



- (51) **International Patent Classification:**
G06F 13/16 (2006.01)
- (21) **International Application Number:**
PCT/US2015/024414
- (22) **International Filing Date:**
6 April 2015 (06.04.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
14/252,673 14 April 2014 (14.04.2014) US
- (71) **Applicant:** MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, WA 98052-6399 (US).
- (72) **Inventors:** QUACH, Nhon, Toai; c/o Microsoft Technology Licensing, LLC, LCA - International Patents (8/1172), One Microsoft Way, Redmond, WA 98052-6399 (US).
CARRIE, Susan; c/o Microsoft Technology Licensing, LLC, LCA - International Patents (8/1172), One Microsoft

Way, Redmond, WA 98052-6399 (US). **ANDREWS, Jeffrey**; c/o Microsoft Technology Licensing, LLC, LCA - International Patents (8/1172), One Microsoft Way, Redmond, WA 98052-6399 (US). **SELL, John**; c/o Microsoft Technology Licensing, LLC, LCA - International Patents (8/1172), One Microsoft Way, Redmond, WA 98052-6399 (US). **PO, Kevin**; c/o Microsoft Technology Licensing, LLC, LCA - International Patents (8/1172), One Microsoft Way, Redmond, WA 98052-6399 (US).

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

[Continued on next page]

(54) **Title:** FINE-GRAINED BANDWIDTH PROVISIONING IN A MEMORY CONTROLLER

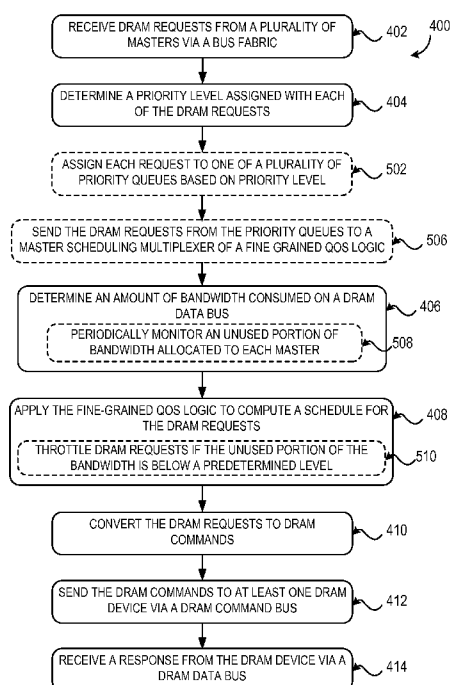


FIG. 4

(57) **Abstract:** Systems and methods for applying a fine-grained QoS logic are provided. The system may include a memory controller, the memory controller configured to receive memory access requests from a plurality of masters via a bus fabric. The memory controller determines the priority class of each of the plurality of masters, and further determines the amount of memory data bus bandwidth consumed by each master on the memory data bus. Based on the priority class assigned to each of the masters and the amount of memory data bus bandwidth consumed by each master, the memory controller applies a fine-grained QoS logic to compute a schedule for the memory requests. Based on this schedule, the memory controller converts the memory requests to memory commands, sends the memory commands to a memory device via a memory command bus, and receives a response from the memory device via a memory data bus.



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

Declarations under Rule 4.17:

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

Published:

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

FINE-GRAINED BANDWIDTH PROVISIONING IN A MEMORY CONTROLLER

BACKGROUND

5 [0001] Dynamic random access memory (DRAM) devices are found in a wide range of electronic devices, including computers, smart phones, tablets, game consoles, and televisions. DRAM devices allow the applications running on these electronic devices to temporarily store instructions and data from the non-volatile storage in the device (e.g. FLASH memory or hard drives) before the instructions are executed by the central
10 processing unit (CPU) or other fixed hardware units of the device.

[0002] For efficiency reasons, DRAM devices are shared by many different master devices, or intellectual-property (IP) logic blocks that can generate read or write traffic to the DRAM memory devices, in an electronic system. Each master device issues memory requests to the DRAM devices, often via a bus fabric and a memory controller. The
15 memory controller is responsible for scheduling these memory requests and determining how the requests are sent to the DRAM devices. The access latency and, therefore, the effective utilization bandwidth of the DRAM devices depend on how efficiently and effectively the memory controller is able to schedule the requests. When the scheduling is inefficient, excessive bandwidth variations can cause failures in hardware and software
20 applications of the system, such as a frame rate drop in a video game or sound popping due to missing audio frames. A common approach to prevent these failures is to utilize a quality of service (QoS) logic determining how the requests are scheduled and to perform extensive testing to ensure that the application will not fail under as many known bandwidth variation scenarios as possible.

