



- (51) International Patent Classification:
G11C 7/06 (2006.01) *G11C 5/14* (2006.01)
- (21) International Application Number:
PCT/US2016/015027
- (22) International Filing Date:
27 January 2016 (27.01.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/112,843 6 February 2015 (06.02.2015) US
- (71) Applicant: MICRON TECHNOLOGY, INC. [US/US];
8000 S Federal Way, Boise, Idaho 83716-9632 (US).
- (72) Inventors: ZAWODNY, Jason T.; 978 E. Covey Run Ct.,
Eagle, Idaho 83616 (US). DOBELSTEIN, Kelley; 7462
4th Avenue NE, Seattle, Washington 98115 (US). FINK-
BEINER, Timothy P.; 5821 S Sweet Gum Way, Boise,
Idaho 83716 (US). MURPHY, Richard C.; 13929 W.
Battenberg Court, Boise, Idaho 83713 (US).
- (74) Agents: BROOKS III, Edward J. et al.; Brooks, Cameron
& Huebsch, PLLC, 1201 Marquette Avenue S, Suite 400,
Minneapolis, Minnesota 55403 (US).
- (81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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).

- 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))

(54) Title: APPARATUSES AND METHODS FOR SCATTER AND GATHER





Fig. 2

(57) Abstract: The present disclosure includes apparatuses and methods related to scatter/gather in a memory device. An example apparatus comprises a memory device that includes an array of memory cells, sensing circuitry, and a memory controller coupled to one another. The sensing circuitry includes a sense amplifier and a compute component configured to implement logical operations. A channel controller is configured to receive a block of instructions, the block of instructions including individual instructions for at least one of a gather operation and a scatter operation. The channel controller is configured to send individual instructions to the memory device and to control the memory controller such that the at least one of the gather operation and the scatter operation is executed on the memory device based on a corresponding one of the individual instructions.



Apparatuses and Methods for Scatter and Gather

Technical Field

[0001] The present disclosure relates generally to semiconductor memory and methods, and more particularly, to apparatuses and methods related to implementing scatter and gather operations in a memory array.

Background

[0002] Memory devices are typically provided as internal, semiconductor, integrated circuits in computers or other electronic systems. There are many different types of memory including volatile and non-volatile memory. Volatile memory can require power to maintain its data (e.g., host data, error data, etc.) and includes random access memory (RAM), dynamic random access memory (DRAM), static random access memory (SRAM), synchronous dynamic random access memory (SDRAM), and thyristor random access memory (TRAM), among others. Non-volatile memory can provide persistent data by retaining stored data when not powered and can include NAND flash memory, NOR flash memory, and resistance variable memory such as phase change random access memory (PCRAM), resistive random access memory (RRAM), and magnetoresistive random access memory (MRAM), such as spin torque transfer random access memory (STT RAM), among others.

[0003] Electronic systems often include a number of processing resources (e.g., one or more processors), which may retrieve and execute instructions and store the results of the executed instructions to a suitable location. A processor can comprise a number of functional units such as arithmetic logic unit (ALU) circuitry, floating point unit (FPU) circuitry, and/or a combinatorial logic block, for example, which can be used to execute instructions by performing logical operations such as AND, OR, NOT, NAND, NOR, and XOR, and invert (e.g., inversion) logical operations on data (e.g., one or more operands). For example, functional unit circuitry may be used to perform arithmetic operations such as addition, subtraction, multiplication, and/or division on operands via a number of logical operations.

[0004] A number of components in an electronic system may be involved in providing instructions to the functional unit circuitry for execution.

The instructions may be generated, for instance, by a processing resource such as a controller and/or host processor. Data (e.g., the operands on which the instructions will be executed) may be stored in a memory array that is accessible by the functional unit circuitry. The instructions and/or data may be retrieved from the memory array and sequenced and/or buffered before the functional unit circuitry begins to execute instructions on the data. Furthermore, as different types of operations may be executed in one or multiple clock cycles through the functional unit circuitry, intermediate results of the instructions and/or data may also be sequenced and/or buffered.

[0005] In many instances, the processing resources (e.g., processor and/or associated functional unit circuitry) may be external to the memory array, and data is accessed via a bus between the processing resources and the memory array to execute a set of instructions. Processing performance may be improved in a processor-in-memory device, in which a processor may be implemented internal and/or near to a memory (e.g., directly on a same chip as the memory array). A processing-in-memory device may save time by reducing and/or eliminating external communications and may also conserve power.

Brief Description of the Drawings

[0006] Figure 1A is a block diagram of an apparatus in the form of an electronic system including a memory device in accordance with a number of embodiments of the present disclosure.

[0007] Figure 1B is another block diagram of an apparatus in the form of an electronic system including a memory device in accordance with a number of embodiments of the present disclosure.

[0008] Figure 2 illustrates a block diagram of an apparatus for gather and/or scatter in a memory device in accordance with a number of embodiments of the present disclosure.

[0009] Figures 3-6 illustrate a number of embodiments for a high speed interface between a memory device and a host via a channel controller in accordance with the present disclosure.

[0010] Figure 7 is a schematic diagram illustrating sensing circuitry to a memory device in accordance with a number of embodiments of the present disclosure.

[0011] Figure 8 is a schematic diagram illustrating sensing circuitry to a memory device in accordance with a number of embodiments of the present disclosure.

[0012] Figure 9 is a logic table illustrating selectable logic operation results implemented by a sensing circuitry shown in Figure 8 in accordance with a number of embodiments of the present disclosure.

Detailed Description

[0013] The present disclosure includes apparatuses and methods related to scatter/gather in a memory array. In one embodiment, the apparatus comprises a memory device coupled to a host via a data bus and a control bus. The memory device includes an array of memory cells and sensing circuitry coupled to the array via a plurality of sense lines. The sensing circuitry includes sense amplifiers and a compute component configured to implement logical operations.

[0014] A memory controller is coupled to the array and sensing circuitry. The memory controller is configured to receive a block of instructions, including a gather and/or scatter operation from a channel controller associated with the host and memory device. The memory controller is further configured to subdivide the block of instructions so that operations before, between and after a local gather operation can be executed locally on the memory device.

[0015] Scatter and/or gather (hereinafter “scatter/gather”), also known as vectored Input/Output (I/O), is a technique used in applications accessing large data sets. When applied to large data sets, memory array access patterns generated by scatter/gather operations can often show very low ordering and locality of successive accesses. Data accesses can be dispersed all over the memory array, which can be inefficient. Embodiments of the present disclosure provide an apparatus and methods for implementing an efficient scatter/gather technique involving a memory device having processing-in-memory (PIM) capabilities.

[0016] Scatter can be generally expressed as:

$$\text{for } (i = 0; i < N; i++) \\ y[idx[i]] = x[i]$$

and gather can be generally expressed as:

$$\text{for } (i = 0; i < N; i++)$$

$$x[i] = y[idx[i]]$$

where i is an index (e.g., counter, pointer) and N represents a number of indices used to store a quantity of data in a memory array. The expression $idx[i]$ represents an index vector comprised of indices. As such, the requested data, $x[i]$, is a function (y) of the index vector. In this manner, the index vector contains information by which pieces of the requested data are stored/found in the memory array.

[0017] For small-sized index vectors, the operations to extract and process the indices do not pose a significant challenge for modern computing system and memory device. However, when the aggregate size of the index vectors become larger than the caching ability of the computing system, performance of the scatter/gather operations decreases significantly due to data being fetched into cache, used a small number of times (e.g., once) and thereafter flushed out of cache having to be re-read into cache again later. However where large quantities of data are being requested, this inefficiency is more than just an inadequate cache size. Giga Updates Per Second (GUPS) is an example of one recognized benchmark to measure performance of a computing system and memory device for scatter/gather operations.

[0018] The index vector (e.g., $idx[i]$) can be thought of as being “pathological” in that the pattern of data elements comprising the requested data, and to which a pattern of index vector elements corresponds, can impact the efficiency and performance of scatter/gather operations. Previous approaches to scatter/gather operations utilized simple “tricks” to optimize the overall scatter/gather operation. However, as the amount of data being requested increases, the quantity of possible patterns in the index vector can also increase. The index vector can take on many index patterns that can confound some simple tricks previously utilized to optimize scatter/gather operations. In order to appreciate the improved scatter/gather techniques and apparatus for implementing such techniques, a discussion of a memory device having PIM capabilities, and associated host, follows.

[0019] According to various embodiments, performance of scatter/gather operations involving a memory device having PIM capabilities can distribute implementation of a scatter/gather operation over multiple sensing circuitries that can implement logical operations within the memory array. Scatter/gather

operations involving a memory device having PIM capabilities can be completed in less time and using less power. Some time and power advantage can be realized by reducing the amount of data that is moved around a computing system to process the requested memory array operations (e.g., reads, writes, etc.).

[0020] A number of embodiments of the present disclosure can provide improved parallelism and/or reduced power consumption in association with performing compute functions as compared to previous systems such as previous PIM systems and systems having an external processor (e.g., a processing resource located external from a memory array, such as on a separate integrated circuit chip). For instance, a number of embodiments can provide for performing fully complete compute functions such as integer add, subtract, multiply, divide, and CAM (content addressable memory) functions without transferring data out of the memory array and sensing circuitry via a bus (e.g., data bus, address bus, control bus), for instance. Such compute functions can involve performing a number of logical operations (e.g., logical functions such as AND, OR, NOT, NOR, NAND, XOR, etc.). However, embodiments are not limited to these examples. For instance, performing logical operations can include performing a number of non-Boolean logic operations such as copy, compare, destroy, etc.

