



(51) International Patent Classification:

G06F 3/06 (2006.01) *G06F 13/16* (2006.01)

(21) International Application Number:

PCT/US2019/051579

(22) International Filing Date:

17 September 2019 (17.09.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/132,875 17 September 2018 (17.09.2018) US

(71) Applicant: **MICRON TECHNOLOGY, INC.** [US/US];
8000 South Federal Way, Boise, Idaho 83716 (US).(72) Inventors: **WANG, Wei**; 3369 Monaghan Street, Dublin,
California 94568 (US). **ZHU, Jiangli**; 1843 Nelson Way,
San Jose, California 95124 (US). **TAI, Ying Yu**; 1439 Todd
Street, Mountain View, California 94040 (US). **MITTAL,**
Samir; 271 Ely Place, Palo Alto, California 94306 (US).(74) Agent: **SHEKHER, Rahul** et al.; One Lowenstein Drive,
Roseland, New Jersey 07068 (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,
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).

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))

(54) Title: SCHEDULING OF READ OPERATIONS AND WRITE OPERATIONS BASED ON A DATA BUS MODE




FIG. 1

(57) **Abstract:** A data bus can be determined to be in a write mode based on a prior operation transmitted over the data bus being a write operation. In response to determining that the data bus is in the write mode, a number of partition queues of a plurality of partition queues that include at least one write operation can be identified. A determination as to whether the number of partition queues of the plurality of partition queues satisfies a threshold number can be made. In response to determining that the number of partition queues satisfies the threshold number, another write operation from the plurality of partition queues can be transmitted over the data bus.

Published:

— *with international search report (Art. 21(3))*

SCHEDULING OF READ OPERATIONS AND WRITE OPERATIONS BASED ON A DATA BUS MODE

TECHNICAL FIELD

[001] The present disclosure generally relates to a memory sub-system, and more specifically, relates to the scheduling of read operations and write operations based on a data bus mode in memory sub-systems.

BACKGROUND

[002] A memory sub-system can be a storage system, such as a solid-state drive (SSD), or a hard disk drive (HDD). A memory sub-system can be a memory module, such as a dual in-line memory module (DIMM), a small outline DIMM (SO-DIMM), or a non-volatile dual in-line memory module (NVDIMM). A memory sub-system can include one or more memory components that store data. The memory components can be, for example, non-volatile memory components and volatile memory components. In general, a host system can utilize a memory sub-system to store data at the memory components and to retrieve data from the memory components.

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 implementations of the disclosure.

[004] **FIG. 1** illustrates an example computing environment that includes a memory sub-system in accordance with some embodiments of the present disclosure.

[005] **FIG. 2** illustrates an example scheduling of operations in accordance with some embodiments of the present disclosure.

[006] **FIG. 3** is a flow diagram of an example method to schedule the issuance of operations when a data bus is in a read mode in accordance with some embodiments.

[007] **FIG. 4** is a flow diagram of an example method to schedule the issuance of operations when a data bus is in a write mode in accordance with some embodiments.

[008] **FIG. 5** is a flow diagram of an example method to schedule read operations and write operations in accordance with some embodiments of the present disclosure.

[009] **FIG. 6** is a block diagram of an example computer system in which implementations of the present disclosure can operate.

DETAILED DESCRIPTION

[0010] Aspects of the present disclosure are directed to the scheduling of read operations and write operations based on a data bus mode. A memory sub-system is also hereinafter referred to as a “memory device.” An example of a memory sub-system is a storage device that is coupled to a central processing unit (CPU) via a peripheral interconnect (e.g., an input/output bus, a storage area network). Examples of storage devices include a solid-state drive (SSD), a flash drive, a universal serial bus (USB) flash drive, and a hard disk drive (HDD). Another example of a memory sub-system is a memory module that is coupled to the CPU via a memory bus. Examples of memory modules include a dual in-line memory module (DIMM), a small outline DIMM (SO-DIMM), a non-volatile dual in-line memory module (NVDIMM), etc. In some embodiments, the memory sub-system can be a hybrid memory/storage sub-system. In general, a host system can utilize a memory sub-system that includes one or more memory components. The host system can provide data to be stored at the memory sub-system and can request data to be retrieved from the memory sub-system.

[0011] In a conventional memory sub-system, read operations and write operations that are received from the host system can be placed into a command queue. The conventional memory sub-system can include a data bus (i.e., an interface) that is used to transmit the read operations and write operations to different partitions of the memory sub-system. A partition of the memory sub-system can be one or more memory components (i.e., memory dies). Thus, the data bus can couple the command queue with multiple partitions in order to transmit data to and from memory components of the partitions.