25 [0003] In the prior art, QoS logic solutions primarily relied on two mechanisms: prioritization and rate limiting. A prioritization mechanism favors requests from masters that have been assigned higher priority by the system over requests from masters that have been assigned a lower priority. To prevent the starvation of low priority masters, however, a rate limiting mechanism is used to prevent high priority masters from exceeding an
30 amount of bandwidth allocated to them by the system.

[0004] However, these prior QoS logic solutions still allowed for wide variations in the effective bandwidth of the DRAM devices, as perceived by the master devices. As such, software applications developed for electronic systems utilizing prior QoS logic solutions were limited to a specific type of DRAM for which the bandwidth variations has

been extensively tested. If a different type of DRAM device, either faster or slower, was used in the system, the QoS logic solution would undergo additional bandwidth variations and the software application would fail to run properly. In the example of game consoles, which typically have long product lifecycles, this drawback prevents console makers from switching to faster, cheaper types of DRAM during a latter period of the lifecycle of a game console for fear that previously developed games will not run properly on the updated hardware.

SUMMARY

[0005] Systems and methods for applying a fine-grained QoS logic are provided.

The system may include a memory controller, the memory controller configured to receive memory access requests from a plurality of masters via a bus fabric. The memory controller determines the priority class of each of the plurality of masters, and further determines the amount of data bus bandwidth consumed by each master on the data bus. Based on the priority class assigned to each of the masters and the amount of data bus bandwidth consumed by each master, the memory controller applies a fine-grained QoS logic to compute a schedule for the memory access requests. Based on this schedule, the memory controller converts the memory access requests to memory access commands, sends these memory access commands to a memory access device via a memory command bus, and receives a response from the memory device via a memory data bus.

[0006] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter. Furthermore, the claimed subject matter is not limited to implementations that solve any or all disadvantages noted in any part of this disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a schematic diagram of a system configured to execute a fine-grained QoS logic to provision the bandwidth of a DRAM data bus, according to one embodiment.

[0008] FIG. 2 is a schematic diagram of the fine-grained QoS logic of the system of FIG. 1, which may be configured to compute an amount of bandwidth consumed by a DRAM request on a DRAM data bus.

[0009] FIG. 3 is a schematic diagram of a second embodiment of a system configured to execute a fine-grained QoS logic to provision the bandwidth of a DRAM

data bus, in which the masters, bus fabric, and DRAM memory controller are formed as a system-on-a-chip.

[0010] FIG. 4 is a flowchart of a method for performing fine-grained bandwidth provisioning in a DRAM memory controller, according to one embodiment.

5 [0011] FIG. 5 is a flowchart illustrating substeps by which steps of the method of FIG. 4 can be accomplished.

[0012] FIG. 6 is a flowchart illustrating substeps by a step of the method of FIG. 4 can be accomplished.

[0013] FIG. 7 is a schematic view showing additional components of the system.

10

DETAILED DESCRIPTION

[0014] Systems and methods for controlling access to memory, such as dynamic random access memory (DRAM), are disclosed herein, to address the challenges described above.

15 [0015] In modern electronic systems, it is typical for the CPU and other hardware devices, often referred to as master devices, to issue requests to the DRAM of the system via a memory controller and various buses. The memory controller may use various queues and a multiplexer to schedule the requests issued from the master devices and may apply a quality of service (QoS) logic in order to determine the most efficient manner to execute the requests.

20 [0016] FIG. 1 depicts an embodiment of a system 10 including a memory controller 12 that is configured to receive dynamic random access memory (DRAM) requests 14 from a plurality of masters 16 via a bus fabric 18, determine a priority level assigned with each of the plurality of masters 16, determine an amount of DRAM data bus bandwidth consumed 48 by each master on a DRAM data bus 30, apply a fine-grained
25 QoS logic 26 to compute a schedule for the DRAM requests 14 based on a priority level assigned to each master 16 and an amount of bandwidth consumed 48 by each master 16, and, based on the schedule, convert the DRAM request 14 to DRAM commands 42, send the DRAM commands 42 to at least one DRAM device 32 via a DRAM command bus 34, and receive a response from the DRAM device 32 via a DRAM data bus 30. The masters
30 16 may be, for example, the CPU and GPU of a video game console and may be requesting to read or write, on the DRAM devices 32, data related to a video game currently being played by a user of the console.