[0021] In previous approaches, data may be transferred from the array and sensing circuitry (e.g., via a bus comprising input/output (I/O) lines) to a processing resource such as a processor, microprocessor, and/or compute engine, which may comprise ALU circuitry and/or other functional unit circuitry configured to perform the appropriate logical operations. However, transferring data from a memory array and sensing circuitry to such processing resource(s) can involve significant power consumption. Even if the processing resource is located on a same chip as the memory array, significant power can be consumed in moving data out of the array to the compute circuitry, which can involve performing a sense line (which may be referred to herein as a digit line or data line) address access (e.g., firing of a column decode signal) in order to transfer data from sense lines onto I/O lines (e.g., local I/O lines), moving the data to the array periphery, and providing the data to the compute function.

[0022] Furthermore, the circuitry of the processing resource(s) (e.g., compute engine) may not conform to pitch rules associated with a memory array. For example, the cells of a memory array may have a $4F^2$ or $6F^2$ cell size, where “F” is a feature size corresponding to the cells. As such, the devices (e.g., logic gates) associated with ALU circuitry of previous PIM systems may not be capable of being formed on pitch with the memory cells, which can affect chip size and/or memory density, for example. A number of embodiments of the present disclosure include sensing circuitry formed on pitch with an array of memory cells and capable of performing compute functions such as gather and scatter operations local to the array of memory cells.

[0023] In the following detailed description of the present disclosure, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration how one or more embodiments of the disclosure may be practiced. These embodiments are described in sufficient detail to enable those of ordinary skill in the art to practice the embodiments of this disclosure, and it is to be understood that other embodiments may be utilized and that process, electrical, and/or structural changes may be made without departing from the scope of the present disclosure. As used herein, designators such as “N”, “M”, etc., particularly with respect to reference numerals in the drawings, indicate that a number of the particular feature so designated can be included. As used herein, “a number of” a particular thing can refer to one or more of such things (e.g., a number of memory arrays can refer to one or more memory arrays). A “plurality of” is intended to refer to more than one of such things.

[0024] The figures herein follow a numbering convention in which the first digit or digits correspond to the drawing figure number and the remaining digits identify an element or component in the drawing. Similar elements or components between different figures may be identified by the use of similar digits. For example, 206 may reference element “06” in Figure 2, and a similar element may be referenced as 606 in Figure 6. As will be appreciated, elements shown in the various embodiments herein can be added, exchanged, and/or eliminated so as to provide a number of additional embodiments of the present disclosure. In addition, as will be appreciated, the proportion and the relative scale of the elements provided in the figures are intended to illustrate certain

embodiments of the present invention, and should not be taken in a limiting sense.

[0025] Figure 1A is a block diagram of an apparatus in the form of a computing system 100 including a memory device 120 in accordance with a number of embodiments of the present disclosure. As used herein, a memory device 120, memory controller 140, channel controller 143, bank arbiter 145, high speed interface (HSI) 141, memory array 130, sensing circuitry 150, and/or additional logic circuitry 170, as shown in Figures 1A and 1B, might also be separately considered an “apparatus.”

[0026] System 100 includes a host 110 coupled (e.g., connected) to memory device 120, which includes a memory array 130. Host 110 can be a host system such as a personal laptop computer, a desktop computer, a digital camera, a smart phone, or a memory card reader, among various other types of hosts. Host 110 can include a system motherboard and/or backplane and can include a number of processing resources (e.g., one or more processors, microprocessors, or some other type of controlling circuitry). The system 100 can include separate integrated circuits or both the host 110 and the memory device 120 can be on the same integrated circuit. The system 100 can be, for instance, a server system and/or a high performance computing (HPC) system and/or a portion thereof. Although the example shown in Figures 1A and 1B illustrates a system having a Von Neumann architecture, embodiments of the present disclosure can be implemented in non-Von Neumann architectures, which may not include one or more components (e.g., CPU, ALU, etc.) often associated with a Von Neumann architecture.

[0027] For clarity, the system 100 has been simplified to focus on features with particular relevance to the present disclosure. The memory array 130 can be a DRAM array, SRAM array, STT RAM array, PCRAM array, TRAM array, RRAM array, NAND flash array, and/or NOR flash array, for instance. The array 130 can comprise memory cells arranged in rows coupled by access lines (which may be referred to herein as word lines or select lines) and columns coupled by sense lines, which may be referred to herein as data lines or digit lines. Although a single array 130 is shown in Figure 1, embodiments are not so limited. For instance, memory device 120 may include a number of arrays 130 (e.g., a number of banks of DRAM cells, NAND flash cells, etc.).

[0028] The memory device 120 includes address circuitry 142 to latch address signals provided over a data bus 156 (e.g., an I/O bus) through I/O circuitry 144. Status, exception, and other data information can be provided from the memory controller 140 on the memory device 120 to a channel controller 143 via a high speed interface (HSI) (both shown in Figure 1B), including an out-of-band bus 157 (shown in Figure 1B), which in turn can be provided from the memory device 120 to the host 110. Address signals are received through address circuitry 142 and decoded by a row decoder 146 and a column decoder 152 to access the memory array 130. Data can be read from memory array 130 by sensing voltage and/or current changes on the data lines using sensing circuitry 150. The sensing circuitry 150 can read and latch a page (e.g., row) of data from the memory array 130. The I/O circuitry 144 can be used for bi-directional data communication with host 110 over the data bus 156. The write circuitry 148 is used to write data to the memory array 130.

[0029] Memory controller 140, e.g., bank control logic and/or sequencer, decodes signals provided by a control bus 154 from the host 110. These signals can include chip enable signals, write enable signals, and address latch signals that are used to control operations performed on the memory array 130, including data read, data write, and data erase operations. In various embodiments, the memory controller 140 is responsible for executing instructions from the host 110. The memory controller 140 can be a state machine, a sequencer, or some other type of controller. The controller 140 can control shifting data (e.g., right or left) in an array, e.g., memory array 130.

[0030] Examples of the sensing circuitry 150 are described further below. For instance, in a number of embodiments, the sensing circuitry 150 can comprise a number of sense amplifiers and a number of compute components, which may serve as, and be referred to herein as, an accumulator and can be used to perform logical operations (e.g., on data associated with complementary data lines).

[0031] In a number of embodiments, the sensing circuitry 150 can be used to perform logical operations using data stored in array 130 as inputs and store the results of the logical operations back to the array 130 without transferring data via a sense line address access (e.g., without firing a column decode signal). As such, various compute functions can be performed using, and

within, sensing circuitry 150 rather than (or in association with) being performed by processing resources external to the sensing circuitry (e.g., by a processor associated with host 110 and/or other processing circuitry, such as ALU circuitry, located on device 120 (e.g., on controller 140 or elsewhere)).

[0032] In various previous approaches, data associated with an operand, for instance, would be read from memory via sensing circuitry and provided to external ALU circuitry via I/O lines (e.g., via local I/O lines and/or global I/O lines). The external ALU circuitry could include a number of registers and would perform compute functions using the operands, and the result would be transferred back to the array via the I/O lines. In contrast, in a number of embodiments of the present disclosure, sensing circuitry 150 is configured to perform logical operations on data stored in memory array 130 and store the result back to the memory array 130 without enabling an I/O line (e.g., a local I/O line) coupled to the sensing circuitry 150. The sensing circuitry 150 can be formed on pitch with the memory cells of the array. Additional peripheral sense amplifier and/or logic 170 can be coupled to the sensing circuitry 150.

[0033] As such, in a number of embodiments, circuitry external to array 130 and sensing circuitry 150 is not needed to perform compute functions as the sensing circuitry 150 can perform the appropriate logical operations to perform such compute functions without the use of an external processing resource. Therefore, the sensing circuitry 150 may be used to compliment and/or to replace, at least to some extent, such an external processing resource (or at least the bandwidth consumption of such an external processing resource).

[0034] However, in a number of embodiments, the sensing circuitry 150 may be used to perform logical operations (e.g., to execute instructions) in addition to logical operations performed by an external processing resource (e.g., host 110). For instance, host 110 and/or sensing circuitry 150 may be limited to performing only certain logical operations and/or a certain number of logical operations.

[0035] Enabling an I/O line can include enabling (e.g., turning on) a transistor having a gate coupled to a decode signal (e.g., a column decode signal) and a source/drain coupled to the I/O line. However, embodiments are not limited to not enabling an I/O line. For instance, in a number of embodiments, the sensing circuitry (e.g., 150) can be used to perform logical operations

without enabling column decode lines of the array; however, the local I/O line(s) may be enabled in order to transfer a result to a suitable location other than back to the array 130 (e.g., to an external register).

[0036] Figure 1B is a block diagram of another apparatus architecture in the form of a computing system 100 including a plurality of memory devices 120-1, . . . , 120-N coupled to a host 110 via a channel controller 143 in accordance with a number of embodiments of the present disclosure. In at least one embodiment the channel controller 143 may be coupled to the plurality of memory devices 120-1, . . . , 120-N in an integrated manner in the form of a module 118. In an alternative embodiment, the channel controller 143 may be integrated with the host 110, as illustrated by dashed lines 111. The channel controller 143 can be coupled to each of the plurality of memory devices 120-1, . . . , 120-N via a control bus 154 as described in Figure 1A which in turn can be coupled to the host 110. The channel controller 143 can also be coupled to each of the plurality of memory devices, 120-1, . . . , 120-N via a data bus 156 as described in Figure 1A which in turn can be coupled to the host 110. In addition, the channel controller 143 can be coupled to each of the plurality of memory devices 120-1, . . . , 120-N via an out-of-band (OOB) bus 157 associated with a high speed interface (HSI) 141, described more in connection with Figures 3-6, that is configured to report status, exception and other data information to the channel controller 143 to exchange with the host 110.