[0012] The conventional memory sub-system can transmit and receive data via the data bus by implementing a bus protocol based on requirements of the data bus and the memory components. For example, when a command is transmitted over the data bus to one of the memory components in one of the partitions, a particular amount of time should elapse before a subsequent operation (e.g., a next read operation or a write operation) can be transmitted to the same partition. Additionally, if a first portion of a memory component is included in a first partition and a second portion of the memory component is included in a second partition and if the first portion of the memory component included in the first partition receives an operation, then a particular amount of time should elapse before a subsequent operation (e.g., a next read operation or a next write operation) can be transmitted to the second portion of the memory component that is included in the second partition. Additionally, the utilization of the data bus can further be based on characteristics of the memory components. For example, for a conventional memory sub-system that utilizes cross point array based memory components,

different amounts of time can elapse between a transmitting of a read operation and a transmitting of a write operation (e.g., different read and write latencies). Other requirements of the bus protocol can include an amount of elapsed time (i.e., a turnaround) for the data bus when transmitting separate write operations or read operations and an amount of elapsed time between utilizing the data bus to perform read operations at different memory components (e.g., different memory die). As such, the implementation of the bus protocol to control the transmitting and receiving of data from read operations and write operations over the data bus can result in a particular amount of time where the data bus is not being used to transmit or receive data (e.g., the bus protocol enforces a particular time to elapse between operations). For example, if the conventional memory sub-system is to perform a read operation followed by a write operation (or vice versa), then the switch between performing the read operation and performing the write operation can result in an amount of time where the utilization of the data bus is reduced.

[0013] Advantages of the present disclosure address the above and other deficiencies by scheduling read operations and write operations based on a data bus mode to reduce the number of switches from read operations to write operations or from write operations to read operations. A data bus mode can specify a prior operation that was transmitted over the data bus. For example, the data bus mode can specify a read mode if the last operation that was transmitted over the data bus was a read operation or the data bus mode can specify a write mode if the last operation that was transmitted over the data bus was a write operation. The scheduling of read operations and write operations can be based on reducing the number of times that the data bus mode switches between the read mode and the write mode.

[0014] If the data bus is at a read mode, then a read operation from an available partition queue (i.e., the portion of the command queue corresponding to an available partition) that includes the most number of read operation for the current working memory component or memory die can be selected. The current working memory component can be the prior memory component (or memory die) that was subjected to the prior read operation. Furthermore, the partition queue can be considered to be available if an operation can be provided to the corresponding partition while satisfying requirements of the bus protocol. If a read operation in an available partition queue is identified for the current working memory component, then the read operation can be transmitted over the data bus. Otherwise, if such a read operation cannot be identified for any available partition queues associated with the current working memory component, then a read operation in an available partition queue with the most number of read operations for any memory component can be identified and transmitted over the data bus. However, if no such read operation can be identified in any available partition queue, then a

write operation from an available partition queue with the most number of write operations on any of the memory components can be identified and transmitted over the data bus. Thus, the data bus can switch from the read mode to the write mode.

[0015] If the data bus is in the write mode, then a write operation from an available partition queue with the most number of write operations on any memory component can be transmitted. If such a write operation is not identified, then the data bus can remain in the write mode or switch to the read mode based on an indication of a number of partitions that are associated with write operations. For example, the indication can identify a number of partition queues that include write operations. If the number of partition queues with write operations satisfies a threshold condition, then the data bus can remain in the write mode until a write operation is transmitted over the data bus. Otherwise, a read operation in the partition queue with the most number of read operations on any memory component can be retrieved and transmitted by the data bus.

[0016] Advantages of the present disclosure include, but are not limited to, an increased performance of a memory sub-system. For example, since the throughput of the bus of the memory sub-system can be increased by reducing the number of times that the data bus mode switches between read modes and write modes, the memory sub-system can be capable of performing more write operations and read operations at the memory components included in the memory sub-system.

[0017] **Fig. 1** illustrates an example computing environment 100 that includes a memory sub-system 110 in accordance with some embodiments of the present disclosure. The memory sub-system 110 can include media, such as memory components 112A to 112N. The memory components 112A to 112N can be volatile memory components, non-volatile memory components, or a combination of such. In some embodiments, the memory sub-system is a storage system. An example of a storage system is a SSD. In some embodiments, the memory sub-system 110 is a hybrid memory/storage sub-system. In general, the computing environment 100 can include a host system 120 that uses the memory sub-system 110. For example, the host system 120 can write data to the memory sub-system 110 and read data from the memory sub-system 110.

[0018] The host system 120 can be a computing device such as a desktop computer, laptop computer, network server, mobile device, or such computing device that includes a memory and a processing device. The host system 120 can include or be coupled to the memory sub-system 110 so that the host system 120 can read data from or write data to the memory sub-system 110.