[0017] When requests 14 are issued by a master 16, the system 10 is configured to transmit the requests through the bus fabric 18. The bus fabric is shown in FIG. 1 as a

single fabric arranged between the master 16 and the memory controller 12. However, the bus fabric 18 may also comprise as many bus segments as needed to accommodate all the traffic patterns from the masters 16. FIG. 3 illustrates a second embodiment of the system 10, in which the bus fabric is divided into a plurality of bus segments, including segments with rate limiters 60 and segments without rate limiters 62. Turning back to FIG. 1, the bus fabric 18 may be configured to apply a QoS Priority Assignment Logic 20 to assign a priority level to each of the requests 14 based on a priority level assigned to each of the masters 16. Priority levels may be assigned to each of the masters 16 by software applications running on the system 10. Such an assignment may occur, for example, once per boot up of system 10, on an application border when an application is run on system 10 or on each boot up of a virtualized environment on system 10. The masters 16 may include real-time master devices that may be sensitive to bandwidth variation and that may cause a degraded user experience if bandwidth variation becomes excessive. Such real time master devices may be assigned higher priority levels by the by software applications running on the system 10.

[0018] After assigning a priority level to each of the requests 14, the bus fabric 18 may be configured to send the request to the memory controller 12. The memory controller 12 may be further configured to assign each DRAM request 14 to one of a plurality of priority queues 22 based on the priority level, and schedule the DRAM requests 14 in each of the plurality of priority queues 22.

[0019] To understand the scheduling parameters that may be used by the priority queues 22 to schedule the requests 14, it is helpful to briefly turn to a discussion of the structure and operation of the DRAM devices 32. The DRAM devices 32 may be configured as arrays of “pages,” each page containing a number of memory cells (for example, between 8 kilobits and 16 kilobits per page). The pages may be arranged in rows and columns, and further grouped into banks, each bank containing tens of thousands of pages. A page in the DRAM devices 32 may be either open or closed depending on the whether or not it has been accessed by one of the requests 14. System 10 may be configured such that all pages in all banks of the DRAM devices 32 are closed when system 10 is booted up. When one of the DRAM requests 14 is issued requesting access to a closed page in a bank of the DRAM device 32, the request 14 will initially result in a “page miss” and the memory controller 12 will have to perform a “page open operation” to open the targeted page. While the page open operation is occurring the memory controller 12 will not be able to access the content of any other pages in that bank, since the memory

controller 12 is typically configured to allow only one open page in each bank at any one time. As a result, during the page open operation, the memory controller 12 may be configured to seek requests 14 targeting pages in other banks.

[0020] When a page is open, subsequent requests 14 to access the same page will initially result in a “page hit” within the memory controller 12 and thus can be completed without performing a page open operation. Thus, memory controller 12 may be configured to perform “same page gathering,” wherein requests 14 requesting access to the same page are grouped together to reduce the need for page open operations.

[0021] When the a subsequent request targets a different page in the same bank as the last issued request (i.e. the last request completed prior to current request), a so-called “page conflict” occurs, and the memory controller 12 must close the currently open page with a “page close operation” before opening the page targeted by the request 14 with a page open operation. Both page open operations and page close operations consume bandwidth on the DRAM data bus 30. Bandwidth consumption on the DRAM data bus 30 may be measured as the number of bus cycles required to complete a given request 14. The number of bus cycles, and therefore the amount of bandwidth consumed 48, on the DRAM data bus 30 per operation is determined by the timing parameters of the specific type of memory used for DRAM devices 32. However, irrespective of the type of memory used, bandwidth consumption 48 on the DRAM data bus 30 varies depending on whether the request 14 has a page hit, a page miss, or a page conflict.

[0022] In addition, DRAM requests 14 may be either read requests or write requests. If a request 14 is a read request, the memory controller 12 is configured to perform a “read operation,” sending a “read command” on the DRAM command bus 34 and receiving “read data” on the DRAM data bus 30 after a predetermined number of clock cycles. If the request 14 is a write request, the memory controller 12 is configured to perform a write operation, sending a write command on the DRAM command bus 34 and sending the write data on the data bus 30 after waiting for a predetermined number of clock cycles. However, when request 14 is a write request and the last issued request was a read request, the memory controller 12 is configured to perform a “read to write bus turnaround” to avoid a data hazard (i.e., a risk of mishandling the data) resulting from a conflict in the use of the data bus 30. For example, if read data is still returning on the DRAM data bus 30 when a write operation begins, a bus collision between the outgoing write data and the incoming read data could occur. To avoid such occurrences during a write operation, prior to beginning the write operation, the memory controller 12 may be

configured to confirm that no previous read data is being returned on the DRAM data bus 30. When there is a conflict, the memory controller 12 may be configured to perform a read to write bus turnaround, in which the write command is delayed to allow the read data to be returned on the data bus 30 before sending the write data to the DRAM device 32.