[0037] As shown in Figure 1B, the channel controller 143 can receive the status and exception information from a high speed interface (HSI) (also referred to herein as a status channel interface) 141 associated with a bank arbiter 145 in each of the plurality of memory devices 120-1, . . . , 120-N. In the example of Figure 1B, each of the plurality of memory devices 120-1, . . . , 120-N can include a bank arbiter 145 to sequence control and data with a plurality of banks, e.g., Bank zero (0), Bank one (1), . . . , Bank six (6), Bank seven (7), etc. Each of the plurality of banks, Bank 0, . . . , Bank 7, can include a memory controller 140 and other components, including an array of memory cells 130 and sensing circuitry 150, peripheral logic 170, etc., as described in connection with Figure 1A.

[0038] For example, each of the plurality of banks, e.g., Bank 0, . . . , Bank 7, in the plurality of memory devices 120-1, . . . , 120-N can include address

circuitry 142 to latch address signals provided over a data bus 156 (e.g., an I/O bus) through I/O circuitry 144. Status and/or exception information can be provided from the memory controller 140 on the memory device 120 to the channel controller 143, using the OOB bus 157, which in turn can be provided from the plurality of memory devices 120-1, . . . , 120-N to the host 110. For each of the plurality of banks, e.g., Bank 0, . . . , Bank 7, address signals can be received through address circuitry 142 and decoded by a row decoder 146 and a column decoder 152 to access the memory array 130. Data can be read from memory array 130 by sensing voltage and/or current changes on the data lines using sensing circuitry 150. The sensing circuitry 150 can read and latch a page (e.g., row) of data from the memory array 130. The I/O circuitry 144 can be used for bi-directional data communication with host 110 over the data bus 156. The write circuitry 148 is used to write data to the memory array 130 and the OOB bus 157 can be used to report status and exception information to the channel controller 143.

[0039] As described next in connection with Figure 2, the channel controller 143 can include one or more local buffers 161 to store an index vector and bin (e.g., separately store) bank commands and indices for the various banks associated with each of the plurality of memory devices 120-1, . . . , 120-N in connection with a gather and/or scatter operation. The channel controller 143 can send commands to the plurality of memory devices 120-1, . . . , 120-N and field return results/data from such operations. As described herein, the return results/data can be returned to the channel controller 143 via the OOB bus 157 associated with the status channel interface on each of the plurality of memory devices 120-1, . . . , 120-N.

[0040] For clarity, the system 100 has been simplified to focus on features with particular relevance to the present disclosure. The memory array 130 can be a DRAM array, SRAM array, STT RAM array, PCRAM array, TRAM array, RRAM array, NAND flash array, and/or NOR flash array, for instance. The array 130 can comprise memory cells arranged in rows coupled by access lines (which may be referred to herein as word lines or select lines) and columns coupled by sense lines, which may be referred to herein as data lines or digit lines. Although a single array 130 is shown in Figure 1, embodiments are

not so limited. For instance, memory device 120 may include a number of arrays 130 (e.g., a number of banks of DRAM cells, NAND flash cells, etc.).

[0041] As in Figure 1A, a memory controller 140, e.g., bank control logic and/or sequencer, associated with any particular bank, Bank 0, . . . , Bank 7, in a given memory device, 120-1, . . . , 120-N, can decode signals provided by control bus 154 from the host 110. These signals can include chip enable signals, write enable signals, and address latch signals that are used to control operations performed on the memory array 130, including data read, data write, and data erase operations. In various embodiments, the memory controller 140 is responsible for executing instructions from the host 110. And, as above, the memory controller 140 can be a state machine, a sequencer, or some other type of controller. That is, the controller 140 can control shifting data (e.g., right or left) in an array, e.g., memory array 130.

[0042] Figure 2 illustrates a block diagram example of an operations flow for a gather and/or scatter operation in an apparatus in accordance with a number of embodiments of the present disclosure. According to various embodiments, scatter operations and gather operations can be similar. Reference may be made to the channel controller 143, bank arbiter 145 coupled to a plurality of memory devices 120-1, . . . , 120-N and high speed interface 141 (example embodiments shown in more detail in Figures 3-6) of Figure 1B together with the memory controller 140, array of memory cells 130 and sensing circuitry 150, including sense amplifiers and a compute component to implement logical functions, as shown in Figures 1A (sensing circuitry shown in more detail in Figures 7 and 8). The following provides a detailed description of a gather operation, with aspects of a scatter operation that are unique from the described gather operation being discussed after the discussion of the gather operation.

[0043] A host 110 including an integrated (in 111 dashed line form) channel controller 143 (e.g., scaler processor), or alternatively a channel controller 143 integrated in module form, can send a block of instructions that include individual instructions (e.g., commands) for at least one of a gather operation and a scatter operation to a plurality of memory devices 120-1, . . . , 120-N (shown in Figure 1B). For example, a read request can involve one or more gather operations. The host can provide the block of instructions to the

memory device as a block within an index (*idx*) vector 200. The index vector 200 can be stored in an array of memory cells 130 in one or more of a plurality of banks of memory cells, e.g., Bank 0, . . . , Bank 7, in one or more memory devices, 120-1, . . . , 120-N. As described above, the array of memory cells 130 can be an array of dynamic random access memory (DRAM) cells. A memory controller, e.g., memory controller 140 to each bank in a plurality of banks, on a given memory device, 120-1, . . . , 120-N, can process the instructions in various ways, including changing an index vector used at the programming level into information that can be used by the given memory device. A bank arbiter 145 to each memory device (e.g., memory chip) can control and pass command instructions to the memory controllers 140 in the plurality of banks and aggregate status information from the banks to report back to the channel controller 143. The index vector can contain data as well location information associated with the data. The location information can be, for example, a series of indices (e.g., “vector” references). The indices can define location within a particular memory device structure, such as bank(s), subarray(s), row(s), and/or column(s).

[0044] In one embodiment, the channel controller 143 can extract an individual instruction comprising indices (e.g., individual vectors) from the index vector 200, and load the indices to a local *idx* buffer 202. The buffer 161 of the channel controller 143 shown in Figure 1B, for example, may include the local *idx* buffer 202. Buffer 161 of the channel controller 143 in Figure 1B may also include the independent device and bank command buffers, e.g., 204-1, 204-2, . . . , 204-N and the reassembly buffer 261 discussed in more detail below.

[0045] Typically, instructions executed in the memory device (e.g., memory chips), or portion of the memory device (e.g., bank), can be limited in scope to data contained in the memory device or portion of the memory device. Since it cannot be certain that a gather operation is confined to data solely within a particular portion of the memory device 120 (e.g., bank), a bank arbiter can orchestrate the gather operation among a plurality of memory controllers 140 associated with a plurality of memory portions (e.g., banks). As shown in Figure 2, the channel controller 143 can store, e.g., “bin”, indices from the *idx* vector according to memory device 120 structure and create independent device and bank command buffers, e.g., 204-1, 204-2, . . . , 204-N, to store individual

instructions among the block of instructions to fetch and return data on an OOB bus 157 associate a high speed interface 141 to each memory device 120. The channel controller 143 can send the individual instruction (e.g., commands) to the memory device to fetch source data, e.g., x-source vector 201 in an array or memory cells in one or more banks in one or more memory devices.

[0046] The memory controller 140 to a given bank can de-reference locations of indices (e.g., individual vectors) in the block of instructions. The memory controller 140 is configured such that individual instructions in the received block of instructions can be subdivided and sent to a memory controller 140 such that operations before, between, and after local gather and/or scatter operations can be executed locally in a particular portion of a given memory device 120. The individual instructions subdivided out of the block of instructions can be referred to herein as a local instruction set.

[0047] The memory controller 140 can manage execution of the local instruction set on the associated portion of the memory device (e.g., bank). Execution can be performed by the sensing circuitry 150, e.g., PIM RAM execution units 250-1, . . . , 250-N in respective banks, performing logical operations and/or functions as the same is described herein. The memory controller 140 can monitor the local instruction set for completion in order to perform the gather operation at an appropriate time to ensure local, in-order execution. In this way, the memory controller 140 can manage an appropriate result for a particular data request, e.g., x-source vector in DRAM 201. Further, the memory controller can ensure that all other operations of a portion of the memory device are complete before starting a gather operation.

[0048] According to various embodiments, the memory controller 140 manages the gather operation to obtain portions of requested data, e.g., x-source vector in DRAM 201, from different portions of the memory device and/or different portions within a bank, e.g., subarray. As the reader will appreciate, the bank arbiter 145 can manage the gather operation from among each memory controller 140 to obtain portions of requested data from among the plurality of banks, e.g., Bank 0, . . . , Bank 7, of the memory device 120. The channel controller 143 can iterate through the indices of the index vector to “stream” 257 a collection of data from the portions of the memory device 120 comprising the requested data, e.g., using the bank arbiter 145, high speed interface 141 and the

out-of-band (OOB) bus 157 to field the requested data to the channel controller 143 and/or return the requested data to the host 110. For example, the channel controller 143 can allocate buffers 161, shown as individual bank command buffers (BANK CMD BUFFER) 204-1, 204-2, . . . , 204-N in Figure 2, to hold particular indices, which are portions of the index vector 200. The buffer 161, individual bank command buffers 204-1, 204-2, . . . , 204-N, can be local to the channel controller 143 and/or integrated buffers at respective portions of the memory device 120. The channel controller 143 may additionally allocate local buffers 161 to hold intermediate results from the x-source vector 201 reads, and can allocate a reassembly buffer 261 to hold the merged results for the requested data, e.g., y-target vector in DRAM 203.