The host system 120 can be coupled to the memory sub-system 110 via a physical host interface. As used herein, “coupled to” generally refers to a connection between components, which can be an indirect communicative connection or direct communicative connection (e.g., without intervening components), whether wired or wireless, including connections such as electrical, optical, magnetic, etc. Examples of a physical host interface include, but are not limited to, a serial advanced technology attachment (SATA) interface, a peripheral component interconnect express (PCIe) interface, universal serial bus (USB) interface, Fibre Channel, Serial Attached SCSI (SAS), etc. The physical host interface can be used to transmit data between the host system 120 and the memory sub-system 110. The host system 120 can further utilize an NVM Express (NVMe) interface to access the memory components 112A to 112N when the memory sub-system 110 is coupled with the host system 120 by the PCIe interface. The physical host interface can provide an interface for passing control, address, data, and other signals between the memory sub-system 110 and the host system 120.

[0019] The memory components 112A to 112N can include any combination of the different types of non-volatile memory components and/or volatile memory components. An example of non-volatile memory components includes a negative-and (NAND) type flash memory. Each of the memory components 112A to 112N can include one or more arrays of memory cells such as single level cells (SLCs) or multi-level cells (MLCs) (e.g., triple level cells (TLCs) or quad-level cells (QLCs)). In some embodiments, a particular memory component can include both an SLC portion and a MLC portion of memory cells. Each of the memory cells can store one or more bits of data (e.g., data blocks) used by the host system 120. Although non-volatile memory components such as NAND type flash memory are described, the memory components 112A to 112N can be based on any other type of memory such as a volatile memory. In some embodiments, the memory components 112A to 112N can be, but are not limited to, random access memory (RAM), read-only memory (ROM), dynamic random access memory (DRAM), synchronous dynamic random access memory (SDRAM), phase change memory (PCM), magneto random access memory (MRAM), negative-or (NOR) flash memory, electrically erasable programmable read-only memory (EEPROM), and a cross-point array of non-volatile memory cells. A cross-point array of non-volatile memory can perform bit storage based on a change of bulk resistance, in conjunction with a stackable cross-gridded data access array. Additionally, in contrast to many flash-based memories, cross-point non-volatile memory can perform a write in-place operation, where a non-volatile memory cell can be programmed without the non-volatile memory cell being previously erased. Furthermore, the memory cells of the memory components 112A to 112N can be grouped as memory pages or data blocks that can refer to a unit of the memory component used to store data.

[0020] The memory system controller 115 (hereinafter referred to as “controller”) can communicate with the memory components 112A to 112N to perform operations such as reading data, writing data, or erasing data at the memory components 112A to 112N and other such operations. The controller 115 can include hardware such as one or more integrated circuits and/or discrete components, a buffer memory, or a combination thereof. The controller 115 can be a microcontroller, special purpose logic circuitry (e.g., a field programmable gate array (FPGA), an application specific integrated circuit (ASIC), etc.), or other suitable processor. The controller 115 can include a processor (processing device) 117 configured to execute instructions stored in local memory 119. In the illustrated example, the local memory 119 of the controller 115 includes an embedded memory configured to store instructions for performing various processes, operations, logic flows, and routines that control operation of the memory sub-system 110, including handling communications between the memory sub-system 110 and the host system 120. In some embodiments, the local memory 119 can include memory registers storing memory pointers, fetched data, etc. The local memory 119 can also include read-only memory (ROM) for storing micro-code. While the example memory sub-system 110 in FIG. 1 has been illustrated as including the controller 115, in another embodiment of the present disclosure, a memory sub-system 110 may not include a controller 115, and may instead rely upon external control (e.g., provided by an external host, or by a processor or controller separate from the memory sub-system).

[0021] In general, the controller 115 can receive commands or operations from the host system 120 and can convert the commands or operations into instructions or appropriate commands to achieve the desired access to the memory components 112A to 112N. The controller 115 can be responsible for other operations such as wear leveling operations, garbage collection operations, error detection and error-correcting code (ECC) operations, encryption operations, caching operations, and address translations between a logical block address and a physical block address that are associated with the memory components 112A to 112N. The controller 115 can further include host interface circuitry to communicate with the host system 120 via the physical host interface. The host interface circuitry can convert the commands received from the host system into command instructions to access the memory components 112A to 112N as well as convert responses associated with the memory components 112A to 112N into information for the host system 120.

[0022] The memory sub-system 110 can also include additional circuitry or components that are not illustrated. In some embodiments, the memory sub-system 110 can include a cache or buffer (e.g., DRAM) and address circuitry (e.g., a row decoder and a column decoder) that can

receive an address from the controller 115 and decode the address to access the memory components 112A to 112N.