5 Similarly, when the request 14 is a read request and the last issued request was a write request, the memory controller 12 is configured to perform a “write to read bus turnaround,” delaying the read command to allow the write data to be written into the memory of the DRAM device 32 before the read data are returned from the DRAM device 32 to the memory controller 12 via the DRAM data bus 30. Both read to write bus
10 turnaround and write to read bus turnaround operations consume a number of bus cycles, and therefore an amount of bandwidth, on the DRAM data bus 30. Thus, the amount of bandwidth consumed 48 on the DRAM data bus 30 required to complete a request 14 may depend on both the request type (i.e. read or write) of the current request 14 being issued by the memory controller 12 to the DRAM devices 32 and the request type of last issued
15 request 14. To minimize this bandwidth consumption, the memory controller 12 may be configured to perform “read/write grouping,” in which requests 14 are grouped based on the request type to reduce the need for read to write bus turnaround and write to read bus turnaround operations.

[0023] Furthermore, if a request 14 is a read request, the master 16 issuing the
20 request 14 may be requesting data that has not yet been written on the DRAM devices 32. To avoid issuing the request 14 before the data has been written, the page address of the request 14 must be compared to the page addresses of all write requests 14 in all priority queues 22. This process is referred to as address conflict resolution, and may be used to ensure data consistency. If there is a match, the read request 14 in question must wait until
25 all earlier write requests 14 are completed.

[0024] Table 1, shown below, describes examples of the bandwidth costs for some common cases of a given request 14 based on the location of the page requested, request type and the state of the DRAM device 32 at the time when the request 14 is made (i.e. request type previously completed).

Previous DRAM Request	Current DRAM Request					
	Same page		Different page			
	Read	Write	Same Bank		Different Bank	
			Read	Write	Read	Write
Read	BL/2	BL/2 + tR2W	tPRE + tACT + BL/2	tR2W + tPRE + tACT + BL/2	<ul style="list-style-type: none"> • If a different page from the targeted one is open: tPRE + tACT + BL/2 • If the targeted page is closed: tACT + BL/2 • If the targeted page is open: BL/2 	<ul style="list-style-type: none"> • If a different page from the targeted one is open: tR2W + tPRE + tACT + BL/2 • If the targeted page is closed: tACT + BL/s • If the targeted page is open: BL/2
Write	BL/2 + tW2R	BL/2	tW2R + tPRE + tACT + BL/2	tPRE + tACT + BL/2	<ul style="list-style-type: none"> • If a different page from the targeted one is open: tW2R + tPRE + tACT + BL/2. • If the targeted page is closed: tACT + BL/2 • If the targeted page is open: BL/2 	<ul style="list-style-type: none"> • If the different page from the targeted one is open: tPRE + tACT + BL/2 • If the targeted page is closed: tACT + BL/2 • If the targeted page is open: BL/2

Table 1: Bandwidth Cost per Request in terms of Bus Cycles

[0025] In Table 1 above, the variable BL indicates the Burst Length of the DRAM device 32; BL/2 is the duration of time the DRAM data bus 30 will be busy carrying the data read from or written to the DRAM device 32 for any double data rate memory devices such as DDRn, low-power DDRn (LPDDRn), or graphics DDRn (GDDRn); the variable tACT is the bandwidth cost of opening a page, which is the time it takes to activate a page as described above; tPRE is the bandwidth cost of closing a page, which is the time it takes to close an open page as described above; tW2R is the cost, expressed as the number of cycles of the DRAM data bus 30 required to turn the DRAM data bus 30 for a read operation following a write operation by performing a write to read bus turnaround as described above; and tR2W is the cost, expressed as the number of cycles of the DRAM data bus 30 required to turn the DRAM data bus 30 for a write operation following a read by performing a read to write bus operation as described above. It is important to note that while different DRAM memory devices (e.g., DDR3, LPDDR3, or GDDR3) may use different names for these parameters, the concepts described in the present disclosure apply to any of the DRAM devices described above. Note that the refresh operations, which will be elaborated on in the following paragraph, were not included in Table 1 as they are considered the overhead cost of the DRAM device 32, not those of the master 16.

[0026] Table 1 represents one possible embodiment of the system 10. The system 10 is not limited to the parameters listed in table 1 and, depending on the final objectives of the QoS scheme, other tables may be constructed according to the principles described herein to achieve the desired results.

[0027] The memory controller 12 may be configured to schedule the DRAM requests 14 in the priority queues 22 in accordance with the above methods so as to minimize the bandwidth consumption on the DRAM data bus 30. Each priority queue 22 may be composed of a plurality of queues 46. The number and size of the priority queues 22 for each priority level, as well as the number of queues 46 within each of the priority queues 22, may vary depending on the implementation. In FIG. 1, N master devices 16 and M priority queues 22 are depicted. The queues 46 within each of priority queues 22 for a given priority level may comprise at least one read queue and one write queue. The memory controller 12 may be configured to schedule the requests 14 within each of the plurality of priority queues 22. The schedule may be master dependent and based on parameters such as address conflict resolution, same page gathering, read/write grouping, and the priority level assigned to each of the requests 14.

