



(51) International Patent Classification:
G06F 13/16 (2006.01)

(21) International Application Number:
PCT/US2020/066029

(22) International Filing Date:
18 December 2020 (18.12.2020)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/728,152 27 December 2019 (27.12.2019) US

(71) Applicant: **ADVANCED MICRO DEVICES, INC.**
[US/US]; 2485 Augustine Drive, Santa Clara, California
95054 (US).

(72) Inventors: **MAGRO, James R.**; 7171 Southwest Parkway,
Austin, Texas 78735 (US). **BALAKRISHNAN, Kedar-**
nath; #102-103, Export Promotion Industrial Park, White-
field, Bangalore 560066 (IN). **BHARGAVA, Ravindra**
N.; 7171 Southwest Parkway, Austin, Texas 78735 (US).
SHEN, Guanhao; 7171 Southwest Parkway, Austin, Texas
78735 (US).

(74) Agent: **KENNEDY, Brandon** et al.; Kennedy Lenart
Spraggins LLP, 797 Sam Bass Road #2559, Round Rock,
Texas 78681 (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, IT, 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, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

(54) Title: STAGING BUFFER ARBITRATION




FIG. 1

(57) Abstract: Staging buffer arbitration includes: storing a plu-
rality of memory access requests in a staging buffer; selecting
a memory access request of the plurality of memory access re-
quests from the staging buffer based on one or more arbitration
rules; and moving the memory access request from the staging
buffer to a command queue.

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

STAGING BUFFER ARBITRATION

BACKGROUND

[0001] Command queues store memory access requests prior to execution. Increased arbitration and storage pressure on the command queue results in a performance decrease.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Figure 1 is a block diagram of an example processor for staging memory access requests according to some embodiments.

[0003] Figure 2 is a flowchart of an example method for staging memory access requests according to some embodiments.

[0004] Figure 3 is a flowchart of an example method for staging memory access requests according to some embodiments.

[0005] Figure 4 is a flowchart of an example method for staging memory access requests according to some embodiments.

[0006] Figure 5 is a flowchart of an example method for staging memory access requests according to some embodiments.

[0007] Figure 6 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0008] Figure 7 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0009] Figure 8 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0010] Figure 9 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0011] Figure 10 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0012] Figure 11 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0013] Figure 12 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

[0014] Figure 13 is a flowchart of an example method for staging buffer arbitration according to some embodiments.

DETAILED DESCRIPTION

[0015] In some embodiments, a method of staging buffer arbitration includes storing a plurality of memory access requests in a staging buffer; selecting a memory access request of the plurality of memory access requests from the staging buffer based on one or more arbitration rules; and moving the memory access request from the staging buffer to a command queue.

[0016] In some embodiments, selecting the memory access request includes selecting a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on one or more of: a bank targeted by another memory access request, a rank targeted by another memory access request, or a memory subchannel targeted by another memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on a Dynamic Random Access Memory page targeted by another memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on a priority value. In some embodiments, selecting the memory access request includes: identifying, in the staging buffer, another memory access request associated with a first page miss; identifying, in the command queue, a queued memory access request associated with a second page miss different from the first page miss; and selecting the memory access request in response to identifying the other memory access request and the queued memory access request. In some embodiments, selecting the memory access request includes identifying, in the staging buffer, another memory access request associated with a first page conflict; identifying, in the command queue, a queued memory access request associated with a second page conflict different from the first page conflict; and selecting the memory access request in response to identifying the other memory access request and the queued memory access request. In some embodiments, selecting the memory access request includes: identifying, in the command queue, a page hit request; and selecting the memory access request based on the memory access being another page hit request.

[0017] In some embodiments, a memory management unit for staging buffer arbitration performs steps including: storing a plurality of memory access requests in a staging buffer; selecting a memory access request of the plurality of memory access requests from the staging buffer based on one or more arbitration rules; and moving the memory access request from the staging buffer to a command queue.

