (according to the cache hit), then to the action 1934, set dirty flag for 128 byte sub-group. If a miss [3], flow proceeds to the actions 1908.

[00144] For the actions 1908 [8], the controller 1302 performs actions [4] mark transaction T status as “retry”, [5] save transaction T, [6] initiate concurrent “fetch” process, [7] initiate concurrent “evict” process, and [8] continue concurrent processing of transactions.

[00145] For the [6] fetch, flow proceeds to the action 1910, [6a] start access of 4 kB group B_T (containing transaction T target) in SCM and transfer to 4 kB fetch buffer 2006 in the controller

1302. Process flow branches in parallel, with fetch is done, to the action 1922 [7c], and continues with action 1912 [6b] transfer 4 kB group B_T from 4 kB buffer in controller 1302 to DRAM 1304. This is followed by action 1914 [6c] return transaction T back to queue to be retried, in the read retry buffers 2022, and action 1916 [6d] terminate fetch process. In some embodiments, once the fetch from SCM into DRAM is completed, the memory controller 1302 has to get the data from the DRAM 1304. The controller could wait for confirmation, or in some versions rely on a predetermined time interval, and then request the data from the DRAM 1304 again with the process of action 1902. In some embodiments, the controller 1302 reads data as the data is moved from SCM to the DRAM cache, thereby preventing the need for a separate DRAM read once the data is in the DRAM cache.

[00146] For the [7] evict, flow proceeds to the action 1918, [7a] select 4 kB group B_E in DRAM cache to evict (group B_E and group B_T have the same A_P[26:12] address in the 16 sets for that address). Next, flow proceeds to the action 1920 [7b] transfer 4 kB group B_E (containing transaction) in from DRAM to 4 kB evict buffer 2004 in controller 1302 along with 32 dirty flags (one dirty flag per 128 bytes of data). Flow proceeds to the action 1922 [7c] wait for B_T to be accessed and transferred from SCM to control by fetch process (optional). Flow proceeds to the action 1924 [7d] transfer 4 kB group B_E (containing transaction) in 4 kB evict buffer 2004 in controller 1302 to SCM. Optionally, the 32 dirty flags can control the writing of each 128 bytes of data. In action 1926, [7e] the evict process is terminated.

[00147] Evict selection options, for various embodiments, are described below.

[7a.1] eviction selection could be random.

[7a.2] several of the 32 128 byte sub-groups of the 16 target 4 kB groups could be sampled. The number of high dirty flags could be used to indicate how much each 4 kB group has been used. The group with the most usage could be chosen. If SCM endurance is an issue, instead the group with zero high dirty flags could be chosen.

[7a.3] Similar to the above, but the touch flag could be used instead to indicate how much the 4 kB block has been used.

[00148] **FIG. 19** is a system action diagram depicting the controller using DRAM with cache mode, performing the actions of **FIG. 18**. In the DRAM system, DRAM 1304 is shown as DRAM 0 and DRAM 1 in first DRAM rank, with more DRAM in a second DRAM rank. Each DRAM 1304 has two channels, channel A and channel B in DRAM 0, channel C and channel D in DRAM 1. SCM memory stack 1312 has a 4 kB group 1202 in SCM (physical) memory. Inside the controller 1302, transaction formatting and steering 2010 is coupled to the DRAM 1304 through interface 2012 (denoted xface). Status/control 2014 is coupled to transaction formatting and steering 2010 for the interfaces 2012 coupled to the DRAM 1304. Transaction

formatting and steering 2010 is coupled to a write queue 2016, read queue 2018, write retry buffers 2020 and read retry buffers 2022. Further transaction formatting and steering 2002 is coupled to the SCM memory stack 1312 through a further interface 2024. A further status 2008 is coupled to the transaction formatting and steering 2002. The transaction formatting and steering 2002 is coupled to evict buffers 2004 and fetch buffers 2006, which further couple to the transaction formatting and steering 2014, for the DRAM 1304.

[00149] Controller sequencing for the DRAM cache, for a cache miss as described above with reference to FIG. 18, is readily followed in FIG. 19 with use of the bracketed numbers correspondingly present in both FIG. 18 and FIG. 19.

[00150] FIG. 20 depicts a touch flag enhancement for the eviction decision process. A touch flag is added to the 16 bit tag flag word 2106, which then has 12 bits for the tag compare and four flag bits, as dirty, valid, parity and touch 2104. Touch 2104 (and dirty) bits are initially low when the 4 kB block is loaded from SCM by the fetch process. A 64 byte read transaction sets the touch 2104 flag high (and adjusts parity) in the local tag flag word 2106 for the 128 byte sub-block that is accessed (similar to when a 64 byte write transaction sets the dirty flag). The touch bit, cleared when the 4 kB block is first written to DRAM cache (i.e., loaded from SCM), is set when the 4 kB block or part of it is read from DRAM cache, to inform the system that the cached 4 kB block is being read and is therefore worth keeping. A 4 kB block that has been written to DRAM cache, so that the touch bit is cleared, but still has the touch bit showing cleared at a later date, may be a candidate for eviction since it is not being read and is therefore not worth keeping. The dirty bit, set by a write transaction, informs the system that cached data has been written to and should therefore be written back to SCM. This is useful in a writeback cache. For the eviction decision process, both the dirty and touch bits have purpose. The touch bit can be used to identify candidates for eviction, as above, and the dirty bit can be used to make sure that the cached data that has been written to (i.e., overwritten) is written back to the SCM, before that location in cache memory can be considered a candidate for eviction.

[00151] FIG. 21A depicts a 4 kB cache line distributed across banks/channels/devices in accordance with present embodiments. In this example, the data and associated tag for each 4 kB cache line are distributed across all the available banks 102 (e.g., 32 banks from two DRAMs and two channels per DRAM, with eight banks 102 per channel). This allows a 4 kB cache line 2202 to be transferred between the DRAMs and the controller as quickly as possible. This makes the fetch and evict processes easier to schedule. FIG. 22A shows how a 4 kB cache line 2202 in physical SCM address space is distributed across banks 102 of multiple channels and multiple devices to minimize transport time during the fetch and evict processes.