[0028] The memory controller 12 may be further configured to send the DRAM requests 14 to a master scheduling multiplexer 24, periodically monitor an unused portion of bandwidth allocated to each master 16, and throttle DRAM requests 14 at the memory controller 12 if an unused portion of the bandwidth drops below a predetermined level.

5 The memory controller 12 may be configured to send the DRAM requests 14 to the master scheduling multiplexer 24 after selecting the best candidate requests 14 from the priority queues 22, based on the scheduling parameters described above.

[0029] The master scheduling multiplexer 24 may be configured to apply the fine-grained QoS logic 26, including a fine-grained QoS bandwidth provisioning logic 28
10 configured to allocate a predetermined amount of bandwidth in the DRAM data bus 30 to each of the plurality of masters 16. The fine-grained QoS logic 26 may be configured to divide the bandwidth of the DRAM device 32 into two categories, useful QoS bandwidth and overhead bandwidth. Overhead bandwidth may be bandwidth consumed by the memory controller 12 to ensure correct operation of the DRAM devices 32 and may
15 include bandwidth consumed for refresh operations. During a refresh operation, the memory controller 12 is configured to suspend all read operations and write operations to all banks of the DRAM device 32, close all open pages, and issue refresh commands to refresh the content of all the pages in the DRAM devices 32. After the refresh operation is complete, the memory controller 12 is configured to resume read and write operations.

20 Useful QoS bandwidth may include the bandwidth consumed 48 by the masters 16 to perform operations such as read and write operations on the DRAM devices 32. The memory controller 12 may be configured to assign a predetermined amount of the useful QoS bandwidth to the masters 16 through the fine-grained QoS bandwidth provisioning logic 28. The amount of bandwidth allocated to each of the masters 16 may be determined
25 by software applications running on the system 10. Such an assignment may occur, for example, once per boot up of the system 10, on an application border when an application is run on system 10, or on each boot up of a virtualized environment on system 10.

[0030] The memory controller 12 may be configured to determine the unused portion of bandwidth for each master 16 by computing the bandwidth consumed 48 by
30 each master 16 on the DRAM data bus 30 for each request 14 and subtracting the bandwidth consumed 48 from the amount of useful QoS bandwidth allocated to each master 16 by the fine-grained QoS bandwidth provisioning logic 28. The memory controller 12 may be configured to measure the bandwidth consumed 48 by a request 14 in terms of the number of bus cycles of the DRAM data bus 30 that are required to complete

the request 14. The number of bus cycles required to complete a request 14 may vary depending on various characteristics of the request 14, such as request type (i.e. read or write), whether the page targeted by the request 14 in the DRAM device 32 is currently open or closed, and whether or not the page targeted by request 14 is located in the same bank as the page targeted by the last issued request.

[0031] The fine-grained QoS logic 26 of the memory controller 12 may be further configured to include a QoS credit accumulator 36. The QoS credit accumulator 36 may be configured to incrementally accumulate credits for each of the plurality of masters 16. The rate of accumulation may correspond to the bandwidth allocated to the master by the fine-grained QoS bandwidth provisioning logic 28, and the total credits for each master 16 may not exceed a predetermined maximum QoS credit value. For example, if a given master 16 is allocated 1 gigabyte per second of bandwidth in the DRAM data bus 30, and the DRAM data bus 30 is running at a clock frequency of 200 MHz (200,000,000 cycles per second), then the QoS credit accumulator 36 would accumulate 5 bytes worth of credit per cycle designated for the master 16 in question. If the master 16 were to issue a request 14, the bandwidth consumed 48 by the request 14 would be subtracted from the credits in the QoS credit accumulator 36 corresponding to the master 16. If the master 16 issues no requests 14, the QoS credit accumulator 36 would continue to accumulate credit at a rate of 5 bytes per cycle until the total credits for the master 16 reached the maximum QoS credit value. The system 10 may be configured to allow a software application running on the system 10 to determine the maximum QoS credit value.

[0032] The fine-grained QoS logic 26 may be configured to further include a fine-grained QoS charge accounting logic 38, which may be configured to calculate the bandwidth 48 on the DRAM data bus consumed by each DRAM memory request, and compute a credit balance for each master by subtracting the calculated bandwidth from the credits in the QoS credit accumulator corresponding to each master. FIG. 2 shows a schematic diagram of the processes that may be included in the operation of the fine-grained QoS charge accounting logic 38. As shown, the fine-grained QoS charge accounting logic 38 may be configured to receive, from the priority queues 22, current request information 50 regarding the currently queued requests. The current request information 50 may include the request type (i.e. read or write) of the requests 14, as well as the page address and the bank of the page targeted by the requests 14 currently being sent from the priority queues 22 to the master scheduling multiplexer 24. In addition, the fine-grained QoS charge accounting logic 38 may be further configured to receive last