[0023] The memory sub-system 110 includes a sequencer component 113 that can be used to issue read operations or write operations. In some embodiments, the controller 115 includes at least a portion of the sequencer component 113. For example, the controller 115 can include a processor 117 (processing device) configured to execute instructions stored in local memory 119 for performing the operations described herein. In some embodiments, the sequencer component 113 is part of the host system 110, an application, or an operating system.

[0024] The sequencer component 113 can transmit read operations and write operations to the memory components 112A to 112N by utilizing a data bus (i.e., an interface) of the memory sub-system 110. The scheduling component 113 can determine particular operations at particular partition queues that can be retrieved and transmitted over the data bus to the various memory components 112A to 112N that correspond to different partitions. Further details with regards to the operations of the sequencer component 113 are described below.

[0025] **Fig. 2** illustrates an example scheduling of operations in accordance with some embodiments of the present disclosure. As shown, the sequencer component 113 can identify read operations and write operations in different partition queues and can transmit the read operations and write operations over a data bus to different partitions. The partition queues P0, P1, and P2 can correspond to a command queue of a memory sub-system. The partition queue P0 can include operations that are to read or write data at the partition P0, the partition queue P1 can include operations that are to read or write data at the partition P1, and the partition queue P2 can include operations that are to read or write data at the partition P2. The sequencer component 113 can identify a particular read operation or a write operation from any of the partition queues that can be transmitted over the data bus to the partitions P0, P1, and P2. As described in further detail below, the sequencer component 113 can identify a particular read operation or a particular write operation from a particular partition queue and can transmit the identified operation over the data bus.

[0026] **Fig. 3** is a flow diagram of an example method 300 to schedule the issuance of operations when a data bus is in a read mode in accordance with some embodiments. The method 300 can be performed by processing logic that can include hardware (e.g., processing device, circuitry, dedicated logic, programmable logic, microcode, hardware of a device, integrated circuit, etc.), software (e.g., instructions run or executed on a processing device), or a combination thereof. In some embodiments, the method 300 is performed by the sequencer

component 113 of **Fig. 1**. Although shown in a particular sequence or order, unless otherwise specified, the order of the processes can be modified. Thus, the illustrated embodiments should be understood only as examples, and the illustrated processes can be performed in a different order, and some processes can be performed in parallel. Additionally, one or more processes can be omitted in various embodiments. Thus, not all processes are required in every embodiment. Other process flows are possible.

[0027] As shown in **Fig. 3**, at operation 310, a processing device determines that a data bus is in a read mode. The data bus can be considered to be in a read mode when a read operation was the prior (i.e., the last) operation that was transmitted over the data bus. In some embodiments, if no operation has been transmitted over the data bus, then the data bus can be considered to be in the read mode. At operation 320, the processing device determines whether a read operation for an available partition queue with the most number of read operations for a current working memory component can be identified and/or, at operation 325, the processing device determines whether a read operation for an available partition queue with the most number of read operations for any memory component can be identified. At operation 330, the processing device determines whether a read operation is identified. If a read operation is identified, then at operation 340, the processing device transmits the read operation over the data bus to the corresponding memory component that is associated with the available partition queue of the read operation. In some embodiments, if a read operation is identified at an available partition queue with the most number of read operations for the current working memory component, then the read operation that is transmitted is the read operation for the current working memory component. In the same or alternative embodiments, if no read operation is identified at an available partition queue for the current working memory component, then another read can be identified at an available partition queue with the most number of read operations for any memory component. Subsequently, the read operation can be transmitted. As such, if a read operation is present at an available partition queue and is to retrieve data at the current working memory component, then such a read operation can be transmitted over the data bus before a read operation that is present in an available partition queue that is to retrieve data at another memory component that is different than the current working memory component.

[0028] Referring to **Fig. 3**, if the read operation is not identified at operation 330, then the processing device determines whether a write operation in an available partition queue with the most write operations can be transmitted over the data bus. For example, if multiple write operations in multiple available partition queues are present, then a write operation at the available partition queue with the highest number of write operations (relative to other available

partition queues) can be identified. For example, at operation 340, the processing device transmits the write operation over the data bus. Otherwise, if a write operation in an available partition queue is not identified, then at operation 360, the processing device waits for a read operation or a write operation that can be transmitted over the data bus. For example, the processing device can return to operation 310.

[0029] As such, if the data bus is in a read mode, a determination can be made as to whether any available partition queue includes a read operation for a current working memory component. If such a read operation is identified, then the read operation from the available partition queue with the most read operations for the current working memory component can be transmitted over the data bus. Otherwise, if such a read operation is not identified, then a determination can be made as to whether any available partition queue includes a read operation for any memory component. If such a read operation is identified, then a read operation from an available partition queue with the most read operations for any memory component can be transmitted over the data bus. Otherwise, a write operation can be transmitted over the data bus. For example, a write operation in an available partition queue that includes the most write operations on any memory component can be transmitted over the data bus.