[0018] In some embodiments, selecting the memory access request includes selecting a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on one or more of: a bank targeted by another memory access request, a rank targeted by another memory access request, or a memory subchannel targeted by another memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on a Dynamic Random Access Memory page targeted by another memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on a priority value. In some embodiments, selecting the memory access request includes: identifying, in the staging buffer, another memory access request associated with a first page miss; identifying, in the command queue, a queued memory access request associated with a second page miss different from the first page miss; and selecting the memory access request in response to identifying the other memory access request and the queued memory access request. In some embodiments, selecting the memory access request includes identifying, in the staging buffer, another memory access request associated with a first page conflict; identifying, in the command queue, a queued memory access request associated with a second page conflict different from the first page conflict; and selecting the memory access request in response to identifying the other memory access request and the queued memory access request. In some embodiments, selecting the memory access request includes: identifying, in the command queue, a page hit request; and selecting the memory access request based on the memory access being another page hit request.

[0019] In some embodiments, an apparatus for staging buffer arbitration includes a processor, the processor including a memory management unit to perform steps including: storing a plurality of memory access requests in a staging buffer; selecting a memory access request of the plurality of memory access requests from the staging buffer based on one or more arbitration rules; and moving the memory access request from the staging buffer to a command queue.

[0020] In some embodiments, selecting the memory access request includes selecting a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on one or more of: a bank targeted by another memory access request, a rank targeted by another memory access

request, or a memory subchannel targeted by another memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on a Dynamic Random Access Memory page targeted by another memory access request. In some embodiments, selecting the memory access request includes selecting the memory access request based on a priority value. In some embodiments, selecting the memory access request includes: identifying, in the staging buffer, another memory access request associated with a first page miss; identifying, in the command queue, a queued memory access request associated with a second page miss different from the first page miss; and selecting the memory access request in response to identifying the other memory access request and the queued memory access request. In some embodiments, selecting the memory access request includes identifying, in the staging buffer, another memory access request associated with a first page conflict; identifying, in the command queue, a queued memory access request associated with a second page conflict different from the first page conflict; and selecting the memory access request in response to identifying the other memory access request and the queued memory access request. In some embodiments, selecting the memory access request includes: identifying, in the command queue, a page hit request; and selecting the memory access request based on the memory access being another page hit request.

[0021] Figure 1 is a block diagram of a non-limiting example processor 100 according to some embodiments. The example processor 100 can be implemented in a variety of computing devices, including mobile devices, personal computers, peripheral hardware components, gaming devices, set-top boxes, and the like. The processor 100 includes a memory management unit 102. The memory management unit 102 receives memory access requests (e.g., requests to read and/or write data to a particular region of memory). The memory management unit 102 also performs the translation of virtual memory addresses in the memory access requests to physical memory addresses in order to perform the memory access request.

[0022] The memory management unit 102 includes a command queue 104. The command queue 104 stores memory access requests as they are received (e.g., from a central processing unit or other component of the processor 100) prior to execution to access Dynamic Random Access Memory 106. Although the Dynamic Random Access Memory 106 is shown as being separate from the processor 100, it is understood that the Dynamic Random Access Memory 106 may include on-chip Dynamic Random Access Memory 106 (e.g., as a component of the processor 100). Where the Dynamic Random Access Memory 106

includes multiple banks, the memory management unit 102 may include multiple command queues 104 each corresponding to a respective bank.

[0023] The memory management unit 102 selects memory access requests for execution from the command queue 104 using one or more schemes, such as first-come-first-served (FCFS), first-ready, first-come-first-served (FR-FCFS), first-in-first-out (FIFO), etc. The memory management unit 102 includes a command queue arbiter 108 that selects memory access requests from the command queue 104 for execution using one or more rules. For example, the one or more rules are based on timing or clock information (e.g., an age of a memory access request). As another example, the one or more rules are based on a page table 110. For example, memory access requests that will result in a page table 110 hit are preferentially selected from the command queue 104 for execution.

[0024] In existing approaches, memory access requests received by a memory management unit 102 are placed directly in a command queue 104 for subsequent execution. To improve performance and relieve pressure on the command queue 104, the memory management unit 102 includes a staging buffer 112. Memory access requests received by the memory management unit 102 are placed in the staging buffer 112. A staging buffer arbiter 114 then selects, based on one or more arbitration rules, memory access requests from the staging buffer 112 for movement to the command queue 104.