[0049] As the channel controller 143 encounters a gather operation in a block of memory device instructions requested by the host 110, the channel controller 143 can implement the gather operation by reading a local index (*idx*) buffer 202 block by block (e.g., index register storing the *idx* vector 200). Each memory controller 140 associated with banks of the memory device 120 can locally store a moderately-sized buffer of relevant data corresponding to a portion of the index vector 200 distributed by the bank arbiter 145 to that particular bank. For example, each memory controller 140 may include a storage area as a buffer and may include a register. The memory controller 140 obtains the portion of requested data, e.g., the x-source vector portion, pointed to by an individual index, relevant to that bank. The memory controller 140 can group requests for portions of the requested data together by location, e.g., subarray, such that the portion of the memory device, e.g., bank, can execute instructions in an efficient manner to gather the necessary data and return it 257 to the channel controller 143 via the OOB bus 157 (also referred to as channel).

[0050] According to various embodiments, a portion of the requested data associated with a particular portion of the memory device 120, e.g., bank, is returned to the channel controller 143 via the out-of-band (OOB) bus 157 to avoid resource contention with further reads of the index vector 200. The OOB bus 157 can also be used to avoid resource contention with further writing of data to the y-target vector 203.

[0051] According to some embodiments, as data is returned to the channel controller via the OOB bus 157, each portion of the requested data can

be tagged 261 as to its intended destination within the requested data, e.g., within the y-target vector 203. The channel controller 143 can assemble the portions of data obtained from various locations into a coherent writes to burst into target locations, e.g., y-target vector in DRAM 203.

[0052] Using the OOB channel to retrieve portions of the requested data avoids use of the data bus 156 shown in Figures 1A and 1B. In this manner, the “irregular” pattern of reads is not performed directly over the data bus 156, but rather is offloaded to the individual memory devices 120-1, . . . , 120-N. The memory devices 120-1, . . . , 120-N can use much less time and energy gathering only the specific data of the channel controller 143, whereas alternatively the channel controller 143 and/or host 110 would be wasting much of the data read over the normal data bursts, e.g., a normal DRAM data burst, if the channel controller 143 and/or host 110 were reading the gather operation information directly over the data bus 156. For example, time and energy will not be used, interrupting data reads over the normal data bursts, to also accommodate reading the data gather operation information over the data bus 156. Thus, according to various embodiments, sensing circuitry 150 gathers portions of the requested information, and each memory controller 140 communicates this information to the channel controller 143 over the OOB bus 157.

[0053] A scatter operation differs from a gather operation in that the portions of the memory device 120, e.g., sensing circuitry 150, operate to store portions of the requested data into the y-target vector 203, e.g., the host requests each memory device 120 to store the requested data into the y-target vector in DRAM. In the scatter operation, the channel controller 143 will be reading locations to be written to from the idx vector 200 and the data to be written from the x-source vector 201. The channel controller 143 will then group write requests by locality, e.g., bin indices according to PIM RAM structures (e.g., individual PIM RAM execution units (memory controllers 140)) and create individual instructions, stored in local command bank buffers 204-1, . . . , 204-N, to send those individual instructions to appropriate memory controllers 140 in respective memory devices 120-1, . . . , 120-N. The channel controller 143 will provide a stream of instructions to be executed by the sensing circuitry 150 of the respective portions of the memory device 120-1, . . . , 120-N that contains embedded data and specific locations such that each memory device 120 can

execute these local write requests independent of the host 110 and/or channel controller 143, e.g., can be executed “offline”. As used herein, “offline” refers to operations implemented by the sensing circuitries 150 of the respective portions of memory devices 120-1, . . . , 120-N rather than by a data read and/or write involving the data bus 156. With respect to scatter operations, the OOB bus 157 can be used to ensure the local individual instructions to each memory device 120/memory controller 140, e.g., local writes, are completed.

[0054] Figures 3-6 illustrates a series of block diagrams for example embodiments of a high speed interface (HSI) 141 between a channel controller 143 and a bank arbiter 145 as shown in Figure 1B. In at least one embodiment, the channel controller 143 can be integrated with the host 110 (as represented by the dashed lines). Alternatively, the channel controller 143 may be part of a module 111 connected to multiple memory devices 120-1, . . . , 120-N as shown in Figure 1B. As the reader will appreciate from reading this disclosure, the high speed interface (HSI) 141 can provide a status channel interface, or out-of-band (OOB) bus 155 between a processor in memory (PIM) device, e.g., memory device 120, and the channel controller 143, separate from a data bus 156 and control bus 154 between the memory device 120 and channel controller 143 and host 110 as shown in Figures 1A and 1B.

[0055] In one example, the HSI 141 is a first in first out (FIFO) style interface to the bank arbiter 145 of a processor in memory (PIM), e.g., memory device 120. The bank arbiter side can be a read FIFO interface with an empty bit, as well as clock and read signals, and a data out bus representing the bank, e.g., 121-1, . . . , 121-M, and row address, as well as the operation type, activate, pre-charge, and/or auto refresh for the PIM 120. Control logic and sequencers as part of memory controller 140 of each bank 121-1, . . . , 121-M (e.g., Bank 0, . . . , Bank 7) can each have a FIFO coming into the bank arbiter 145 for pushing status information. The bank arbiter 145 pulls status information off each of these FIFOs (e.g., 16 FIFOs for 2 each by 8 banks) and pushes the status information back to the host 110 via the HSI 141 and channel controller 143. The FIFOs may be checked in round robin fashion using a counter.

[0056] The HSI or status channel interface, e.g., OOB, 141 can be a dedicated interface between the memory device 120 and the channel controller 143. By way of example, and not by way of limitation, this interface may be

composed of sixteen (16) data outputs, 4 data strobes, and a rank enable input. In this example, when the memory device 120 is configured as a x64 DQ part, the status output pins are bused to each memory device 120-1, . . . , 120-M on a given channel of the channel controller 143. For example, the channel controller 143 may be equipped with four (4) rank enable outputs per channel. Rank Enable zero (0) may be connected to all memory devices 120-1, . . . , 120-M on the first rank, Rank Enable one (1) can be connected to all memory devices on the second rank, and so on as shown in the various example embodiments in Figures 3-6.

[0057] Figure 3 illustrates an example of an OOB bus (HSI) 341 between channel controller 343 and the bank arbiter 345 of a memory device 320 in the x64DQ one (1) rank configuration. As shown in the example of Figure 3, one rank enable bit (pin) 347 is provided from the channel controller 343 to the memory device 320. In this example, 16 SQ (status outputs/pins) 349 are received from the memory device 320 for the channel and 4 SQS (status output strobes) 351 may be received.

[0058] Figure 4 illustrates an example of an OOB bus (HSI) 441 between channel controller 443 and the bank arbiter 445 of a memory device 420 in the x64DQ four (4) rank configuration. In this example, each of four (4) separate rank enable bits 447 are shown (0, 1, 2, and 3) on the channel controller 443 independently connected for one channel to four memory devices 420-1, 420-2, . . . , 420-4. In this example, 16 SQ (status outputs) 449 are still received, but now from the four memory devices 420-1, 420-2, . . . , 420-4 for the channel. Likewise, four (4) SQS (status output strobes) 451 are received on the channel for the four memory devices 420-1, 420-2, . . . , 420-4.

[0059] In this case, the bandwidth of the status channel is divided just as the bandwidth of a double data rate (DDR) interface is divided, being effectively time domain multiplexed to avoid contention on the OOB 441. Each memory device 420, e.g., 420-1, 420-2, . . . , 420-4, is granted access to the uni-directional status bus 441 in turn as its rank enable input is activated. When x64 mode is selected (fuse option or test option), each memory device 420 may transmit one data packet at an active edge of the rank enable input. Status output bits 449 can be streamed out, 1 bit per clock cycle, synchronized to the SQS strobe pins 451. The first bit may be a preamble bit. In one example, if the preamble bit is a logic

0, the memory device, e.g., 420-1, 420-2, . . . , 420-4, will transmit a 128 bit status message on the 16 DQs on eight (8) successive clocks. At the end of this message the SQ and SQS pins will transition to a tri-state mode, and the interface circuitry will wait for the next active edge on the rank enable input to transmit additional messages. If there are no messages waiting in a memory controller's (e.g., memory controller 440 of a given memory device 420) output queue, the sense of the preamble bit will be logic 1. The SQ and SQS drivers will be turned off on the next successive clock. The memory device will go back to waiting for another active edge on the rank enable input.

[0060] Figure 5 illustrates an example of an OOB bus (HSI) 541 between channel controller 543 and the bank arbiter 545 of a memory device 520 in an x32DQ one (1) rank configuration. In the case of a x32 DQ configuration on the memory 520, the status channel can be split, as shown in Figure 5, just as the data channel can be split, with half of the status outputs 540 (as well as half of the DQs and DQSs) going to each of the memory devices 520 in the architecture.

[0061] In this example, the rank enable inputs 547 for the first rank are all tied together. Thus they will all begin transmitting their status messages on the lower eight (8) SQ pins 549. The protocol is largely the same as in the x64 case, with data being sent following a preamble on successive active clock edges of the SQSs. One deviation from the prior example is that since the data bus is half as wide for each memory device 520, the message length will increase from eight (8) data clocks to sixteen (16), to accommodate the same 128 bit message size.

[0062] Figure 6 illustrates an example of an OOB bus (HSI) 641 between channel controller 643 and the bank arbiter 645 of a memory device 620 in an x16 DQ four (4) rank configuration. In the case of a x16 DQ configuration on the memory 620, the status channel can be divided, as shown in Figure 6, into four (4) parts with each memory device, e.g., 620-1, 620-5, 620-9 and 620-13, connected to the channel controller 643 via the lowest four (4) bits of a given memory device's SQ bus 649. In the example, the number of clock cycles per message again grows, now to thirty-two (32) to accommodate the same 128 bit message size.