[0030] **Fig. 4** is a flow diagram of an example method 400 to schedule the issuance of operations when a data bus is in a write mode in accordance with some embodiments. The method 400 can be performed by processing logic that can include hardware (e.g., processing device, circuitry, dedicated logic, programmable logic, microcode, hardware of a device, integrated circuit, etc.), software (e.g., instructions run or executed on a processing device), or a combination thereof. In some embodiments, the method 400 is performed by the sequencer component 113 of **Fig. 1**. Although shown in a particular sequence or order, unless otherwise specified, the order of the processes can be modified. Thus, the illustrated embodiments should be understood only as examples, and the illustrated processes can be performed in a different order, and some processes can be performed in parallel. Additionally, one or more processes can be omitted in various embodiments. Thus, not all processes are required in every embodiment. Other process flows are possible.

[0031] As shown in **Fig. 4**, at operation 410, a processing device determines that a data bus is in a write mode. The data bus can be considered to be in a write mode when a write operation was the prior (i.e., the last) operation that was transmitted over the data bus. At operation 420, the processing device determines that a write operation in an available partition queue is not identified. For example, a write operation is not present in any available partition queue. In some embodiments, if a write operation is present in multiple partition queues, then the write operation

that is present in the available partition queue with the most write operations can be transmitted over the data bus. At operation 430, the processing device identifies a number of partition queues that include write operations. At operation 440, the processing device determines whether the number of partition queues exceeds a threshold number of partition queues. The number of partition queues can be considered to exceed the threshold number of partition queues when the number of partition queues is equal to or exceeds the threshold number. If the number of partition queues satisfies the threshold number, then, at operation 450, the processing device waits until a write operation in an available partition queue is identified. For example, the processing device can pause transmitting an operation over the data bus until a partition queue with a write operation becomes available and then the write operation can be transmitted over the data bus. Otherwise, if the number of partition queues does not satisfy the threshold number, then, at operation 460, the processing device transmits a read operation in a partition queue with the most read operations if such a read operation is available. For example, the read operation can be transmitted over the data bus. Otherwise, if the read operation is not available, then the processing device waits to transmit an operation and returns to operation 410.

[0032] As such, when a data bus is in the write mode, a write operation in the available partition queue with the most write operations on any memory component can be transmitted over the data bus. Otherwise, a read operation or a write operation is transmitted over the data bus based on a number of partition queues that include a write operation.

[0033] **Fig. 5** is a flow diagram of an example method 500 to schedule read operations and write operations in accordance with some embodiments of the present disclosure. The method 500 can be performed by processing logic that can include hardware (e.g., processing device, circuitry, dedicated logic, programmable logic, microcode, hardware of a device, integrated circuit, etc.), software (e.g., instructions run or executed on a processing device), or a combination thereof. In some embodiments, the method 500 is performed by the sequencer component 113 of **Fig. 1**. Although shown in a particular sequence or order, unless otherwise specified, the order of the processes can be modified. Thus, the illustrated embodiments should be understood only as examples, and the illustrated processes can be performed in a different order, and some processes can be performed in parallel. Additionally, one or more processes can be omitted in various embodiments. Thus, not all processes are required in every embodiment. Other process flows are possible.

[0034] As shown in **Fig. 5**, at operation 510, a processing device determines whether a data bus is in a read mode. The data bus can be in a read mode when the prior operation that was transmitted over the data bus to a partition was a read operation. Otherwise, if the prior operation