[0025] In some embodiments, the arbitration rules are based on a Dynamic Random Access Memory 106 page targeted by a memory access request. For example, a memory access request targeting a Dynamic Random Access Memory 106 page that is open is preferentially selected for movement to the command queue 104 as overhead required in closing and opening pages is reduced. As another example, a memory access request targeting a Dynamic Random Access Memory 106 page that is also targeted by another memory access request in the command queue 104, and therefore will be open when the selected memory access request is executed, is preferentially selected.

[0026] In some embodiments, the arbitration rules are based on a bank group rotation or rank rotation. For example, where the Dynamic Random Access Memory 106 includes multiple banks, memory access requests are selected from the staging buffer 112 for addition to the command queue 104 such that consecutively added requests do not target a same bank. As an example, a memory access request targeting a first bank is moved to the command queue 104, then a memory access request targeting a second bank is moved to the command queue 104. Another memory access request targeting the first bank is then be added to the command queue 104, etc. In some embodiments, memory access requests are selected to

target different ranks within the same or different banks or to alternatively target different subchannels of Dynamic Random Access Memory 106 (e.g., sub-channel balancing).

[0027] In some embodiments, the arbitration rules are based on a request type for the memory access requests (e.g., read or write). As there is computational overhead in switching between Dynamic Random Access Memory 106 reads and writes, read requests and/or write requests are grouped together as part of a “burst” of requests for movement to the command queue 104. Thus, a group of read requests and/or a group of write requests may be executed consecutively.

[0028] In some embodiments, the memory management unit 102 determines that the staging buffer 112 is full. Accordingly, the memory management unit 102 stores a received memory access request directly in the command queue 104 without storing the received memory access request in the staging buffer 112.

[0029] For further explanation, Figure 2 sets forth a flow chart illustrating an exemplary method for staging memory access requests according to embodiments of the present disclosure that includes receiving 202 (e.g., by a memory management unit 102 of a processor 100) a memory access request 204 directed to Dynamic Random Access Memory 106. The memory access request 204 includes a request to read or write data to or from Dynamic Random Access Memory 106. The memory access request 204 is received via a data fabric or other interconnect coupling the memory management unit 102 to a central processing unit or other component.

[0030] The method of Figure 2 also includes storing 206 (e.g., by the memory management unit 102) the memory access request 204 in a staging buffer 112. The method of Figure 2 also includes moving 208 the memory access request 204 from the staging buffer 112 to a command queue 104. Moving 208 the memory access request 204 includes removing the memory access request 204 from the staging buffer 112 and storing the memory access request 204 in the command queue 104. In some embodiments, the memory access request 204 is moved in response to a number of memory access requests 204 in the staging buffer 112 meeting a threshold. In some embodiments, the memory access request 204 is moved in response to a number of memory access requests 204 in the command queue 104 falling below a threshold. In some embodiments, the memory access request 204 is moved in response to an age of the memory access request 204 (e.g., a time at which the memory access request 204 was received) meeting a threshold. In some embodiments the memory access request 204 is moved in response to one or more arbitration rules being satisfied.

[0031] For further explanation, Figure 3 sets forth a flow chart illustrating an exemplary method for staging memory access requests according to embodiments of the present disclosure that includes receiving 202 (e.g., by a memory management unit 102 of a processor 100) a memory access request 204 directed to Dynamic Random Access Memory 106; storing 206 the memory access request 204 in a staging buffer 112; and moving 208 the memory access request 204 from the staging buffer 112 to a command queue 104.

[0032] The method of Figure 3 differs from Figure 2 in that the method of Figure 3 also includes selecting 302 (e.g., by the memory management unit 102 of the processor 100) the memory access request 204 from the command queue 104. The memory access request 204 is selected from the command queue 104 using one or more schemes, such as first-come-first-served (FCFS), first-ready, first-come-first-served (FR-FCFS), first-in-first-out (FIFO), etc. In some embodiments, the memory management unit 102 includes a command queue arbiter 108 that selects memory access requests from the command queue 104 for execution using one or more rules. For example, the one or more rules are based on timing or clock information (e.g., an age of a memory access request). As another example, the one or more rules are based on a page table 110. For example, memory access requests that will result in a page table 110 hit are preferentially selected from the command queue 104 for execution.