[00152] FIG. 21B depicts 4 kB cache line tags and data for 16 sets (ways) located in the same

DRAM row in accordance with present embodiments. In this example, the transfer blocks of data (e.g., two 64 byte transfer blocks) for a way are located in the same row. A way is the multiple sets of data with the same lower and middle address, but different upper address. The tags for a way are located in a single column (e.g. 16 x 16 bits). This generates a tag after a row/column access and tag compare, and the data after two more column accesses. The row access of the tag is overlapped with the row access of the data. FIG. 21B shows how the 16 aliased 4 kB cache lines in physical SCM address space are accessed to find the one with the matching 39:27 address. In the example, these 16 4 kB cache lines 2202 occupy the 16 sets at the address $A_P[26:12]$. They have the same $A_P[26:12]$ row address but different $A_P[39:27]$ match addresses. Note that each row of the tag match contains 64 columns of 16 x 16 bit tags, associated with 64 rows in the data bank 102.

[00153] The methods, systems and devices described above may be implemented in computer systems, or stored by computer systems. The methods described above may also be stored on a non-transitory computer readable medium. Devices, circuits, and systems described herein may be implemented using computer-aided design tools available in the art, and embodied by computer-readable files containing software descriptions of such circuits. This includes, but is not limited to one or more elements of DRAMs and their components. These software descriptions may be: behavioral, register transfer, logic component, transistor, and layout geometry-level descriptions. Moreover, the software descriptions may be stored on storage media or communicated by carrier waves.

[00154] Data formats in which such descriptions may be implemented include, but are not limited to: formats supporting behavioral languages like C, formats supporting register transfer level (RTL) languages like Verilog and VHDL, formats supporting geometry description languages (such as GDSII, GDSIII, GDSIV, CIF, and MEBES), and other suitable formats and languages. Moreover, data transfers of such files on machine-readable media may be done electronically over the diverse media on the Internet or, for example, via email. Note that physical files may be implemented on machine-readable media such as: 4 mm magnetic tape, 8 mm magnetic tape, 3-1/2 inch floppy media, CDs, DVDs, hard drives, solid-state drives (SSD), and so on.

[00155] The above description of illustrated embodiments of the disclosure, including what is described in the Abstract, is not intended to be exhaustive or to limit the disclosure to the precise forms disclosed. While specific embodiments of, and examples for, the disclosure are described herein for illustrative purposes, various equivalent modifications are possible within the scope of the disclosure, as those skilled in the relevant art will recognize. Other embodiments may have layers in different orders, additional layers or fewer layers than the illustrated embodiments.

[00156] Various operations are described as multiple discrete operations, in turn, in a manner that is most helpful in understanding the present disclosure, however, the order of description should not be construed to imply that these operations are necessarily order dependent. In particular, these operations need not be performed in the order of presentation.

[00157] The terms “over,” “above” “under,” “between,” and “on” as used herein refer to a relative position of one material layer or component with respect to other layers or components. For example, one layer deposited above or over or under another layer may be directly in contact with the other layer or may have one or more intervening layers. Moreover, one layer deposited between two layers may be directly in contact with the two layers or may have one or more intervening layers. In contrast, a first layer “on” a second layer is in direct contact with that second layer. Similarly, unless explicitly stated otherwise, one feature deposited between two features may be in direct contact with the adjacent features or may have one or more intervening layers.

[00158] The words “example” or “exemplary” are used herein to mean serving as an example, instance, or illustration. Any aspect or design described herein as “example” or “exemplary” is not necessarily to be construed as preferred or advantageous over other aspects or designs. Rather, use of the words “example” or “exemplary” is intended to present concepts in a concrete fashion. As used in this application, the term “or” is intended to mean an inclusive “or” rather than an exclusive “or.” That is, unless specified otherwise, or clear from context, “X includes A or B” is intended to mean any of the natural inclusive permutations. That is, if X includes A; X includes B; or X includes both A and B, then “X includes A or B” is satisfied under any of the foregoing instances. In addition, the articles “a” and “an” as used in this application and the appended claims may generally be construed to mean “one or more” unless specified otherwise or clear from context to be directed to a singular form. Moreover, use of the term “an embodiment” or “one embodiment” or “an embodiment” or “one embodiment” throughout is not intended to mean the same embodiment or embodiment unless described as such. The terms “first,” “second,” “third,” “fourth,” etc. as used herein are meant as labels to distinguish among different elements and may not necessarily have an ordinal meaning according to their numerical designation.

CLAIMS

What is claimed is:

1. A memory controller comprising:

cache read logic to send a cache read command across a first memory bus to a first memory component, the cache read command including: a command opcode, a row address for a row which stores a tag field associated with data, and a tag compare field;

cache results logic to receive a response from the first memory component across the first memory bus with the data or an indication of a miss resulting from comparing the tag compare field with the tag field stored in the row of the first memory component; and

cache fetch logic to send a cache fetch command across a second memory bus to a second memory component, in response to the indication of the miss, the cache fetch command to cause the second memory component to send the data associated with the tag field stored in the row of the first memory component to the memory controller.

2. The memory controller of claim 1, wherein the cache fetch logic is to write the data received from the second memory component in response to the cache fetch command, into the first memory component.

3. The memory controller of claim 1, wherein the cache read logic is to retry the cache read command to the first memory component.

4. The memory controller of claim 1 wherein:

the first memory component comprises a dynamic random access memory (DRAM) that is configurable as a cache memory and has the row which stores the tag field; and

the second memory component comprises storage class memory (SCM).

5. The memory controller of claim 1, wherein the cache read command comprises:

a first command comprising a first tag field and a row address field;

a second command comprising a second tag field and a column address field; and

a third command comprising a redundant tag field.

6. The memory controller of claim 1, wherein the cache read command comprises:

a first command comprising a row address field;

a second command comprising two tag fields; and

a third command comprising a redundant tag field.

7. The memory controller of claim 1, wherein the cache read command comprises:

a first command comprising a row address field;

a second command comprising a first tag field and a column address field; and

a third command comprising a second tag field.