issued request information 52 from a DRAM protocol engine 44. The last issued request information 54 may include the request type (i.e. read or write) of the request 14, as well as the page address and the bank of the page targeted by the request 14 that was last sent to the DRAM devices 32. The last issued request information 54 may further include whether the request 14 last issued to the DRAM devices 32 had a page hit, page miss, or page conflict. The fine-grained QoS charge accounting logic 38 may be further configured to use one or both of the current request information 50 and last issued request information 54 to calculate the bandwidth consumed 48 on the DRAM data bus 30 by a request 14 and subtract that amount from the credits in the QoS credit accumulator 36 corresponding to the master 16 that issued the request 14. The fine-grained QoS charge accounting logic 38 may also be configured to use global page status information 58 received from a DRAM protocol engine 44. The global page status information 58 may include the current status of all pages in the DRAM devices 32. The fine-grained QoS charge accounting logic 38 may use global page status information 58 to determine which requests 14 would have a page hit, a page miss, or a page conflict if issued to the DRAM devices 32. The fine-grained QoS charge accounting logic 38 may be further configured to consider the global overhead bandwidth consumed by refresh 56 and the credit gain per clock cycle for refresh 54 before subtracting the bandwidth consumed 48 by the request 14 from the total credits in the QoS credit accumulator 36. Because the system 10 may be configured to allocate only useful QoS bandwidth to masters 16, the QoS charge accounting logic 38 may be configured to provide a credit gain per clock cycle to compensate for the bandwidth consumed by refresh operations, or other forms of overhead bandwidth consumption. After the QoS charge accounting logic 38 has subtracted the bandwidth consumed 48 by the last issued request 14, the fine-grained QoS credit accumulator 36 may be further configured to send the current balance in the QoS credit accumulator 36 for each of the masters 16 to a QoS scheduling logic 40, discussed below.

[0033] Turning back to FIG. 1, the fine-grained QoS logic 26 may be further configured to include a fine-grained QoS scheduling logic 40, configured to schedule the DRAM requests 14 based on the priority class of each request, and to throttle the DRAM requests 14 if the credit balance in the QoS credit accumulator 36 for the master 16 issuing the DRAM request 14 is below a predetermined minimum QoS credit value. The fine-grained QoS scheduling logic 40 may be configured to receive a current credit balance for each of the masters 16 corresponding to each of the queued requests 14 from the fine-grained QoS charge accounting logic 38. The fine-grained QoS scheduling logic 40 may

be further configured to select the request 14 with the highest priority level and a credit balance above the predetermined minimum QoS credit value. If a request 14 has a credit balance below the minimum QoS credit value, the fine-grained QoS scheduling logic 40 may be configured to throttle the request 14, irrespective of the priority level of the request 14. If multiple requests 14 have the same priority level and credit balances above the minimum QoS credit value, scheduling may be performed in a round-robin manner, or according to another suitable default scheduling algorithm. If all available requests 14 have credit balances below the minimum QoS credit value, the request 14 with the highest priority level may be selected. The system 10 may be configured to allow a software application running on the system 10 to determine the minimum QoS credit value.

[0034] The memory controller 12 may be further configured to include a DRAM protocol engine 44 configured to receive DRAM requests 14 from the master scheduling multiplexer 24, send DRAM requests 14 to the DRAM devices 32 via a DRAM command bus 34. The DRAM protocol engine 44 may be configured to receive the DRAM request 14 that have been selected by the master scheduling multiplexer 24 based on the schedule created by the fine-grained QoS logic 26. The memory controller 12 may be further configured to convert the DRAM requests 14 to DRAM commands 42 prior to sending the DRAM requests 14 to the DRAM protocol engine 44. The DRAM protocol engine 44 may be further configured to receive a response from the DRAM devices 32 and send information about the last issued request 14 to the fine-grained QoS charge accounting logic 38. Optionally, the protocol engine 44 may be configured to receive a response from the DRAM devices 32, which may be in the form of an interrupt to inform the memory controller 12 that the DRAM request 14 was completed successfully. The protocol engine 44 may be further configured to send the last issued request information 54, as discussed shown in FIG. 2, to the fine-grained QoS charge accounting logic 38. This information may include the request type (i.e. read or write) of the request 14, the page address and bank of the page targeted by the request 14, and whether the request 14 had a page hit, a page miss, or a page conflict.