[0033] The method of Figure 3 differs from Figure 2 in that the method of Figure 3 also includes executing 304 (e.g., by the memory management unit 102), the memory access request 204. Executing 304 the memory access request 204 includes reading data from a Dynamic Random Access Memory 106 address specified in the memory access request 204 and/or writing data to a Dynamic Random Access Memory 106 address specified in the memory access request 204.

[0034] For further explanation, Figure 4 sets forth a flow chart illustrating an exemplary method for staging memory access requests according to embodiments of the present disclosure that includes receiving 202 (e.g., by a memory management unit 102 of a processor 100) a memory access request 204 directed to Dynamic Random Access Memory 106; storing 206 the memory access request 204 in a staging buffer 112; and moving 208 the memory access request 204 from the staging buffer 112 to a command queue 104.

[0035] The method of Figure 4 differs from Figure 2 in that the method of Figure 4 also includes receiving 402 another memory access request 404. The method of Figure 4 further differs from Figure 2 in that the method of Figure 4 also includes determining 406 that the staging buffer 112 is full. The staging buffer 112 includes a predefined amount of memory for storing a predefined maximum number of memory access requests. Accordingly,

determining 406 that the staging buffer 112 is full includes determining that the staging buffer 112 is storing the predefined maximum number of memory access requests.

[0036] The method of Figure 4 further differs from Figure 2 in that the method of Figure 4 also includes storing 408 the other memory access request 404 in the command queue 104 without storing the other memory access request 404 in the staging buffer 112. Thus, the staging buffer 112 is bypassed when full.

[0037] For further explanation, Figure 5 sets forth a flow chart illustrating an exemplary method for staging memory access requests according to embodiments of the present disclosure that includes receiving 202 (e.g., by a memory management unit 102 of a processor 100) a memory access request 204 directed to Dynamic Random Access Memory 106; storing 206 the memory access request 204 in a staging buffer 112; and moving 208 the memory access request 204 from the staging buffer 112 to a command queue 104.

[0038] The method of Figure 5 differs from Figure 2 in that the method of Figure 5 also includes selecting 502 (e.g., by the memory management unit 102, by a staging buffer arbiter 114 of the memory management unit 102), based on one or more arbitration rules, the memory access request 204 from a plurality of memory access requests in the staging buffer 112 for moving to the command queue 104.

[0039] In some embodiments, the arbitration rules are based on a Dynamic Random Access Memory 106 page targeted by a memory access request. For example, a memory access request targeting a Dynamic Random Access Memory 106 page that is open is preferentially selected for movement to the command queue 104 as overhead required in closing and opening pages is reduced. As another example, a memory access request targeting a Dynamic Random Access Memory 106 page that is also targeted by another memory access request in the command queue 104, and therefore will be open when the selected memory access request is executed, is preferentially selected.

[0040] In some embodiments, the arbitration rules are based on a bank group rotation or rank rotation. For example, where the Dynamic Random Access Memory 106 includes multiple banks, memory access requests are selected from the staging buffer 112 for addition to the command queue 104 such that consecutively added requests do not target a same bank. As an example, a memory access request targeting a first bank is moved to the command queue 104, then a memory access request targeting a second bank is moved to the command queue 104. Another memory access request targeting the first bank is then added to the command queue 104, etc. In some embodiments, memory access requests are selected to target different ranks within the same or different banks. Memory access requests are also selected

to alternatively target different subchannels of Dynamic Random Access Memory 106 (e.g., sub-channel balancing).

[0041] In some embodiments, the arbitration rules are based on a request type for the memory access requests (e.g., read or write). As there is computational overhead in switching between Dynamic Random Access Memory 106 reads and writes, read requests and/or write requests are grouped together as part of a “burst” of requests for movement to the command queue 104. Thus, a group of read requests and/or a group of write requests are executed consecutively.