8. A memory controller, comprising:

logic having a first operating mode of the memory controller, to read from a memory component having at least one data row and at least one tag row different from and associated with the at least one data row to implement a cache having multiple ways to hold a data group using a first address and timing for accessing the data group; and

the logic having a second operating mode of the memory controller, to read from the memory component using a second addressing and timing for accessing data,

wherein the first address and timing differs from the second address and timing.

9. The memory controller of claim 8, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the logic operable to transmit in the first operating mode, to the DRAM, a command/address (CA) packet comprising a chip select (CS), a bank address field, a row address field, a tag compare field, and a command opcode.

10. The memory controller of claim 8, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the logic operable to transmit in the first operating mode, to the DRAM, a command/address (CA) packet comprising a chip select (CS) for device or rank selection on a channel, a bank address field, a row address field, a tag compare field, and a command opcode, wherein tag compare fields of two successive column commands are shared.

11. The memory controller of claim 8, further comprising:

the logic to repurpose, in the second operating mode, one or more unused tag rows in the memory component for other use.

12. The memory controller of claim 8, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the logic operable to transmit in the first operating mode, to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field.

13. A memory system, comprising:

a memory component having at least one tag row and at least one data row and multiple ways to hold a data group as a cache-line or cache-block;

a memory controller that is connectable to the memory component to implement a cache and operable with the memory controller and the memory component in each of a plurality of operating modes comprising a first operating mode and a second operating mode having

differing addressing and timing requirements for accessing the data group;

the first operating mode having placement of each of at least two ways of a data group in differing rows in the memory component, with tag access and data access not overlapped; and

the second operating mode having placement of all ways of a data group in a same row in the memory component, with tag access and data access overlapped.

14. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the memory controller operable to transmit, to the DRAM, a command/address (CA) packet comprising a chip select (CS), a bank address field, a row address field, a tag compare field, and a command opcode, to perform a tag row activation operation and subsequent data row operation in the DRAM wherein the tag row activation operation and the data row operation are not overlapped in the first operating mode and wherein the tag row activation operation and the data row operation are overlapped in the second operating mode.

15. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the memory controller operable to transmit, to the DRAM, a command/address (CA) packet comprising a chip select (CS) for device or rank selection on a channel, a bank address field, a row address field, a tag compare field, and a command opcode, to perform a tag column access operation and subsequent data column access operation wherein the tag compare field of the tag column access operation and the tag compare field of the data column access operation are shared.

16. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the memory controller to repurpose one or more unused tag rows in the DRAM for other use comprising ECC code or encryption key.

17. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the memory controller operable to transmit, to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field, wherein the chip select (CS) for device or rank select on channel in the DRAM depends upon the sub-group field in one mode, and the chip select for device or rank select on channel in the

DRAM depends upon the group field in another mode.

18. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the memory controller operable to transmit, to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field, wherein a bank address field for the DRAM depends upon the sub-group field in one mode, and the bank address field for the DRAM depends upon the group field in another mode.

19. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row; and

the memory controller operable to transmit, to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field, wherein a row address field depends upon the sub-group field, the row address field depends upon the group field, and a column address field depends upon the sub-group field.

20. The memory system of claim 13, further comprising:

the memory component comprising a dynamic random access memory (DRAM) having the at least one tag row and the at least one data row;

the memory controller operable to transmit, to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field;

the DRAM is operable to form a row address field indirectly from a portion of the tag field in one mode; and

the DRAM is operable to form a column address indirectly from a further portion of the tag field in another mode.

21. A method, comprising:

selecting a first operating mode or a second operating mode from a plurality of operating modes of which a memory controller and a dynamic random access memory (DRAM) are capable; and

operating the memory controller coupled to the DRAM in the selected first operating mode or second operating mode, as a cache having multiple ways to hold a data group as a cache-line or a cache-block;

wherein the first operating mode and the second operating mode have differing addressing and timing for accessing the data group, and wherein the DRAM has at least one data row and at least one tag row different from and associated with the at least one data row.

22. The method of claim 21, wherein:

the first operating mode comprises placement each of at least two ways of a data group in differing rows in the DRAM, with tag access and data access not overlapped; and

the second operating mode comprises placement of all ways of a data group in a same row in the DRAM, with tag access and data access overlapped.

23. The method of claim 21, further comprising:

transmitting, from the memory controller to the DRAM, a command/address (CA) packet comprising a chip select (CS), a bank address field, a row address field, a tag compare field, and a command opcode; and

performing, in the DRAM, a tag row activation operation and subsequent data row operation, wherein the tag row activation operation and the data row operation are not overlapped in the first operating mode and wherein the tag row activation operation and the data row operation are overlapped in the second operating mode.

24. The method of claim 21, further comprising:

transmitting, from the memory controller to the DRAM, a command/address (CA) packet comprising a chip select (CS) for device or rank selection on a channel, a bank address field, a row address field, a tag compare field, and a command opcode; and

performing, in the DRAM, a tag column access operation and subsequent data column access operation, wherein the tag compare field of the tag column access operation is also used for the tag compare field of the data column access operation.

25. The method of claim 21, further comprising:

transmitting, from the memory controller to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field, wherein the chip select (CS) for device or rank select on channel in the DRAM depends upon the sub-group field in one mode, and the chip select for device or rank select on channel in the DRAM depends upon the group field in another mode.

26. The method of claim 21, further comprising:

transmitting, from the memory controller to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field, wherein a bank address field for the DRAM depends upon the sub-group field in one mode, and the bank address field for the DRAM depends upon the group field in another mode.

27. The method of claim 21, further comprising:

transmitting, from the memory controller to the DRAM, a command/address (CA) packet comprising a physical address with a sub-group field, a group field, and a tag field, wherein a row address field depends upon the sub-group field, the row address field depends upon the group field, and a column address field depends upon the sub-group field..