[0035] FIG. 3 illustrates a second embodiment of a system 100 configured to execute a fine-grained QoS logic to provision the bandwidth of a DRAM data bus. System 100 includes a system-on-a-chip (SoC) 101, on which are formed the masters 16, bus fabric 18, and DRAM memory controller 12, which in turn communicates with DRAM devices 32. The DRAM devices 32 may be formed on the SoC 101 or as independent devices. The various components formed on the system on chip function similarly as

described above, and will not be re-described for the sake of brevity. As described above, in this embodiment the bus fabric is divided into a plurality of bus segments, including segments with rate limiters 60 and segments without rate limiters 62. It will be appreciated that the system on chip 101 utilizes the rate limiters 62 to ensure a minimum fixed amount of bandwidth on the bus fabric is reserved for each master, while ensuring that each master does not exceed this fixed amount by limiting the maximum bandwidth consumed by each master. This avoids scheduling conflicts and promotes efficient access to the DRAM devices from the rate limited masters. If a new master is added to such a system that passes only through a rate unlimited segment of the bus fabric, such as bus segment M, it can consume too much bandwidth on the bus fabric and or degraded performance of the other masters which were designed with the expectation of having a fixed amount of reserved bandwidth. Thus, it will be appreciated that, while the rate limiters enforce some scheduling control, since no rate limiters are configured on bus segments M, DRAM requests passing through the rate-unlimited bus segments M in particular can benefit from the priority class and fine grained QoS scheduling applied by the scheduling multiplexer 24, which in turn will prevent inadvertent consumption of the reserved bandwidth by masters connected through bus segment M. This may improve the performance of all masters.

[0036] Turning now to FIG. 4, a flowchart of a method 400 for performing fine-grained bandwidth provisioning in a memory controller is depicted. The methods described hereafter may be implemented on the hardware of system 10, described above with references to FIGS. 1 and 2, or on other suitable hardware. It will be appreciated that suitable hardware on which the methods described herein may be performed include laptop and desktop personal computers, smartphones, video games consoles, tablet computing devices, smart televisions, etc. In some implementations of the method, the steps in FIG. 4 in dashed lines may be omitted, while in other implementations, the steps in dashed lines are included.

[0037] The method 400 may include at 402 receiving dynamic random access memory (DRAM) requests from a plurality of masters via a bus fabric. At 404, the method 400 may include determining a priority level assigned with each of the DRAM requests.

[0038] At 502, the method 400 may include assigning each DRAM request to one of a plurality of priority queues based on the priority level.

[0039] At 506, the method 400 may include sending the DRAM requests to a master scheduling multiplexer.

[0040] At 406, the method 400 may include determining an amount of DRAM data bus bandwidth consumed by each master on a DRAM data bus. At 508, it is shown that this may be accomplished at least in part by periodically monitoring an unused portion of bandwidth allocated to each master.

5 [0041] At 408, the method 400 may include applying a fine-grained quality of service (QoS) logic to compute a schedule for the DRAM requests based on a priority class assigned to each master and an amount of DRAM data bus bandwidth consumed by each master. At 510, it is shown that the schedule may be computed to throttle the DRAM requests at the memory controller if the unused portion of the bandwidth drops below a
10 predetermined level.

[0042] At 410, based upon the schedule created at step 408, the method 400 may include converting the DRAM requests to DRAM commands. At 412, the method 400 may include sending the DRAM commands to at least one DRAM device via a DRAM command bus. At 414, the method 400 may include receiving a response from the DRAM
15 device via a DRAM data bus.

[0043] FIG. 5 shows a flowchart illustrates example substeps by which steps 508, 510 and 412 of method 400 of FIG. 4 can be accomplished. Step 508 of method 400 may be accomplished by performing substeps 602-608. At 602, the method may include applying a fine-grained QoS bandwidth provisioning logic to allocate a predetermined
20 amount of bandwidth in the DRAM data bus to each of the plurality of masters. At 604, the method may include applying a QoS credit accumulator to incrementally accumulate credits for each of the plurality of masters, the rate of accumulation corresponding to the bandwidth allocated to the master by the fine-grained QoS bandwidth provisioning logic, and the total credits for each master not exceeding a predetermined maximum QoS credit
25 value. At 606, the method may include applying a fine-grained QoS charge accounting logic to calculate a bandwidth on the DRAM data bus consumed by each DRAM memory request. At 608, the method may include computing a credit balance for each master by subtracting the calculated bandwidth 48 from the credits in the QoS credit accumulator corresponding to each master.

30 [0044] Step 510 of method 400 may be accomplished by substep 610. As shown at substep 610, applying a fine-grained scheduling logic and throttling DRAM requests at the memory controller if an unused portion of the bandwidth drops below a predetermined level are achieved by applying a fine-grained QoS scheduling logic to schedule the DRAM requests based on the priority class of each request, and to compute the schedule so as to

throttle the DRAM requests if the credit balance in the QoS credit accumulator for the master issuing the DRAM request is below a predetermined minimum QoS credit value.

[0045] Step 412 of method 400 may be accomplished by substeps 612-614. At 612, the method may include sending the DRAM requests from the master scheduling multiplexer to the DRAM protocol engine. At 614, the method may include sending the DRAM requests from the DRAM protocol engine to at least one of the plurality of DRAM devices.