[0042] For further explanation, Figure 6 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112. The memory access requests include a request to read or write data to or from Dynamic Random Access Memory 106. The memory access requests are via a data fabric or other interconnect coupling the memory management unit 102 to a central processing unit or other component.

[0043] The method of Figure 6 also includes selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112. For example, a staging buffer arbiter 114 selects the memory access request 608 based on the one or more arbitration rules. The arbitration rules are applied to various attributes of the memory access requests stored in the staging buffer 112, memory access requests stored in a command queue 104, a page table 110, and/or other attributes. For example, the arbitration rules are based on request type of memory access requests in the staging buffer 112 and/or command queue 104, a currently open Dynamic Random Access Memory 106 page, bank groups targeted by the memory access requests in the staging buffer 112 and/or command queue 104, refresh state of a bank or page targeted by memory access requests in the staging buffer 112 and/or command queue 104, and/or sub-channels targeted by memory access requests in the staging buffer 112 and/or command queue 104.

[0044] The method of Figure 6 also includes moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104. Moving 610 the memory access request 608 includes deleting the memory access request 608 from the staging buffer 112 and/or freeing a portion of the staging buffer 112 storing the memory access request 608 for subsequent overwriting. Moving 610 the memory access request 608 also includes adding the memory access request 608 to the command queue 104. Thus, the memory access request 608 is later executed from the command queue 104 by the memory management unit 102.

[0045] For further explanation, Figure 7 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0046] The method of Figure 7 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes selecting 702 a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request 608. A memory access request burst includes a plurality of memory access requests of the same type (e.g., read or write). The memory access requests in the memory access request burst are selected for movement to the command queue 104 consecutively and/or at least partially simultaneously such that the memory access requests in the memory access request burst are later executed consecutively and/or at least partially simultaneously. For example, a burst of read requests are executed without executing an intervening write request. As another example, a burst of write requests are executed without executing an intervening read request. As switching between executing read and write requests to Dynamic Random Access Memory 106 costs computational overhead, this computational overhead is avoided by executing multiple memory access requests of a same request type. Accordingly, the memory access request 608 is selected based on other memory access requests of the same request type having been added to the command queue 104 and or based on other memory access requests of the same request type being stored in the staging buffer 112 that are subsequently added to the command queue 104 as part of the memory access request burst.

[0047] For further explanation, Figure 8 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0048] The method of Figure 8 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access

requests from the staging buffer 112 includes selecting 802 the memory access request 608 based on one or more of: a bank targeted by another memory access request, a rank targeted by a memory access request, or a memory subchannel targeted by another memory access request. For example, in some embodiments, memory access requests are added to the command queue 104 such that executed memory access requests alternately target different Dynamic Random Access Memory 106 ranks or banks (e.g., rank balancing, bank balancing). In other embodiments, memory access requests are added to the command queue 104 such that the executed memory access requests target Dynamic Random Access Memory 106 subchannels in a balanced approach. Accordingly, the memory access request 608 is selected based on a rank, bank, or subchannel targeted by a memory access request already added to the command queue 104 (e.g., a queued memory access command targeting a different rank, bank, or subchannel). The memory access request 608 is also selected based on a rank, bank, or subchannel targeted by a memory access request in the staging buffer 112 that is later added to the command queue 104 (e.g., a staged memory access command targeting a different rank, bank, or subchannel).

[0049] For further explanation, Figure 9 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0050] The method of Figure 9 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes selecting 902 the memory access request 608 based on a Dynamic Random Access Memory 106 page targeted by another memory access request. If an executed memory access request targets a page that is not currently open, overhead occurs in closing the currently open page and opening the targeted page. Executing memory access requests targeting a same (e.g. open) page reduces this overhead.

Accordingly, in some embodiments, the memory access request 608 is selected based on a page targeted by an already executed memory access request (e.g., an already open page). In some embodiments, the memory access request 608 is selected based on a page targeted by a memory access request stored in the command queue 104 to be executed prior to the selected memory access request 308 such that the targeted page will be open when the selected

memory access request 608 is executed. In some embodiments, the memory access request 608 is selected based on a page targeted by another memory access request stored in the staging buffer 112 that will be subsequently selected for movement to the command queue 104 such that the targeted page is open when the other memory access request is executed.