FIG. 1A

Tags in Separate Tag Mat




FIG. 1B




FIG. 2A




FIG. 2B

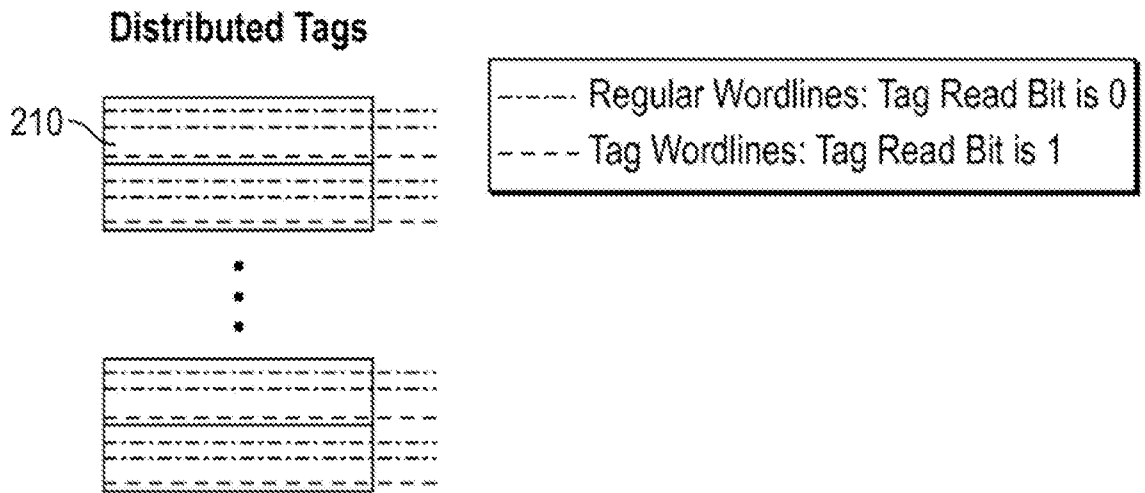


FIG. 2C

DRAM with Cache Mode - Logical Detail




FIG. 3A

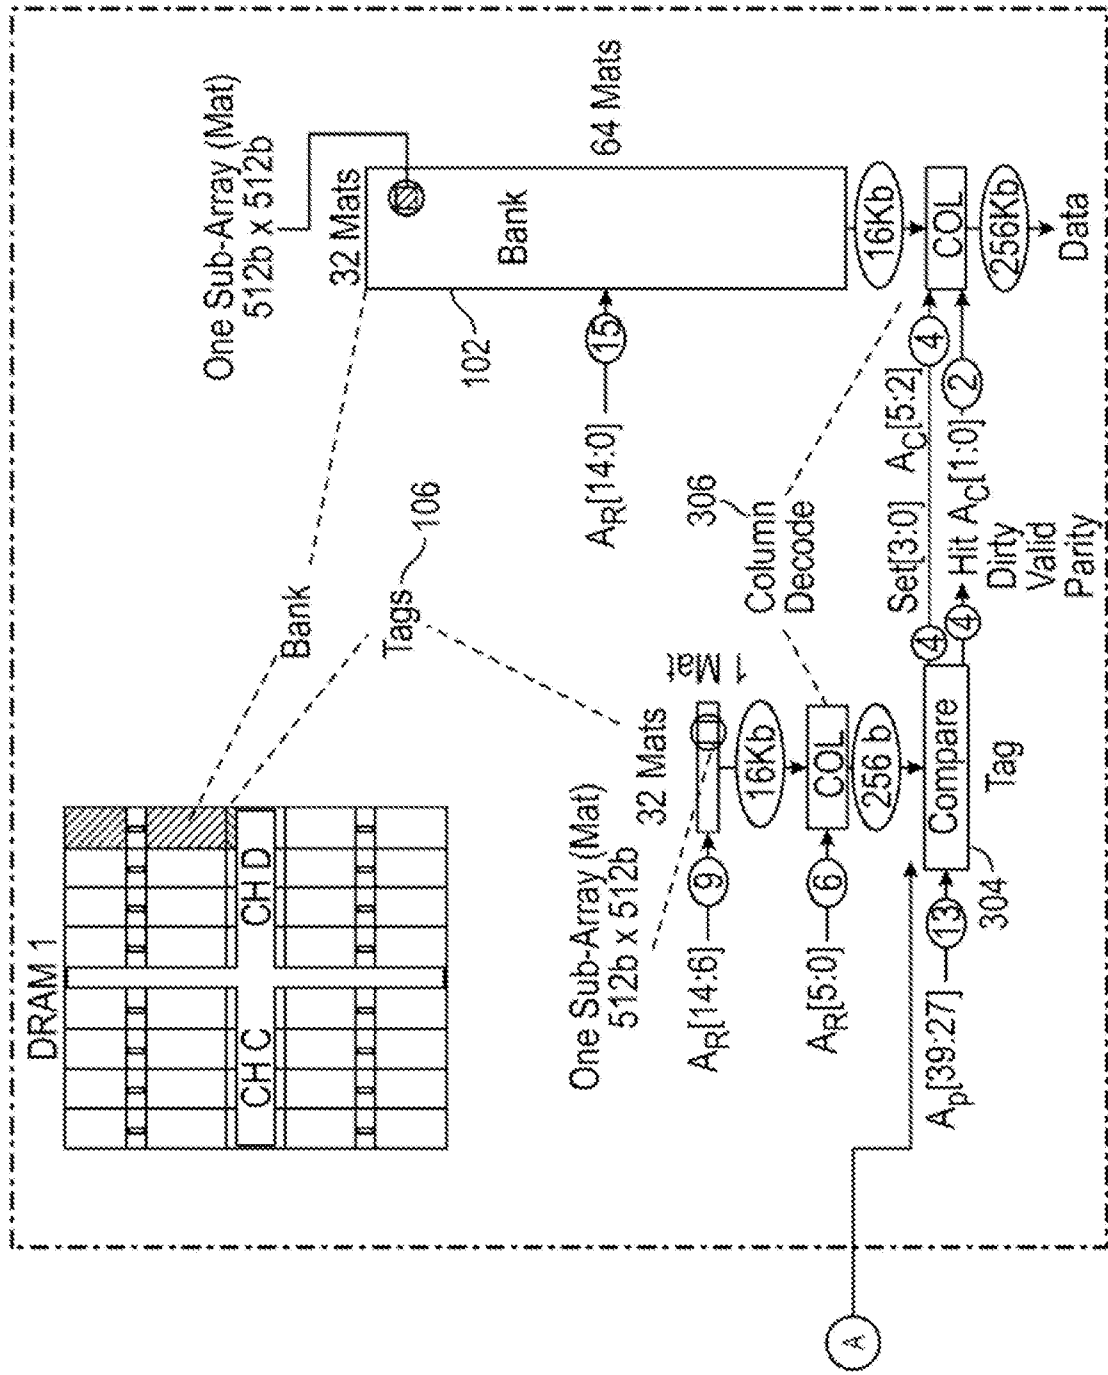


FIG. 3B




FIG. 3C




FIG. 4A