(i.e., the last operation) that was transmitted over the data bus was a write operation, then the data bus can be determined to be in a write mode. At operation 511, the processing device determines whether a read operation (RD) in an available partition queue with the most number of read operations on a current working memory component can be identified. A partition queue can be considered to be available when transmitting an operation to the partition queue satisfies the data bus protocol requirements. For example, a partition queue can be considered to be available when transmitting a read operation or write operation for the partition queue over the data bus would not violate any of the requirements of the data bus protocol. Otherwise, the partition queue can be considered to be unavailable when transmitting such operations for the partition queue over the data bus would violate at least one requirement of the data bus protocol. The current working memory component can refer to the memory component that includes the partition specified by or associated with the last or prior operation transmitted over the data bus. At operation 512, the processing device determines whether such a read operation is found. For example, a read operation that is specifying a partition that is included in the working memory component can be identified. If such a read operation is found or identified, then at operation 514, the processing device transmits the read operation over the data bus if the request can be issued. In some embodiments, the read operation can be issued over the data bus if the data bus is not currently transmitting another command. The read operation can then be transmitted after the data bus becomes available. In some embodiments, if the read operation cannot be issued, then the processing device waits to transmit an operation as specified by operation 525. Otherwise, if such a read operation cannot be found at operation 512, then at operation 513, the processing device determines whether a read operation in an available partition queue with the most number of read operations on any memory component can be retrieved. At operation 514, the processing device determines whether such a read operation can be found. For example, the determination can be whether any partition queue includes at least one read operation. If such a read operation can be found, then at operation 530, then the processing device transmits the read operation over the data bus if the read operation can be transmitted. Otherwise, at operation 515, if the read operation at operation 514 is not found, then the processing device determines whether a write operation can be found. For example, a determination can be made as to whether any available partition queue includes a write operation and the partition queue with the most number of write operations can be identified. A particular write operation from the write operations in the identified partition queue can then be selected. In some embodiments, the write operation that is oldest (e.g., earliest issued by a host system) can be selected. If such a write operation is found, then at operation 530, the processing device transmits the write operation over the data bus if the write operation can be issued. Otherwise, at operation 525, the processing

device waits to transmit an operation. For example, the processing device can wait until a write operation at an available partition queue can be identified. After an operation has been issued at operation 530, the processing device updates, at operation 531, the data bus to be in the read mode or the write mode based on the transmitted operation. For example, the data bus can remain or be changed to be in the read mode if the last operation that was transmitted over the data bus was a read operation (i.e., a read operation) and the data bus can remain or be changed to be in the write mode if the last operation that was transmitted over the data bus was a write operation (i.e., a write operation). Subsequently, at operation 510, the processing device determines whether the data bus is in the read mode as previously described.

[0035] Referring to **Fig. 5**, if at operation 510 the processing device determines that the data bus is not in a read mode, then the data bus can be in the write mode. At operation 520, the processing device determines whether a write operation in an available partition queue with the most write operations on the working memory component can be identified. At operation 521, a determination is made as to whether the write operation is found. If the write operation is found, then the write operation can be transmitted over the data bus if the write operation can be issued (at operation 530). Otherwise, if the write operation is not found, then, at operation 522, the processing device determines whether a number of partition queues with write operations satisfy a threshold number of partition queues. The number of partition queues can satisfy the threshold number of partition queues when the number of partition queues that includes at least one write operation is the same as or exceeds the threshold number of partition queues. If the number of partition queues with write operations satisfies the threshold, then at operation 525, the processing device waits to transmit an operation. For example, the processing device can wait until a partition queue with at least one write operation becomes available and can be transmitted over the data bus. Otherwise, if the number of partition queues with write requests does not satisfy the threshold, then at operation 523, the processing device determines whether a read operation in an available partition queue with the most read operations on any working memory component can be retrieved. At operation 524, the processing device determines whether such a read operation is found. If the read operation is not found, then at operation 525, the processing device waits to transmit an operation. Otherwise, if the read operation is found, then at operation 530, the read operation is transmitted over the data bus.

[0036] **Fig. 6** illustrates an example machine of a computer system 600 within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, can be executed. In some embodiments, the computer system 600 can correspond to a host system (e.g., the host system 120 of **Fig. 1**) that includes, is coupled to, or utilizes a memory

sub-system (e.g., the memory sub-system 110 of **Fig. 1**) or can be used to perform the operations of a controller (e.g., to execute an operating system to perform operations corresponding to the sequencer component 113 of **Fig. 1**). In alternative embodiments, the machine can be connected (e.g., networked) to other machines in a LAN, an intranet, an extranet, and/or the Internet. The machine can operate in the capacity of a server or a client machine in client-server network environment, as a peer machine in a peer-to-peer (or distributed) network environment, or as a server or a client machine in a cloud computing infrastructure or environment.

[0037] The machine can be a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a server, a network router, a switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0038] The example computer system 600 includes a processing device 602, a main memory 604 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) such as synchronous DRAM (SDRAM) or Rambus DRAM (RDRAM), etc.), a static memory 606 (e.g., flash memory, static random access memory (SRAM), etc.), and a data storage system 618, which communicate with each other via a bus 630.

[0039] Processing device 602 represents one or more general-purpose processing devices such as a microprocessor, a central processing unit, or the like. More particularly, the processing device can be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets, or processors implementing a combination of instruction sets. Processing device 602 can also be one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. The processing device 602 is configured to execute instructions 626 for performing the operations and steps discussed herein. The computer system 600 can further include a network interface device 608 to communicate over the network 620.