[0051] For further explanation, Figure 10 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0052] The method of Figure 10 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes selecting 1002 the memory access request 1002 based on a priority value. In some embodiments, the priority value is an explicit priority value assigned to memory access requests 608 (e.g., a priority tier). In other embodiments, the priority value is calculated based on an attribute of memory access requests, such as an age of the memory access requests (e.g., a time at which a given memory access request was generated or received by the memory management unit 102).

[0053] For further explanation, Figure 11 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0054] The method of Figure 11 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes identifying 1102, in the staging buffer 112, another memory access request associated with a first page miss. In other words, execution of the other memory access request will result in a page miss and corresponding computational overhead. For example, the page table 110 is accessed to determine that execution of the other memory access request will result in a page miss.

[0055] The method of Figure 11 further differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 also includes identifying 1104, in the command queue 104, a queued memory access request associated with a second page miss different from the first page miss. For example, the queued memory access request is identified as having a same request type and targeting a same Dynamic Random Access Memory bank as the other memory access request in the staging buffer, but will result in a different row page miss.

[0056] The method of Figure 11 further differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 also includes selecting 1106, in response to identifying the other memory access request and the queued memory access request, the memory access request 608. In other words, the memory access request 608 is preferentially selected over the other memory access request in the staging buffer 112 associated with the first page miss.

[0057] For further explanation, Figure 12 sets forth a flow chart illustrating an exemplary method for staging buffer arbitration according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0058] The method of Figure 12 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes identifying 1202, in the staging buffer 112, another memory access request associated with a first page conflict. In other words, execution of the other memory access request will result in a page conflict and corresponding computational overhead. For example, the page table 110 is accessed to determine that execution of the other memory access request will result in a page conflict.

[0059] The method of Figure 12 further differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 also includes identifying 1204, in the command queue 104, a queued memory access request associated with a second page conflict different from the first page conflict. For example, the queued memory access request is identified as having a same request type and targeting a same Dynamic Random Access Memory bank as

the other memory access request in the staging buffer 112, but will result in a different row page conflict.

[0060] The method of Figure 12 further differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 also includes selecting 1206, in response to identifying the other memory access request and the queued memory access request, the memory access request 608. In other words, the memory access request 608 is preferentially selected over the other memory access request in the staging buffer 112 associated with the first page conflict.

[0061] For further explanation, Figure 13 sets forth a flow chart illustrating an exemplary method for staging memory access requests according to embodiments of the present disclosure that includes storing 602 (e.g., by a memory management unit 102 of a processor 100) a plurality of memory access requests in a staging buffer 112; selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112; and moving 610 the memory access request 608 from the staging buffer 112 to a command queue 104.

[0062] The method of Figure 13 differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes identifying 1302, in the command queue 104, a page hit request. A page hit request includes a memory access request targeting a currently open page of dynamic random access memory 106. The method of Figure 13 further differs from Figure 6 in that selecting 606, based on one or more arbitration rules, a memory access request 608 of the plurality of memory access requests from the staging buffer 112 includes selecting 1304, based on the memory access request 608 being another page hit request, the memory access request 608. In other words, the memory access request 608 is preferentially selected for movement to the command queue 104 over other memory access requests that would result in a page miss. Thus, the staging buffer 112 will hold page conflict requests. In some implementations, the memory access request 608 is selected for movement to the command queue 104 such that the command queue 104 preferentially holds no more than one memory access request for each bank.

[0063] In view of the explanations set forth above, readers will recognize that the benefits of staging buffer arbitration according to embodiments of the present disclosure include:

- Improved performance of a computing system by relieving command queue pressure through the use of an additional staging buffer.
- Improved performance of a computing system by optimally selecting memory access requests for addition to the command queue to reduce computational overhead.