FIG. 4B




FIG. 5B

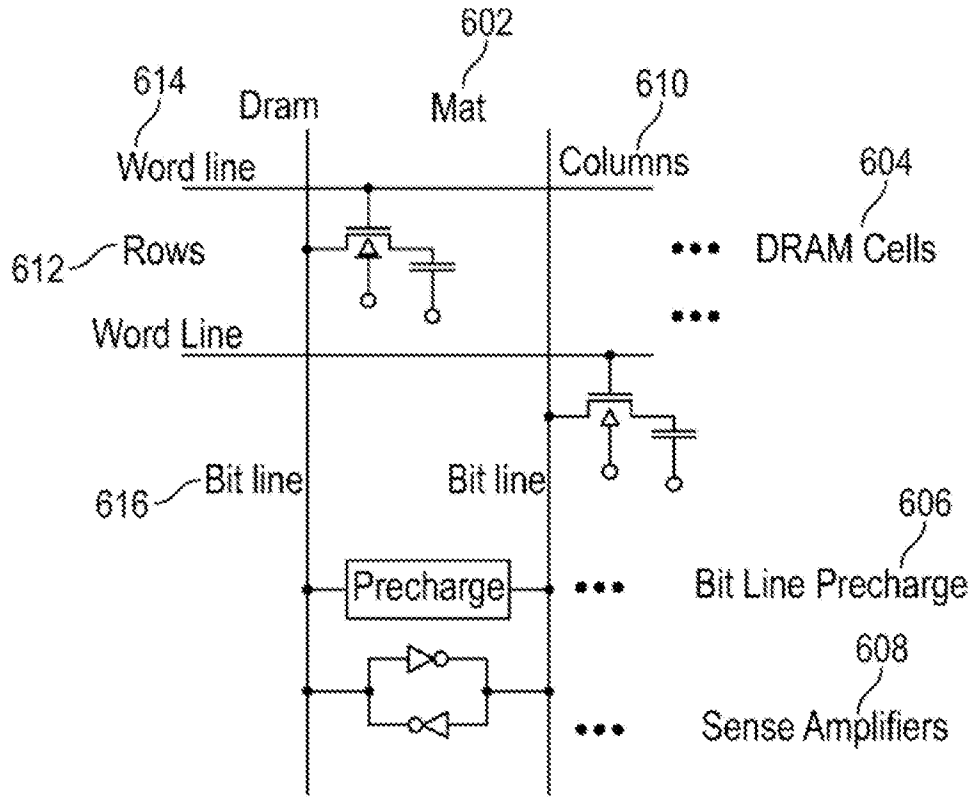


FIG. 6

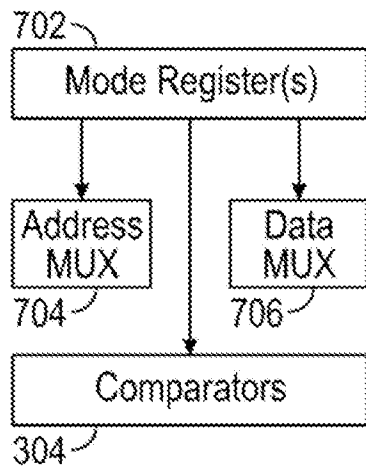


FIG. 7

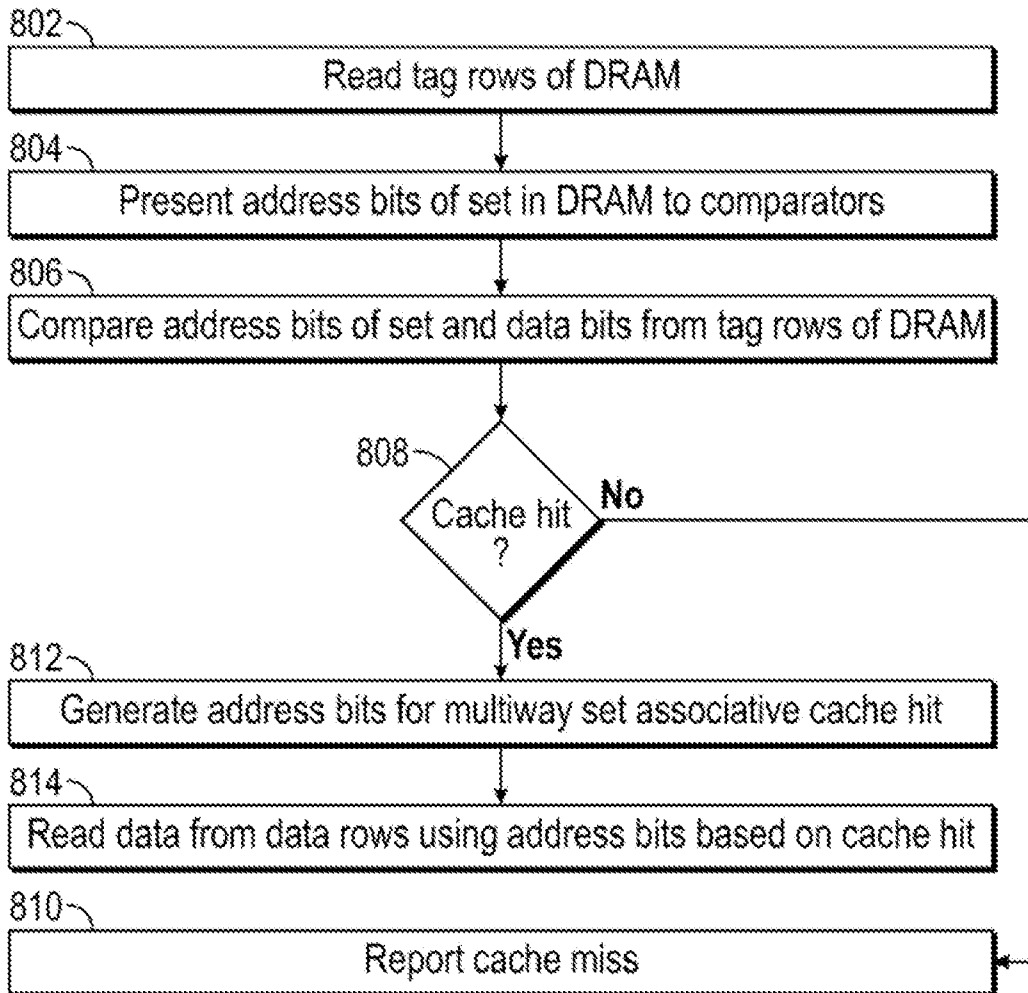


FIG. 8A

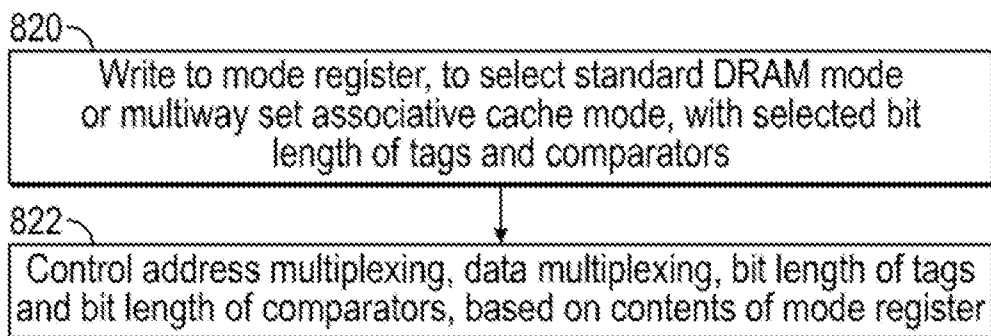


FIG. 8B




FIG. 9