[0040] The data storage system 618 can include a machine-readable storage medium 624 (also known as a computer-readable medium) on which is stored one or more sets of instructions 626 or software embodying any one or more of the methodologies or functions described herein. The instructions 626 can also reside, completely or at least partially, within the main memory

604 and/or within the processing device 602 during execution thereof by the computer system 600, the main memory 604 and the processing device 602 also constituting machine-readable storage media. The machine-readable storage medium 624, data storage system 618, and/or main memory 604 can correspond to the memory sub-system 110 of **Fig. 1**.

[0041] In one embodiment, the instructions 626 include instructions to implement functionality corresponding to a sequencer component (e.g., the sequencer component 113 of **Fig. 1**). While the machine-readable storage medium 624 is shown in an example embodiment to be a single medium, the term “machine-readable storage medium” should be taken to include a single medium or multiple media that store the one or more sets of instructions. The term “machine-readable storage medium” shall also be taken to include any medium that is capable of storing or encoding a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present disclosure. The term “machine-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media, and magnetic media.

[0042] Some portions of the preceding detailed descriptions have been presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the ways used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0043] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. The present disclosure can refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage systems.

[0044] The present disclosure also relates to an apparatus for performing the operations herein. This apparatus can be specially constructed for the intended purposes, or it can include a general purpose computer selectively activated or reconfigured by a computer program stored in

the computer. Such a computer program can be stored in a computer readable storage medium, such as, but not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, each coupled to a computer system bus.

[0045] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems can be used with programs in accordance with the teachings herein, or it can prove convenient to construct a more specialized apparatus to perform the method. The structure for a variety of these systems will appear as set forth in the description below. In addition, the present disclosure is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages can be used to implement the teachings of the disclosure as described herein.

[0046] The present disclosure can be provided as a computer program product, or software, that can include a machine-readable medium having stored thereon instructions, which can be used to program a computer system (or other electronic devices) to perform a process according to the present disclosure. A machine-readable medium includes any mechanism for storing information in a form readable by a machine (e.g., a computer). In some embodiments, a machine-readable (e.g., computer-readable) medium includes a machine (e.g., a computer) readable storage medium such as a read only memory (“ROM”), random access memory (“RAM”), magnetic disk storage media, optical storage media, flash memory components, etc.

[0047] In the foregoing specification, embodiments of the disclosure have been described with reference to specific example embodiments thereof. It will be evident that various modifications can be made thereto without departing from the broader spirit and scope of embodiments of the disclosure as set forth in the following claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense.

CLAIMS

What is claimed is:


1. A method comprising:
 - determining that a data bus is in a write mode based on a prior operation transmitted over the data bus being a write operation;
 - in response to determining that the data bus is in the write mode, identifying a number of partition queues of a plurality of partition queues that include at least one write operation;
 - determining, by a processing device, whether the number of partition queues of the plurality of partition queues satisfies a threshold number; and
 - in response to determining that the number of partition queues satisfies the threshold number, transmitting another write operation from the plurality of partition queues over the data bus.
2. The method of claim 1, further comprising:
 - in response to determining that the number of partition queues does not satisfy the threshold number, transmitting a read operation from the plurality of partition queues over the data bus.
3. The method of claim 2, further comprising:
 - determining that a particular partition queue of the plurality of partition queues is available, wherein the particular partition queue is available when transmitting the read operation from the particular partition queue over the data bus satisfies a data bus protocol.
4. The method of claim 2, further comprising:
 - in response to transmitting the read operation from the plurality of partition queues over the data bus, changing a status of the data bus from the write mode to a read mode.
5. The method of claim 1, wherein the determining of whether the number of partition queues of the plurality of partition queues satisfies the threshold number is in response to a write operation not being identified at an available partition queue.
6. The method of claim 1, wherein the prior operation corresponds to a last operation that was transmitted over the data bus.

7. The method of claim 1, wherein the write operation is associated with a particular partition queue of the plurality of partition queues with a higher number of write operations than other partition queues of the plurality of partition queues.
8. A system comprising:
 - a memory component; and
 - a processing device, operatively coupled with the memory component, to:
 - determine that a data bus is in a read mode based on a prior operation transmitted over the data bus being a read operation;
 - in response to determining that the data bus is in the read mode, identify a plurality of partition queues associated with the data bus;
 - identify a particular read operation included in a particular partition queue of the plurality of partition queues, the particular partition queue including a highest number of read operations for a current working memory component of a plurality of memory components associated with the plurality of partition queues; and
 - transmit the particular read command from the particular partition queue over the data bus.
9. The system of claim 8, wherein the current working memory component corresponds to a particular memory component specified by the prior operation transmitted over the data bus.
10. The system of claim 8, wherein the prior operation corresponds to a last operation that was transmitted over the data bus.
11. The system of claim 8, wherein the particular partition queue is an available partition queue, and wherein the particular partition queue is the available partition queue when transmitting the particular read operation satisfies a data bus protocol.
12. The system of claim 8, wherein each partition queue of the plurality of partition queues corresponds to the plurality of memory components, and wherein the data bus is coupled with the plurality of memory components.
13. The system of claim 8, wherein the transmitting of the particular read command from the particular partition queue over the data bus is based on a data bus protocol.