[0064] Exemplary embodiments of the present disclosure are described largely in the context of a fully functional computer system for staging buffer arbitration. Readers of skill in the art will recognize, however, that the present disclosure also can be embodied in a computer program product disposed upon computer readable storage media for use with any suitable data processing system. Such computer readable storage media can be any storage medium for machine-readable information, including magnetic media, optical media, or other suitable media. Examples of such media include magnetic disks in hard drives or diskettes, compact disks for optical drives, magnetic tape, and others as will occur to those of skill in the art. Persons skilled in the art will immediately recognize that any computer system having suitable programming means will be capable of executing the steps of the method of the disclosure as embodied in a computer program product. Persons skilled in the art will recognize also that, although some of the exemplary embodiments described in this specification are oriented to software installed and executing on computer hardware, nevertheless, alternative embodiments implemented as firmware or as hardware are well within the scope of the present disclosure.

[0065] The present disclosure can be a system, a method, and/or a computer program product. The computer program product can include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present disclosure.

[0066] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium can be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded

device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals *per se*, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0067] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network can include copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0068] Computer readable program instructions for carrying out operations of the present disclosure can be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions can execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer can be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection can be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) can execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present disclosure.

[0069] Aspects of the present disclosure are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0070] These computer readable program instructions can be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions can also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein includes an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0071] The computer readable program instructions can also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0072] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams can represent a module, segment, or portion of instructions, which includes one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block can occur out of the order noted in the figures. For example, two blocks shown in succession can, in fact, be executed substantially concurrently, or the blocks can sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special

purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0073] It will be understood from the foregoing description that modifications and changes can be made in various embodiments of the present disclosure. The descriptions in this specification are for purposes of illustration only and are not to be construed in a limiting sense. The scope of the present disclosure is limited only by the language of the following claims.

CLAIMS

What is claimed is:

1. A method of staging buffer arbitration, the method comprising:
storing a plurality of memory access requests in a staging buffer;
selecting a memory access request of the plurality of memory access requests from the staging buffer based on one or more arbitration rules; and
moving the memory access request from the staging buffer to a command queue.
2. The method of claim 1, wherein selecting the memory access request comprises selecting a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request.
3. The method of claim 1, wherein selecting the memory access request comprises selecting the memory access request based on one or more of: a bank targeted by another memory access request, a rank targeted by another memory access request, or a memory subchannel targeted by another memory access request.
4. The method of claim 1, wherein selecting the memory access request comprises selecting the memory access request based on a Dynamic Random Access Memory page targeted by another memory access request.
5. The method of claim 1, wherein selecting the memory access request comprises selecting the memory access request based on a priority value.
6. The method of claim 1, wherein selecting the memory access request comprises:
identifying, in the staging buffer, another memory access request associated with a first page miss;
identifying, in the command queue, a queued memory access request associated with a second page miss different from the first page miss; and
selecting the memory access request in response to identifying the other memory access request and the queued memory access request.
7. The method of claim 1, wherein selecting the memory access request comprises:
identifying, in the staging buffer, another memory access request associated with a first page conflict;
identifying, in the command queue, a queued memory access request associated with a second page conflict different from the first page conflict; and
selecting the memory access request in response to identifying the other memory access request and the queued memory access request.
8. The method of claim 1, wherein selecting the memory access request comprises:

- identifying, in the command queue, a page hit request; and
selecting the memory access request based on the memory access being another page hit request.
9. A memory management unit configured to perform steps comprising:
storing a plurality of memory access requests in a staging buffer;
selecting a memory access request of the plurality of memory access requests from the staging buffer based on one or more arbitration rules; and
moving the memory access request from the staging buffer to a command queue.
 10. The memory management unit of claim 9, wherein selecting the memory access request comprises selecting a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request.
 11. The memory management unit of claim 9, wherein selecting the memory access request comprises selecting the memory access request based on one or more of: a bank targeted by another memory access request, a rank targeted by another memory access request, or a memory subchannel targeted by another memory access request.
 12. The memory management unit of claim 9, wherein selecting the memory access request comprises selecting the memory access request based on a Dynamic Random Access Memory page targeted by another memory access request.
 13. The memory management unit of claim 9, wherein selecting the memory access request comprises selecting the memory access request based on a priority value.
 14. The memory management unit of claim 9, wherein selecting the memory access request comprises:
identifying, in the staging buffer, another memory access request associated with a first page miss;
identifying, in the command queue, a queued memory access request associated with a second page miss different from the first page miss; and
selecting the memory access request in response to identifying the other memory access request and the queued memory access request.
 15. The memory management unit of claim 9, wherein selecting the memory access request comprises:
identifying, in the staging buffer, another memory access request associated with a first page conflict;
identifying, in the command queue, a queued memory access request associated with a second page conflict different from the first page conflict; and