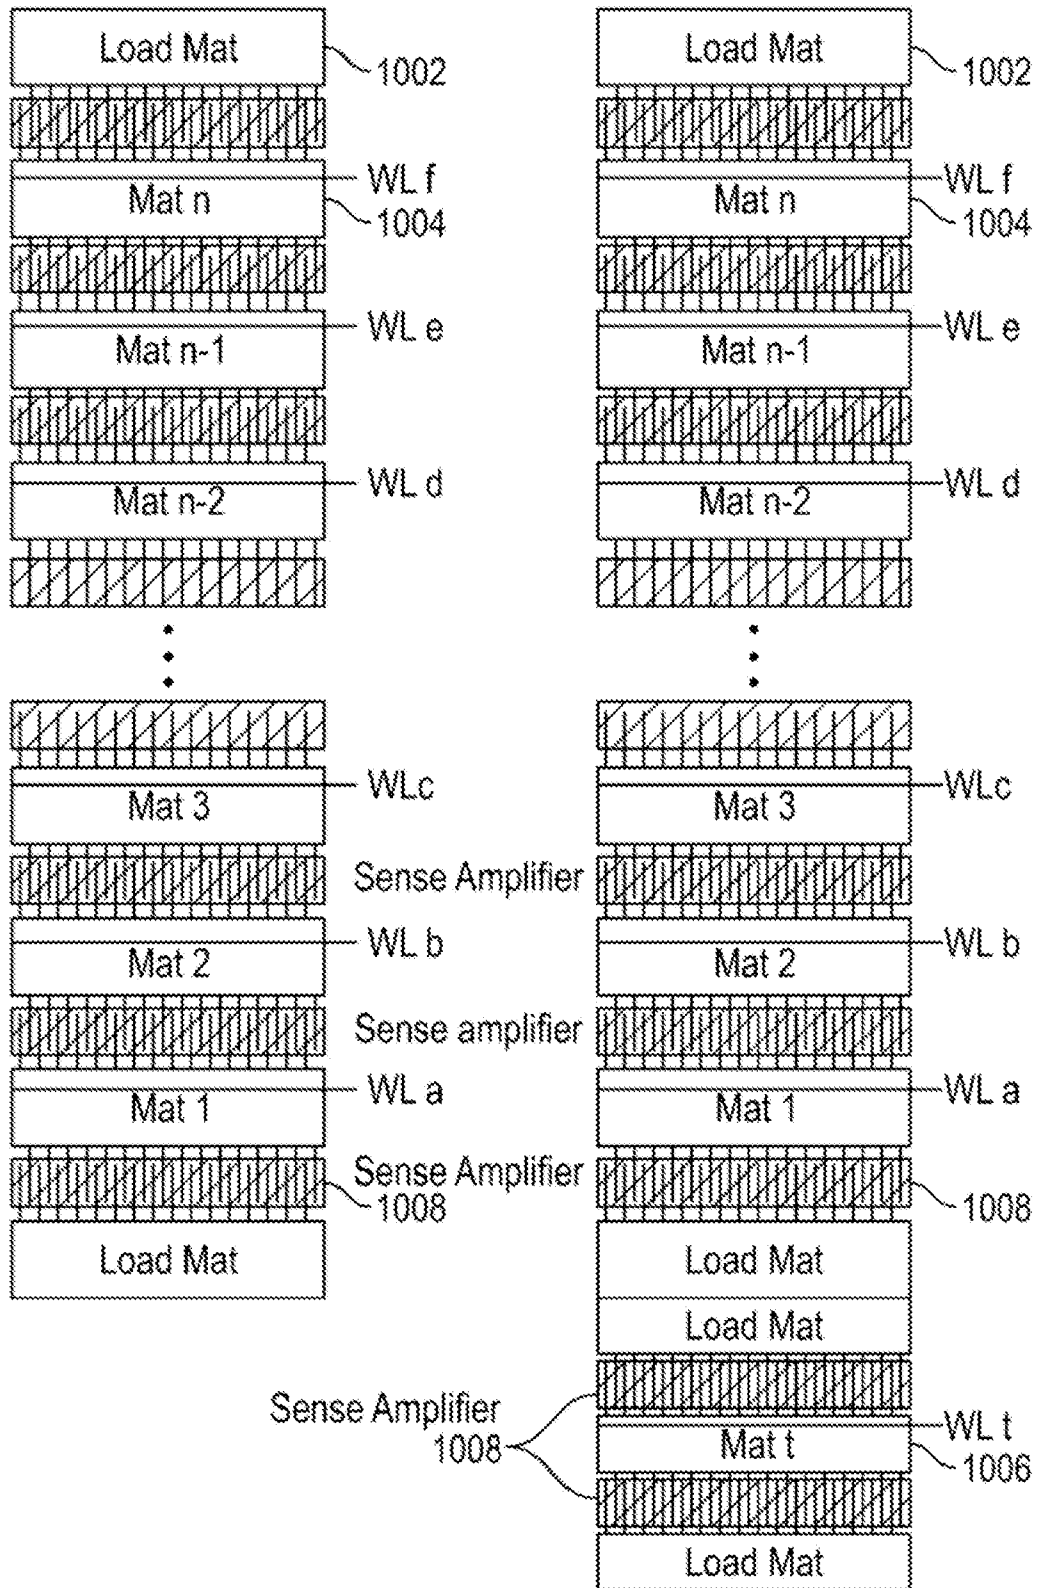


FIG. 10




FIG. 11

Controller Using DRAM with Cache Mode-System Configuration Example




FIG. 12




FIG. 13A




FIG. 13B




FIG. 13C




FIG. 13D




FIG. 14A




FIG. 14B

Controller Utilizing DRAM with Cache Mode - Configuration Example Z-A