[0063] Figure 7 is a schematic diagram illustrating sensing circuitry 750 in accordance with a number of embodiments of the present disclosure. The

sensing circuitry 750 can correspond to sensing circuitry 150 shown in Figures 1A and 1B. The sense amplifier 706 of sensing circuitry 750 can correspond to sense amplifiers 206 shown in Figure 2, and the compute component 731 of sensing circuitry 750 can correspond to compute component 231 shown in Figure 2, for example.

[0064] A memory cell comprises a storage element (e.g., capacitor) and an access device (e.g., transistor). For instance, a first memory cell comprises transistor 702-1 and capacitor 703-1, and a second memory cell comprises transistor 702-2 and capacitor 703-2, etc. In this example, the memory array 730 is a DRAM array of 1T1C (one transistor one capacitor) memory cells. In a number of embodiments, the memory cells may be destructive read memory cells (e.g., reading the data stored in the cell destroys the data such that the data originally stored in the cell is refreshed after being read).

[0065] The cells of the memory array 730 can be arranged in rows coupled by word lines 704-X (Row X), 704-Y (Row Y), etc., and columns coupled by pairs of complementary sense lines (e.g., data lines DIGIT(n-1)/DIGIT(n-1)_, DIGIT(n)/DIGIT(n)_, DIGIT(n+1)/DIGIT(n+1)_). The individual sense lines corresponding to each pair of complementary sense lines can also be referred to as data lines 705-1 (D) and 705-2 (D_) respectively. Although only one pair of complementary data lines are shown in Figure 7, embodiments of the present disclosure are not so limited, and an array of memory cells can include additional columns of memory cells and/or data lines (e.g., 4,096, 8,192, 16,384, etc.).

[0066] Memory cells can be coupled to different data lines and/or word lines. For example, a first source/drain region of a transistor 702-1 can be coupled to data line 705-1 (D), a second source/drain region of transistor 702-1 can be coupled to capacitor 703-1, and a gate of a transistor 702-1 can be coupled to word line 704-X. A first source/drain region of a transistor 702-2 can be coupled to data line 705-2 (D_), a second source/drain region of transistor 702-2 can be coupled to capacitor 703-2, and a gate of a transistor 702-2 can be coupled to word line 704-Y. The cell plate, as shown in Figure 7, can be coupled to each of capacitors 703-1 and 703-2. The cell plate can be a common node to which a reference voltage (e.g., ground) can be applied in various memory array configurations.

[0067] The memory array 730 is coupled to sensing circuitry 750 in accordance with a number of embodiments of the present disclosure. In this example, the sensing circuitry 750 comprises a sense amplifier 706 and a compute component 731 corresponding to respective columns of memory cells (e.g., coupled to respective pairs of complementary data lines). The sense amplifier 706 can be coupled to the pair of complementary sense lines 705-1 and 705-2. The compute component 731 can be coupled to the sense amplifier 706 via pass gates 707-1 and 707-2. The gates of the pass gates 707-1 and 707-2 can be coupled to logical operation selection logic 713.

[0068] The logical operation selection logic 713 can be configured to include pass gate logic for controlling pass gates that couple the pair of complementary sense lines un-transposed between the sense amplifier 706 and the compute component 731 (as shown in Figure 7) and/or swap gate logic for controlling swap gates that couple the pair of complementary sense lines transposed between the sense amplifier 706 and the compute component 731 (as is discussed later with respect to Figure 6, for example). The logical operation selection logic 713 can also be coupled to the pair of complementary sense lines 705-1 and 705-2. The logical operation selection logic 713 can be configured to control continuity of pass gates 707-1 and 707-2 based on a selected logical operation, as described in detail below for various configurations of the logical operation selection logic 713.

[0069] The sense amplifier 706 can be operated to determine a data value (e.g., logic state) stored in a selected memory cell. The sense amplifier 706 can comprise a cross coupled latch, which can be referred to herein as a primary latch. In the example illustrated in Figure 7, the circuitry corresponding to sense amplifier 706 comprises a latch 715 including four transistors coupled to a pair of complementary data lines D 705-1 and D_ 705-2. However, embodiments are not limited to this example. The latch 715 can be a cross coupled latch (e.g., gates of a pair of transistors, such as n-channel transistors (e.g., NMOS transistors) 727-1 and 727-2 are cross coupled with the gates of another pair of transistors, such as p-channel transistors (e.g., PMOS transistors) 729-1 and 729-2). The cross coupled latch 715 comprising transistors 727-1, 727-2, 729-1, and 729-2 can be referred to as a primary latch.

[0070] In operation, when a memory cell is being sensed (e.g., read), the voltage on one of the data lines 705-1 (D) or 705-2 (D_̅) will be slightly greater than the voltage on the other one of data lines 705-1 (D) or 705-2 (D_̅). An ACT signal and the RNL* signal can be driven low to enable (e.g., fire) the sense amplifier 706. The data lines 705-1 (D) or 705-2 (D_̅) having the lower voltage will turn on one of the PMOS transistor 729-1 or 729-2 to a greater extent than the other of PMOS transistor 729-1 or 729-2, thereby driving high the data line 705-1 (D) or 705-2 (D_̅) having the higher voltage to a greater extent than the other data line 705-1 (D) or 705-2 (D_̅) is driven high.

[0071] Similarly, the data line 705-1 (D) or 705-2 (D_̅) having the higher voltage will turn on one of the NMOS transistor 727-1 or 727-2 to a greater extent than the other of the NMOS transistor 727-1 or 727-2, thereby driving low the data line 705-1 (D) or 705-2 (D_̅) having the lower voltage to a greater extent than the other data line 705-1 (D) or 705-2 (D_̅) is driven low. As a result, after a short delay, the data line 705-1 (D) or 705-2 (D_̅) having the slightly greater voltage is driven to the voltage of the supply voltage V_{CC} through source transistor 711, and the other data line 705-1 (D) or 705-2 (D_̅) is driven to the voltage of the reference voltage (e.g., ground) through the sink transistor 713. Therefore, the cross coupled NMOS transistors 727-1 and 727-2 and PMOS transistors 729-1 and 729-2 serve as a sense amplifier pair, which amplify the differential voltage on the data lines 705-1 (D) and 705-2 (D_̅) and operate to latch a data value sensed from the selected memory cell. As used herein, the cross coupled latch of sense amplifier 706 may be referred to as a primary latch 715.

[0072] Embodiments are not limited to the sense amplifier 706 configuration illustrated in Figure 7. As an example, the sense amplifier 706 can be current-mode sense amplifier and/or single-ended sense amplifier (e.g., sense amplifier coupled to one data line). Also, embodiments of the present disclosure are not limited to a folded data line architecture such as that shown in Figure 7.

[0073] The sense amplifier 706 can, in conjunction with the compute component 731, be operated to perform various logical operations using data from an array as input. In a number of embodiments, the result of a logical operation can be stored back to the array without transferring the data via a data line address access (e.g., without firing a column decode signal such that data is

transferred to circuitry external from the array and sensing circuitry via local I/O lines). As such, a number of embodiments of the present disclosure can enable performing logical operations and compute functions associated therewith using less power than various previous approaches. Additionally, since a number of embodiments eliminate the need to transfer data across I/O lines in order to perform compute functions (e.g., between memory and discrete processor), a number of embodiments can enable an increased parallel processing capability as compared to previous approaches.

[0074] The sense amplifier 706 can further include equilibration circuitry 714, which can be configured to equilibrate the data lines 705-1 (D) and 705-2 (D_̄). In this example, the equilibration circuitry 714 comprises a transistor 724 coupled between data lines 705-1 (D) and 705-2 (D_̄). The equilibration circuitry 714 also comprises transistors 725-1 and 725-2 each having a first source/drain region coupled to an equilibration voltage (e.g., $V_{DD}/2$), where V_{DD} is a supply voltage associated with the array. A second source/drain region of transistor 725-1 can be coupled data line 705-1 (D), and a second source/drain region of transistor 725-2 can be coupled data line 705-2 (D_̄). Gates of transistors 724, 725-1, and 725-2 can be coupled together, and to an equilibration (EQ) control signal line 726. As such, activating EQ enables the transistors 724, 725-1, and 725-2, which effectively shorts data lines 705-1 (D) and 705-2 (D_̄) together and to the an equilibration voltage (e.g., $V_{CC}/2$).

[0075] Although Figure 7 shows sense amplifier 706 comprising the equilibration circuitry 714, embodiments are not so limited, and the equilibration circuitry 714 may be implemented discretely from the sense amplifier 706, implemented in a different configuration than that shown in Figure 7, or not implemented at all.

[0076] As described further below, in a number of embodiments, the sensing circuitry (e.g., sense amplifier 706 and compute component 731) can be operated to perform a selected logical operation and initially store the result in one of the sense amplifier 706 or the compute component 731 without transferring data from the sensing circuitry via an I/O line (e.g., without performing a data line address access via activation of a column decode signal, for instance).

[0077] Performance of logical operations (e.g., Boolean logical functions involving data values) is fundamental and commonly used. Boolean logic functions are used in many higher level functions. Consequently, speed and/or power efficiencies that can be realized with improved logical operations, which can translate into speed and/or power efficiencies of higher order functionalities. Described herein are apparatuses and methods for performing logical operations without transferring data via an input/output (I/O) line and/or without transferring data to a control component external to the array. Depending on memory array architecture, the apparatuses and methods for performing the logical operations may not require amplification of a sense line (e.g., data line, digit line, bit line) pair.

[0078] As shown in Figure 7, the compute component 731 can also comprise a latch, which can be referred to herein as a secondary latch 764. The secondary latch 764 can be configured and operated in a manner similar to that described above with respect to the primary latch 715, with the exception that the pair of cross coupled p-channel transistors (e.g., PMOS transistors) comprising the secondary latch can have their respective sources coupled to a supply voltage (e.g., V_{DD}), and the pair of cross coupled n-channel transistors (e.g., NMOS transistors) of the secondary latch can have their respective sources selectively coupled to a reference voltage (e.g., ground), such that the secondary latch is continuously enabled. The configuration of the compute component is not limited to that shown in Figure 7 at 731, and various other embodiments are described further below.