14. A non-transitory computer readable medium comprising instructions, which when executed by a processing device, cause the processing device to perform operations comprising:
- determining that a data bus is in a write mode based on a prior operation transmitted over the data bus being a write operation;
 - in response to determining that the data bus is in the write mode, identifying a number of partition queues of a plurality of partition queues that include at least one write operation;
 - determining whether the number of partition queues of the plurality of partition queues satisfies a threshold number; and
 - in response to determining that the number of partition queues satisfies the threshold number, transmitting another write operation from the plurality of partition queues over the data bus.
15. The non-transitory computer readable medium of claim 14, wherein the operations further comprise:
- in response to determining that the number of partition queues does not satisfy the threshold number, transmitting a read operation from the plurality of partition queues over the data bus
16. The non-transitory computer readable medium of claim 15, wherein the operations further comprise:
- determining that a particular partition queue of the plurality of partition queues is available, wherein the particular partition queue is available when transmitting the read operation from the particular partition queue over the data bus satisfies a data bus protocol.
17. The non-transitory computer readable medium of claim 15, wherein the operations further comprise:
- in response to transmitting the read operation from the plurality of partition queues over the data bus, changing a status of the data bus from the write mode to a read mode.
18. The non-transitory computer readable medium of claim 14, wherein the determining of whether the number of partition queues of the plurality of partition queues satisfies the threshold number is in response to a write operation not being identified at an available partition queue.
19. The non-transitory computer readable medium of claim 14, wherein the prior operation corresponds to a last operation that was transmitted over the data bus.

20. The non-transitory computer readable medium of claim 14, wherein the write operation is associated with a particular partition queue of the plurality of partition queues with a higher number of write operations than other partition queues of the plurality of partition queues.

1/6

**FIG. 1**

2/6

**FIG. 2**

3/6




FIG. 3

4/6





FIG. 4



6/6




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2019/051579**A. CLASSIFICATION OF SUBJECT MATTER****G06F 3/06(2006.01)i, G06F 13/16(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 3/06; G06F 012/00; G06F 12/00; G06F 13/28; G06F 7/00; H04N 5/76; G06F 13/16

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: memory, data bus, queue, switch, priority

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5948081 A (JOSEPH E. FOSTER) 07 September 1999 See columns 2-3; claims 1, 15; and figure 4.	1-20
A	KR 10-0531324 B1 (LG ELECTRONICS INC.) 28 November 2005 See claims 1-14; and figure 3.	1-20
A	US 2003-0088751 A1 (SHENG-CHUNG WU et al.) 08 May 2003 See paragraphs [0023]-[0033]; and figures 4-5.	1-20
A	US 7047374 B2 (SUNEETA SAH et al.) 16 May 2006 See claims 1-8; and figures 9-15.	1-20
A	WO 2013-032816 A1 (APPLE INC.) 07 March 2013 See paragraphs [0005]-[0006], [0042]-[0056]; and figures 1-3.	1-20



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

03 January 2020 (03.01.2020)

Date of mailing of the international search report

03 January 2020 (03.01.2020)

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/US2019/051579

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 05948081 A	07/09/1999	None	
KR 10-0531324 B1	28/11/2005	KR 10-2005-0077519 A	03/08/2005
US 2003-0088751 A1	08/05/2003	TW 561396 B	11/11/2003
US 7047374 B2	16/05/2006	US 06912612 B2	28/06/2005
		US 2003-163649 A1	28/08/2003
		US 2003-177320 A1	18/09/2003
		US 2004-022094 A1	05/02/2004
WO 2013-032816 A1	07/03/2013	AU 2012-216395 A1	21/03/2013
		AU 2012-216395 B2	09/01/2014
		BR 102012021854 A2	08/10/2013
		CN 102968289 A	13/03/2013
		CN 102968289 B	20/04/2016
		EP 2565795 A2	06/03/2013
		EP 2565795 A3	04/09/2013
		EP 2565795 B1	20/05/2015
		JP 2013-058205 A	28/03/2013
		JP 5499124 B2	21/05/2014
		KR 10-1408434 B1	17/06/2014
		KR 10-2013-0024860 A	08/03/2013
		TW 201319814 A	16/05/2013
		TW I493350 B	21/07/2015
		US 09021146 B2	28/04/2015
		US 2013-0054875 A1	28/02/2013