FIG. 14C






FIG. 14D



Command Format Example 3-(All Configs)




FIG. 15C

4Gb/s DRAM Links - 64B RD (Read) Timing Details for MISS




FIG. 16




FIG. 18

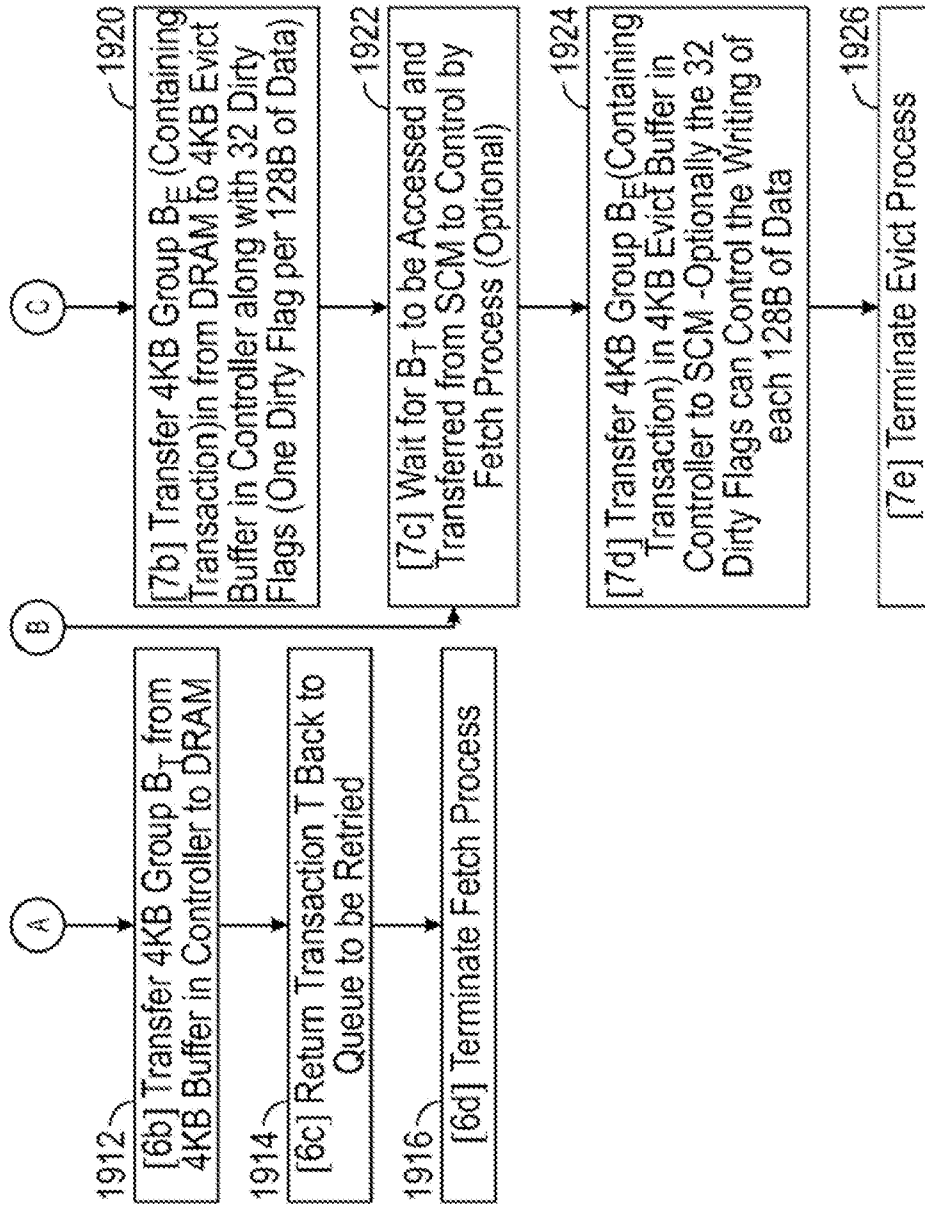


FIG. 18
(Continued)