[0079] Figure 8 is a schematic diagram illustrating sensing circuitry capable of implementing an XOR logical operation in accordance with a number of embodiments of the present disclosure. Figure 8 shows a sense amplifier 806 coupled to a pair of complementary sense lines 805-1 and 805-2, and a compute component 831 coupled to the sense amplifier 806 via pass gates 807-1 and 807-2. The sense amplifier 806 shown in Figure 8 can correspond to sense amplifier 706 shown in Figure 7. The compute component 831 shown in Figure 8 can correspond to compute component 231 shown in Figure 2, for example. The logical operation selection logic 813 shown in Figure 8 can correspond to logical operation selection logic 713 shown in Figure 7, for example.

[0080] The gates of the pass gates 807-1 and 807-2 can be controlled by a logical operation selection logic signal, Pass. For example, an output of the logical operation selection logic can be coupled to the gates of the pass gates 807-1 and 807-2. The compute component 831 can comprise a loadable shift register configured to shift data values left and right.

[0081] According to the embodiment illustrated in Figure 8, the compute components 831 can comprise respective stages (e.g., shift cells) of a loadable shift register configured to shift data values left and right. For example, as illustrated in Figure 8, each compute component 831 (e.g., stage) of the shift register comprises a pair of right-shift transistors 881 and 886, a pair of left-shift transistors 889 and 890, and a pair of inverters 887 and 888. The signals PHASE 1R, PHASE 2R, PHASE 1L, and PHASE 2L can be applied to respective control lines 882, 883, 891 and 892 to enable/disable feedback on the latches of the corresponding compute components 831 in association with performing logical operations and/or shifting data in accordance with embodiments described herein.

[0082] The sensing circuitry shown in Figure 8 also shows a logical operation selection logic 813 coupled to a number of logic selection control input control lines, including ISO, TF, TT, FT, and FF. Selection of a logical operation from a plurality of logical operations is determined from the condition of logic selection control signals on the logic selection control input control lines, as well as the data values present on the pair of complementary sense lines 805-1 and 805-2 when the isolation transistors are enabled via the ISO control signal being asserted.

[0083] According to various embodiments, the logical operation selection logic 813 can include four logic selection transistors: logic selection transistor 862 coupled between the gates of the swap transistors 842 and a TF signal control line, logic selection transistor 852 coupled between the gates of the pass gates 807-1 and 807-2 and a TT signal control line, logic selection transistor 854 coupled between the gates of the pass gates 807-1 and 807-2 and a FT signal control line, and logic selection transistor 864 coupled between the gates of the swap transistors 842 and a FF signal control line. Gates of logic selection transistors 862 and 852 are coupled to the true sense line through isolation transistor 850-1 (having a gate coupled to an ISO signal control line).

Gates of logic selection transistors 864 and 854 are coupled to the complementary sense line through isolation transistor 850-2 (also having a gate coupled to an ISO signal control line).

[0084] Data values present on the pair of complementary sense lines 805-1 and 805-2 can be loaded into the compute component 831 via the pass gates 807-1 and 807-2. The compute component 831 can comprise a loadable shift register. When the pass gates 807-1 and 807-2 are OPEN, data values on the pair of complementary sense lines 805-1 and 805-2 are passed to the compute component 831 and thereby loaded into the loadable shift register. The data values on the pair of complementary sense lines 805-1 and 805-2 can be the data value stored in the sense amplifier 806 when the sense amplifier is fired. The logical operation selection logic signal, Pass, is high to OPEN the pass gates 807-1 and 807-2.

[0085] The ISO, TF, TT, FT, and FF control signals can operate to select a logical function to implement based on the data value (“B”) in the sense amplifier 806 and the data value (“A”) in the compute component 831. In particular, the ISO, TF, TT, FT, and FF control signals are configured to select the logical function to implement independent from the data value present on the pair of complementary sense lines 805-1 and 805-2 (although the result of the implemented logical operation can be dependent on the data value present on the pair of complementary sense lines 805-1 and 805-2. That is, the ISO, TF, TT, FT, and FF control signals select the logical operation to implement directly since the data value present on the pair of complementary sense lines 805-1 and 805-2 is not passed through logic to operate the gates of the pass gates 807-1 and 807-2.

[0086] Additionally, Figure 8 shows swap transistors 842 configured to swap the orientation of the pair of complementary sense lines 805-1 and 805-2 between the sense amplifier 813-7 and the compute component 831. When the swap transistors 842 are OPEN, data values on the pair of complementary sense lines 805-1 and 805-2 on the sense amplifier 806 side of the swap transistors 842 are oppositely-coupled to the pair of complementary sense lines 805-1 and 805-2 on the compute component 831 side of the swap transistors 842, and thereby loaded into the loadable shift register of the compute component 831.

[0087] The logical operation selection logic signal Pass can be activated (e.g., high) to OPEN the pass gates 807-1 and 807-2 (e.g., conducting) when the ISO control signal line is activated and either the TT control signal is activated (e.g., high) with data value on the true sense line is “1” or the FT control signal is activated (e.g., high) with the data value on the complement sense line is “1.”

[0088] The data value on the true sense line being a “1” OPENS logic selection transistors 852 and 862. The data value on the complimentary sense line being a “1” OPENS logic selection transistors 854 and 864. If the ISO control signal or either the respective TT/FT control signal or the data value on the corresponding sense line (e.g., sense line to which the gate of the particular logic selection transistor is coupled) is not high, then the pass gates 807-1 and 807-2 will not be OPENed by a particular logic selection transistor.

[0089] The logical operation selection logic signal PassF can be activated (e.g., high) to OPEN the swap transistors 842 (e.g., conducting) when the ISO control signal line is activated and either the TF control signal is activated (e.g., high) with data value on the true sense line is “1,” or the FF control signal is activated (e.g., high) with the data value on the complement sense line is “1.” If either the respective control signal or the data value on the corresponding sense line (e.g., sense line to which the gate of the particular logic selection transistor is coupled) is not high, then the swap transistors 842 will not be OPENed by a particular logic selection transistor.

[0090] The Pass* control signal is not necessarily complementary to the Pass control signal. It is possible for the Pass and Pass* control signals to both be activated or both be deactivated at the same time. However, activation of both the Pass and Pass* control signals at the same time shorts the pair of complementary sense lines together, which may be a disruptive configuration to be avoided.

[0091] The sensing circuitry illustrated in Figure 8 is configured to select one of a plurality of logical operations to implement directly from the four logic selection control signals (e.g., logical operation selection is not dependent on the data value present on the pair of complementary sense lines). Some combinations of the logic selection control signals can cause both the pass gates 807-1 and 807-2 and swap transistors 842 to be OPEN at the same time, which shorts the pair of complementary sense lines 805-1 and 805-2 together.

According to a number of embodiments of the present disclosure, the logical operations which can be implemented by the sensing circuitry illustrated in Figure 8 can be the logical operations summarized in the logic tables shown in Figure 9.

[0092] Figure 9 is a logic table illustrating selectable logic operation results implemented by a sensing circuitry shown in Figure 8 in accordance with a number of embodiments of the present disclosure. The four logic selection control signals (e.g., TF, TT, FT, and FF), in conjunction with a particular data value present on the complementary sense lines, can be used to select one of plural logical operations to implement involving the starting data values stored in the sense amplifier 806 and compute component 831. The four control signals, in conjunction with a particular data value present on the complementary sense lines, controls the continuity of the pass gates 807-1 and 807-2 and swap transistors 842, which in turn affects the data value in the compute component 831 and/or sense amplifier 806 before/after firing. The capability to selectably control continuity of the swap transistors 842 facilitates implementing logical operations involving inverse data values (e.g., inverse operands and/or inverse result), among others.

[0093] Logic Table 9-1 illustrated in Figure 9 shows the starting data value stored in the compute component 631 shown in column A at 944, and the starting data value stored in the sense amplifier 606 shown in column B at 945. The other 3 column headings in Logic Table 9-1 refer to the continuity of the pass gates 807-1 and 807-2, and the swap transistors 842, which can respectively be controlled to be OPEN or CLOSED depending on the state of the four logic selection control signals (e.g., TF, TT, FT, and FF), in conjunction with a particular data value present on the pair of complementary sense lines 805-1 and 805-2. The “Not Open” column corresponds to the pass gates 807-1 and 807-2 and the swap transistors 842 both being in a non-conducting condition, the “Open True” corresponds to the pass gates 807-1 and 807-2 being in a conducting condition, and the “Open Invert” corresponds to the swap transistors 842 being in a conducting condition. The configuration corresponding to the pass gates 807-1 and 807-2 and the swap transistors 842 both being in a conducting condition is not reflected in Logic Table 8-1 since this results in the sense lines being shorted together.

[0094] Via selective control of the continuity of the pass gates 807-1 and 807-2 and the swap transistors 842, each of the three columns of the upper portion of Logic Table 9-1 can be combined with each of the three columns of the lower portion of Logic Table 9-1 to provide $3 \times 3 = 9$ different result combinations, corresponding to nine different logical operations, as indicated by the various connecting paths shown at 975. The nine different selectable logical operations that can be implemented by the sensing circuitry 850 are summarized in Logic Table 9-2 illustrated in Figure 9, including an XOR logical operation.

[0095] The columns of Logic Table 9-2 illustrated in Figure 9 show a heading 980 that includes the state of logic selection control signals. For example, the state of a first logic selection control signal is provided in row 976, the state of a second logic selection control signal is provided in row 977, the state of a third logic selection control signal is provided in row 978, and the state of a fourth logic selection control signal is provided in row 979. The particular logical operation corresponding to the results is summarized in row 947.