- selecting the memory access request in response to identifying the other memory access request and the queued memory access request.
16. The memory management unit of claim 9, wherein selecting the memory access request comprises:
- identifying, in the command queue, a page hit request; and
- selecting the memory access request based on the memory access being another page hit request.
17. An apparatus comprising a processor, the processor comprising a memory management unit configured to:
- store a plurality of memory access requests in a staging buffer;
- select a memory access request from the plurality of memory access request from the staging buffer based on one or more arbitration rules; and
- move the memory access request from the staging buffer to a command queue.
18. The apparatus of claim 17, wherein selecting the memory access request comprises selecting a memory access request burst of a same request type, wherein the memory access request burst includes the memory access request.
19. The apparatus of claim 17, wherein selecting the memory access request comprises selecting the memory access request based on one or more of: a bank targeted by another memory access request, a rank targeted by another memory access request, or a memory subchannel targeted by another memory access request.
20. The apparatus of claim 17, wherein selecting the memory access request comprises selecting the memory access request based on a Dynamic Random Access Memory page targeted by another memory access request.
21. The apparatus of claim 17, wherein selecting the memory access request comprises selecting the memory access request based on a priority value.
22. The apparatus of claim 17, wherein selecting the memory access request comprises:
- identifying, in the staging buffer, another memory access request associated with a first page miss;
- identifying, in the command queue, a queued memory access request associated with a second page miss different from the first page miss; and
- selecting the memory access request in response to identifying the other memory access request and the queued memory access request.
23. The apparatus of claim 17, wherein selecting the memory access request comprises:

identifying, in the staging buffer, another memory access request associated with a first page conflict;
identifying, in the command queue, a queued memory access request associated with a second page conflict different from the first page conflict; and
selecting the memory access request in response to identifying the other memory access request and the queued memory access request.

24. The apparatus of claim 17, wherein selecting the memory access request comprises:
identifying, in the command queue, a page hit request; and
selecting the memory access request based on the memory access being another page hit request.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2020/066029

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F13/16

ADD.

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

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2018/173649 A1 (KYRYCHYNSKYI ROSTYSLAV [CA] ET AL) 21 June 2018 (2018-06-21) paragraphs [0018], [0019] paragraph [0027] - paragraph [0029] paragraph [0035] - paragraph [0039] figures 1-3 -----	1-24
A	US 2018/018105 A1 (MAGRO JAMES R [US] ET AL) 18 January 2018 (2018-01-18) paragraph [0037] - paragraph [0047] figures 5,6 -----	1-24
A	US 2011/179240 A1 (SUKONIK VITALY [IL] ET AL) 21 July 2011 (2011-07-21) figures 1,3 paragraph [0098] - paragraph [0119] ----- - / - -	1-24



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

10 March 2021

Date of mailing of the international search report

22/03/2021

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

van der Meulen, E

INTERNATIONAL SEARCH REPORT

International application No PCT/US2020/066029

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 8 271 746 B1 (HUTSELL BRIAN D [US] ET AL) 18 September 2012 (2012-09-18) the whole document -----	1-24

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2020/066029

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2018173649	A1	21-06-2018	NONE

US 2018018105	A1	18-01-2018	CN 109690508 A 26-04-2019
			KR 20190019203 A 26-02-2019
			US 2018018105 A1 18-01-2018
			WO 2018013159 A1 18-01-2018

US 2011179240	A1	21-07-2011	CN 102129412 A 20-07-2011
			TW 201145023 A 16-12-2011
			US 2011179240 A1 21-07-2011
			US 2014115254 A1 24-04-2014

US 8271746	B1	18-09-2012	NONE