Eviction Decision Process- "Touch" Flag Enhancement




FIG. 20

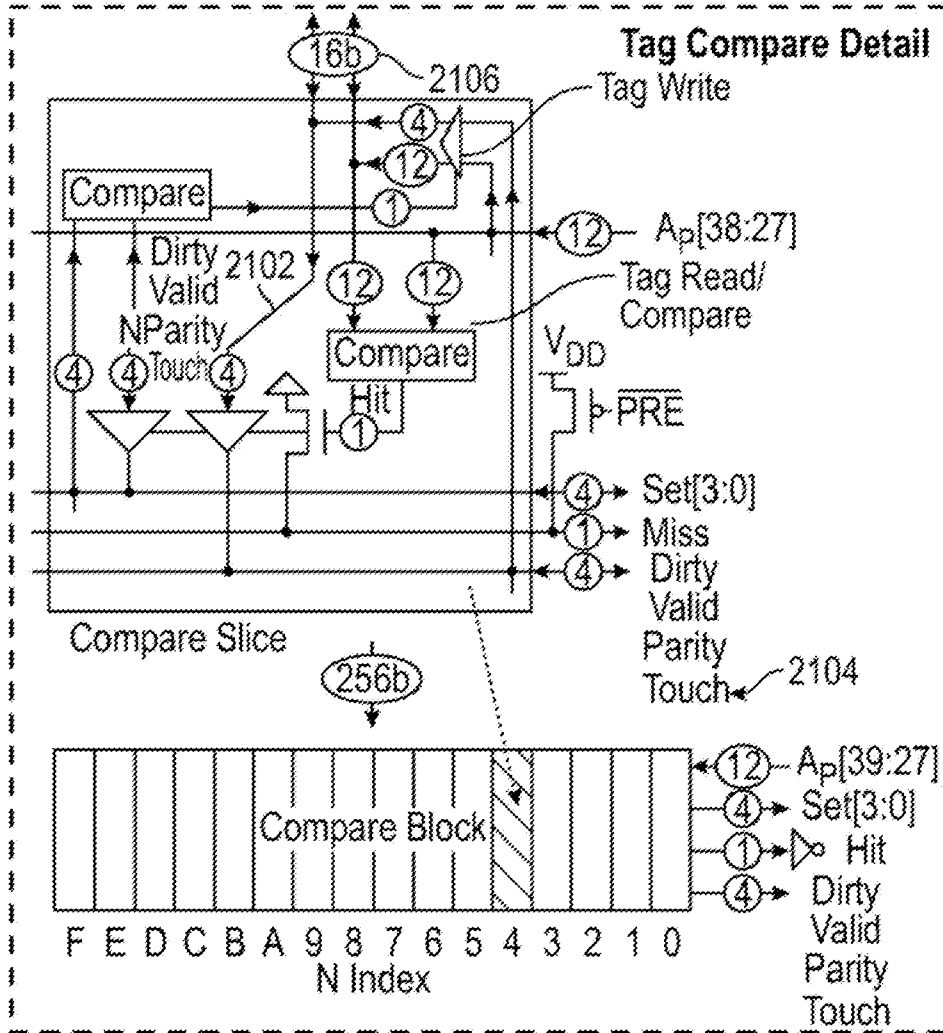


FIG. 20
(Continued)

4KB Cache Line Tags and Data for 16 Sets (Ways) Located in Same DRAM Row




FIG. 21B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 20/22907

A. CLASSIFICATION OF SUBJECT MATTER

IPC - G06F 12/08 (2020.01)

CPC - G06F 12/0893; G06F 12/0895; G06F 12/0802; G06F 12/0877

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)

See Search History document

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

See Search History document

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2013/0138892 A1 (LOH et al.) 30 May 2013 (30.05.2013) Fig 1, 4, 5, abstract, para [0022], [0030], [0031], [0062]-[0069], [0072]-[0074]	1-7
Y	US 2013/0318307 A1 (RABINOVITCH et al.) 28 November 2013 (28.11.2013) abstract, para [0031]	1-7
Y	US 2018/0267911 A1 (RAMBUS INC.) 20 September 2018 (20.09.2018) Fig 1, abstract, para [0039]-[0041]	4
Y	US 5,617,347 A (LAURITZEN) 01 April 1997 (01.04.1997) abstract, col 6, ln 11-18	5, 6

 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

"D" document cited by the applicant in the international application

"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

09 July 2020

Date of mailing of the international search report

28 JUL 2020

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Lee Young

Telephone No. PCT Helpdesk: 571-272-4300

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 20/22907

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be searched, the appropriate additional search fees must be paid.

Group I: Claims 1-7 drawn to a memory controller including details of cache results logic and cache fetch logic.

Group II: Claims 8-27 drawn to a memory controller, memory system, or method including details of a first and second operating mode.

--see extra sheet

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-7

- Remark on Protest**
- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
 - The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
 - No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 20/22907

Continuation of Box No V -- Observations where unity of invention is lacking

The inventions listed as Groups I and II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

Special Technical Features:

Group I includes the special technical feature of cache results logic to receive a response from the first memory component across the first memory bus with the data or an indication of a miss resulting from comparing the tag compare field with the tag field stored in the row of the first memory component; and cache fetch logic to send a cache fetch command across a second memory bus to a second memory component, in response to the indication of the miss, the cache fetch command to cause the second memory component to send the data associated with the tag field stored in the row of the first memory component to the memory controller, not included in the other group.

Group II includes the special technical features of a first operating mode and a second operating mode having differing addressing and timing requirements for accessing a data group, not included in the other group.

Common Technical Features:

The only technical features shared by Groups I-II that would otherwise unify the groups, are that both groups are related to a memory controller including a memory cache and a row address for a row which stores a tag field associated with data. However, these shared technical features do not represent a contribution over prior art, because the shared technical features are disclosed by US 2013/0138892 A1 to Loh et al. (hereinafter 'Loh').

Loh discloses a memory controller including a memory cache and a row address for a row which stores a tag field associated with data (abstract).

As the shared technical features were known in the art at the time of the invention, they cannot be considered special technical features that would otherwise unify the groups.

Therefore, Groups I-II lack unity under PCT Rule 13.