[0096] While example embodiments including various combinations and configurations of sensing circuitry, sense amplifiers, compute component, dynamic latches, isolation devices, and/or shift circuitry have been illustrated and described herein, embodiments of the present disclosure are not limited to those combinations explicitly recited herein. Other combinations and configurations of the sensing circuitry, sense amplifiers, compute component, dynamic latches, isolation devices, and/or shift circuitry disclosed herein are expressly included within the scope of this disclosure.

[0097] Although specific embodiments have been illustrated and described herein, those of ordinary skill in the art will appreciate that an arrangement calculated to achieve the same results can be substituted for the specific embodiments shown. This disclosure is intended to cover adaptations or variations of one or more embodiments of the present disclosure. It is to be understood that the above description has been made in an illustrative fashion, and not a restrictive one. Combination of the above embodiments, and other embodiments not specifically described herein will be apparent to those of skill in the art upon reviewing the above description. The scope of the one or more embodiments of the present disclosure includes other applications in which the above structures and methods are used. Therefore, the scope of one or more

embodiments of the present disclosure should be determined with reference to the appended claims, along with the full range of equivalents to which such claims are entitled.

[0098] In the foregoing Detailed Description, some features are grouped together in a single embodiment for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the disclosed embodiments of the present disclosure have to use more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment. Thus, the following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separate embodiment.

What is claimed is:

1. An apparatus, comprising:
 - a memory device, wherein the memory device includes:
 - an array of memory cells;
 - sensing circuitry coupled to the array, the sensing circuitry including a sense amplifier and a compute component configured to implement logical operations; and
 - a controller coupled to the array and sensing circuitry, the controller configured to:
 - receive a block of instructions, the block of instructions including a set of instructions for at least one of a gather operation and a scatter operation; and
 - subdivide the block of instructions so that the at least one of the gather operation and the scatter operation gather operation can be executed locally on the memory device based on the subdivided block of instructions.
2. The apparatus of claim 1, wherein the array of memory cells are dynamic random access memory (DRAM) cells.
3. An apparatus, comprising:
 - a memory device, wherein the memory device includes:
 - an array of memory cells;
 - sensing circuitry coupled to the array, the sensing circuitry including a sense amplifier and a compute component configured to implement logical operations; and
 - a memory controller coupled to the array and sensing circuitry;
 - and
 - a channel controller configured to receive a block of instructions, the block instructions including individual instructions for at least one of a gather operation and a scatter operation and to control the memory controller so that the at least one of the gather operation and the scatter operation is executed

on the memory device based on a corresponding one of the individual instructions.

4. The apparatus of claim 3, wherein:

the array of memory cells includes a plurality of banks of memory cells, wherein each bank includes a respective memory controller, array, and sensing circuitry;

wherein the channel controller is coupled to each of the respective memory controllers; and

wherein the memory controller is configured to:

subdivide the block of instructions to manage execution of local individual instructions on banks in the memory device;

monitor completion of the gather operation to ensure in-order execution of the gather operation and to ensure that other memory device operations are complete before starting the gather operation among the plurality of banks of memory cells.

5. The apparatus of claims 3-4, wherein the array of memory cells includes a plurality of banks of memory cells, and wherein the channel controller is configured to:

change vector references in the received block of instructions by reading an index (idx) vector array block by block, the idx vector array stored in the memory device;

store a buffer of idx data locally among the plurality of banks; and

group requests together by locality among the plurality of banks so that array and sensing circuitry execution elements can gather and return requests to the channel controller.

6. The apparatus of claim 5, wherein the channel controller is configured to:

stream a collection of data from the plurality of banks by iterating through the idx array; and

allocate local buffers to hold:

a section of the idx array;

intermediate results from x vector reads; and

merged results for a y-target vector.

7. The apparatus of claim 5, wherein the apparatus includes a plurality of memory devices coupled to the channel controller, wherein the channel controller is configured to:

read a buffer of indices in the idx vector array stored to a local buffer;

store the buffer of indices to a local idx buffer;

bin independent bank command buffers for the plurality of banks; and

create individual instructions to communicate to respective ones of the plurality of memory devices.

8. The apparatus of claim 7, wherein the channel controller communicates with the plurality of memory devices via an out-of-band bus.

9. The apparatus of claim 5, wherein the channel controller is configured to: perform a scatter operation by:

reading locations to be written to from the idx vector and data to be written from x vectors into a target y vector; and

grouping write requests by locality.

10. The apparatus of claim 9, wherein the channel controller is configured to: receive a stream of instructions at the memory device to be executed, wherein the stream of instructions contains embedded data and specific locations in the plurality of banks;

send the stream of instructions to the plurality of banks to be executed by the plurality of banks; and

only receive notification, using an out-of-band bus configured to provide status information of the memory device, when writes are completed.

11. An apparatus, comprising:

a channel controller; and

a memory device coupled to the channel controller, wherein the memory device includes:

an array of memory cells;

sensing circuitry coupled to the array, the sensing circuitry including a sense amplifier and a compute component configured to implement logical operations; and

a memory controller coupled to the array and the sensing circuitry, wherein the memory controller is configured to communicate with the channel controller via an out-of-band bus.

12. The apparatus of claim 11, wherein the out-of-bound bus is configured to provide status information of the memory device.

13. The apparatus of claim 11, wherein the channel controller is configured to provide gather and/or scatter instruction commands to the memory controller.

14. The apparatus of claims 11-13, wherein the channel controller is configured to store an index vector (idx) for a gather and/or scatter operation in the array of memory cells to be operated upon by the memory controller.

15. The apparatus of claim 14, wherein the out-of-band (OOB) bus is associated with a status channel interface on the memory device and configured to provide status messages, including a write complete message, to the channel controller in sixteen (16) or fewer clock cycles.

16. The apparatus of claim 14, wherein the out-of-band bus is associated with a status channel interface on the memory device and is configured to provide a 128 bit status message to the channel controller in eight (8) clock cycles for a x64 DQ configuration.

17. The apparatus of claims 11-13, wherein a plurality of memory devices are coupled to the channel controller, and the channel controller is configured with at least four (4) rank enable outputs per channel.

18. The apparatus of claims 11-13, wherein the out-of-band (OOB) bus is associated with a status channel interface on the memory device and configured to provide sixteen (16) status (SQ) outputs per status output strobe.

19. The apparatus of claims 11-13, wherein the memory controller is configured to control the sensing circuitry to implement logical functions including AND, OR, NOT, XOR, NAND, and NOR logical functions.
20. The apparatus of claim 19, wherein the memory controller is configured to control the sensing circuitry to perform non-Boolean logic operations including copy, compare and erase.
21. The apparatus of claims 11-13, wherein the out-of-band (OOB) bus is a 16 bit double data rate 4 (DDR4) type high speed interface to the memory device.
22. An apparatus, comprising:
a channel controller; and
a memory device coupled to the channel controller, wherein the memory device includes:
an array of memory cells;
sensing circuitry coupled to the array, the sensing circuitry including a sense amplifier and a compute component configured to implement logical operations; and
a memory controller coupled to the array and the sensing circuitry; and
wherein the channel controller is configured to:
receive a block of instructions to write to the array, the block instructions including individual instructions for at least one of a gather operation and a scatter operation;
send individual instructions to the memory device;
control the memory controller such that at least one of the gather operation and the scatter operation is execute on the memory device based on a corresponding one of the individual instructions; and
communicate with the memory device via an out-of-band bus.

23. The apparatus of claim 22, wherein the out-of-band bus is configured to provide a word of information to the channel controller that includes a status update.
24. The apparatus of claim 22, wherein the memory controller is configured to execute the block instruction to cause the sensing circuitry to perform the gather operation locally on a plurality of banks on the memory device.
25. The apparatus of claims 22-24, wherein the channel controller is configured to:
- send block instructions for the gather operation to a plurality of memory devices to be performed locally on each respective memory device; and
 - receive return requested data via the out-of-band (OOB) bus to merge collected results.
26. The apparatus of claim 25, wherein the data bus and the out-of-band (OOB) bus are configured to be a same type interface, the data bus being configured to have a greater bit-width than the out-of-band (OOB) bus, and wherein the data bus is a 64-bit DDR4 type PHY interface.
27. The apparatus of claims 22-24, wherein the out-of-band (OOB) bus is configured to have a bandwidth of at least 17.07 Gbps.
28. The apparatus of claims 22-24, wherein the out-of-band (OOB) bus is configured to:
- use a data strobe signal that is distinct from a data strobe signal used by the data bus;
 - use a same clock signal as the data bus; and
 - use a same clock enable signal as the data bus.
29. A method for operating a memory device:
- receiving individual instructions associated with at least one of a gather operation and a scatter operation at respective banks in a plurality of banks on the memory device;

controlling execution of the individual instructions associated with the at least one of the gather operation and the scatter operation on the memory device by a channel controller;

executing the individual instructions using a memory controller on the memory device to perform the at least one of the gather operation and the scatter operation on the respective banks; and

returning requested data results to the channel controller via an out-of-band (OOB) bus.

30. The method of claim 29, wherein the method comprises:

receiving to a channel controller connected to a plurality of memory devices a gather operation index (idx) vector;

binning indices from the idx vector to a plurality of bank command buffers local to the channel controller for each of the plurality of memory devices;

sending individual instructions to the plurality of banks to perform the gather operation.

31. The method of claim 30, wherein the method comprises:

receiving via the OOB bus requested data from the gather operation from the plurality of memory devices to the channel controller; and

merging the received requested data into a local buffer associated with the channel controller.

32. The method of claim 31, wherein the method further comprises operating the plurality of memory devices to perform a scatter operation, wherein the method comprises:

reading locations to be written to from command instructions received from the idx vector and data to be written from an x-source vector using sensing circuitry associated with an array of memory cells on each respective bank of the plurality of memory devices to perform the scatter operation local to the bank.

33. The method of claim 32, wherein the method comprises:

using the channel controller to group write requests by locality for respective banks among the plurality of memory devices;

providing individual instructions associated with the scatter operation to respective banks of the plurality of memory devices such that the individual instructions can be executed locally by respective banks using a memory controller of the bank; and

communicating completion of write requests associated with the scatter operation to the channel controller from respective banks using the OOB bus.




Fig. 1A




Fig. 1B




Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7

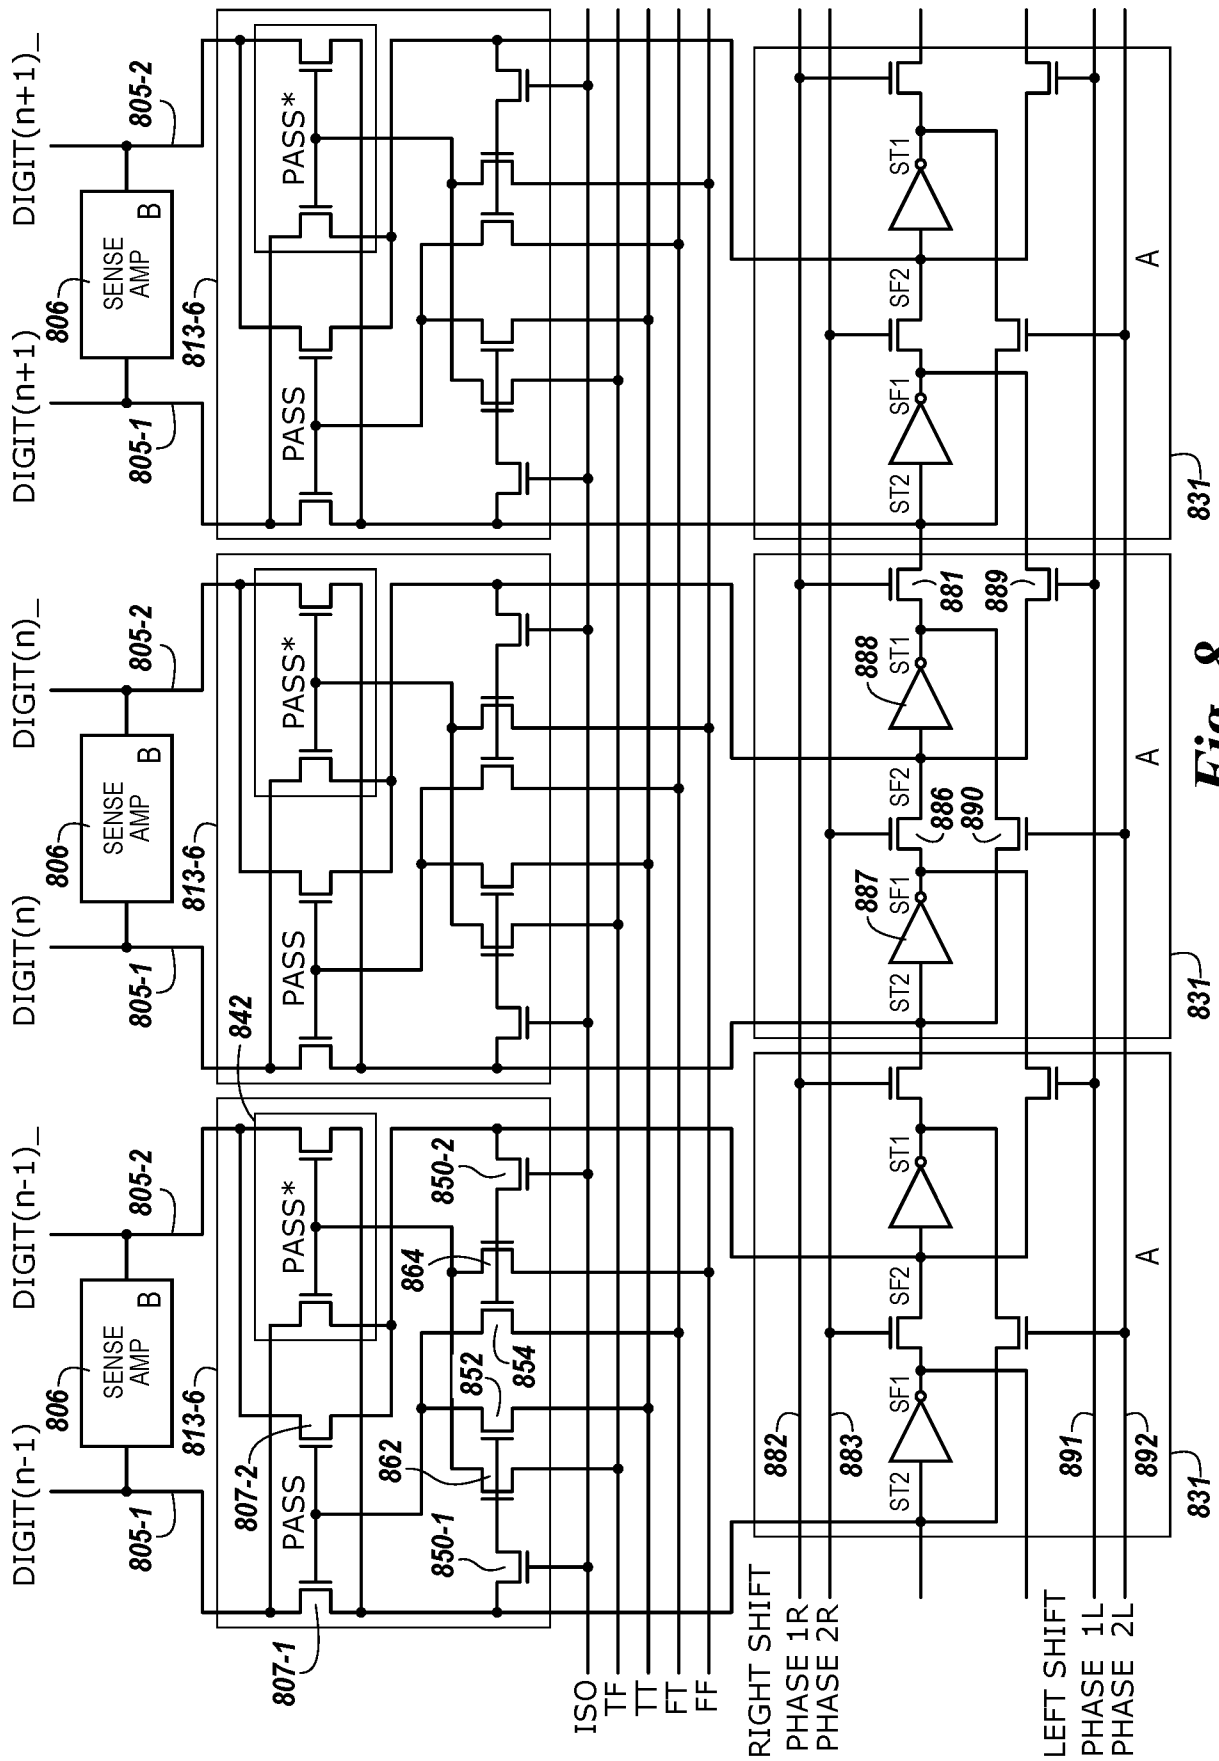


Fig. 8




Fig. 9

A. CLASSIFICATION OF SUBJECT MATTER**G11C 7/06(2006.01)i, G11C 5/14(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G11C 7/06; G11C 7/10; G11C 5/06; G06F 12/00; H03K 19/177; G11C 5/02; G11C 5/14Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: memory, cell, array, logic, sensing, circuit, controller, gather, scatter, instruction, operation, out-of-band, DRAM**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014-0082306 A1 (KAZUHIKO KAJIGAYA) 20 March 2014 See paragraphs [0063], [0065]-[0066], [0079]-[0080]; and figures 5-6, 12.	1-3
A		4-33
A	US 05638317 A (HIEP V. TRAN) 10 June 1997 See column 5, lines 30-35, 38-40; claim 9; and figures 3b-3c.	1-33
A	US 2012-0210052 A1 (ROBERT NORMAN) 16 August 2012 See paragraphs [0038]-[0039]; and figure 2.	1-33
A	US 2014-0176187 A1 (NUWAN S. JAYASENA et al.) 26 June 2014 See paragraphs [0021]-[0034]; and figure 1.	1-33
A	US 2004-0236920 A1 (GAD S. SHEAFFER) 25 November 2004 See paragraphs [0017]-[0018]; and figure 2.	1-33

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

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

19 May 2016 (19.05.2016)

Date of mailing of the international search report

31 May 2016 (31.05.2016)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

BYUN, Sung Cheal

Telephone No. +82-42-481-8262



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/015027

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0082306 A1	20/03/2014	JP 2011-165298 A US 2011-0179239 A1 US 8621135 B2 US 8966153 B2	25/08/2011 21/07/2011 31/12/2013 24/02/2015
US 05638317 A	10/06/1997	JP 03361825 B2 JP 06-089575 A	07/01/2003 29/03/1994
US 2012-0210052 A1	16/08/2012	US 2010-161918 A1 US 2010-162065 A1 US 8164970 B2	24/06/2010 24/06/2010 24/04/2012
US 2014-0176187 A1	26/06/2014	US 2015-155876 A1 US 8922243 B2	04/06/2015 30/12/2014
US 2004-0236920 A1	25/11/2004	None	