[0046] FIG. 6 shows a flowchart that illustrates example substeps 650, 652 by which step 406 of method 400 of FIG. 4 can be accomplished. At 650, the method may include receiving information from the DRAM protocol engine about the last issued request. At 652, the method 601 may include calculating a bandwidth on the DRAM data bus consumed by each DRAM memory request based on information about a last issued request.

[0047] Briefly turning back to FIG. 1, system 10 may also be used to model the performance of one or more software programs that were originally designed to run on a different system with less total memory bandwidth than system 10. To do so, a portion of the requests 14 from the masters 16 may be tagged as resulting from the execution of these programs. The memory controller 12 can calculate the bandwidth and QoS that the masters 16 would receive if the programs were running on the system for which they were originally designed. Using this calculation, the memory controller 12 can then allocate bandwidth and QoS to the requests 14 on system 10. Thus, the computing device may be further configured to tag a portion of the memory requests from the plurality of masters as resulting from an execution of one or more software programs originally designed to run on a system with less total memory bandwidth than the computing device, and to calculate a bandwidth and a quality-of-service (QoS) the plurality of masters would receive if the one or more software programs were running on the system with less total memory bandwidth. Further, the computing device may be configured to allocate the bandwidth and the QoS to the memory requests from the plurality of masters; and, as a result, and model a performance of the one or more software programs on the system with less total memory bandwidth.

[0048] The systems and methods described above have various useful industrial applications. For example, a first application is to allow a faster memory type to yield the same bandwidth variation as a slower memory type. For example, the system may be a video game console designed using, for example, Double Data Rate 3 (DDR3) memory

running at a clock frequency of 2133Mbps. During the lifecycle of the video game console, it may be advantageous (e.g. for performance or cost reasons) to change the memory to, for example, Double Data Rate 4 (DDR4) running at a clock frequency of 2400Mbps. If such a change is made during the lifecycle of a video game console, there is a risk that software designed for use with DDR3 will not run properly with DDR4. However, the fine-grained QoS logic, as described above, can enable the amount of bandwidth on the DRAM data bus to be provisioned to each master device such that the provisioning with DDR4 is the same as it was with DDR3. Thus, the fine-grained QoS logic can restrict the a future generation of memory device, such as DDR4, such that its bandwidth variation is substantially similar to a prior generation of memory device, such as DDR3, and in doing so allow software designed for the prior generation of memory device to function on a system using a later generation of memory device without experiencing additional bandwidth variation.

[0049] Another example a use case scenario for the systems and methods described above is to reduce the amount of bandwidth variation experienced by a master device when additional masters device are added to the system. If the system is a video game console, for example, additional masters may be added even after the console has been launched and software has been designed for it. Additional master device might undesirably introduce additional bandwidth perturbations to the system. However, the fine-grained QoS logic can be used to isolate the master devices and schedule the requests generated therefrom in a logical manner based on priority class and the fine-grained QoS parameters discussed above, thereby shielding the system from any added bandwidth perturbation that would otherwise be caused by adding new master devices. To accomplish this, the master devices can be assigned a fixed amount of bandwidth on the DRAM data bus and, so long as there is adequate bandwidth to support all the master devices, the bandwidth assigned to each master will remain unchanged even when new masters are added. This results in a significant reduction in the amount of testing needed to ensure that the master device will behave in the same manner as before the additional masters were added. In such an application, the fine-grained QoS logic is used to reserve a fixed amount of bandwidth for each of the master devices in the system, and to use priority class and fine grained QoS techniques to limit the bandwidth consumed by the new master from interfering with the fixed reserved bandwidths of the existing masters.

[0050] A third example is to use the systems and methods described above to partition a physical DRAM device into many logically independent DRAM devices.

Master devices can be assigned to these logically independent DRAM devices and thereby be shielded from the bandwidth perturbation caused by other master devices assigned to the other partitions. Once such a partition is made, master devices can be added and removed without disturbing the master devices on the other partitions.

5 [0051] It will be appreciated that the systems and methods above described in the context of DRAM memory devices may also be applied to other types of memory devices, including static random access memory (SRAM), FLASH memory, magnetic disc memory, tape memory, and other kinds of volatile and non-volatile memory devices. Thus, the DRAM data bus described above may be a memory data bus, the DRAM
10 memory requests described above may be memory requests, the DRAM commands described above may be memory commands, and the DRAM command bus described above may be a memory command bus, these various memory components being configured to operate with a particular type of memory device, such as those listed above in this paragraph.

15 [0052] FIG. 7 illustrates the computing environment in which the above described systems and methods may be implemented. As illustrated in FIG. 7, computing system 700 may include not only host system 718 with a processor 720, non-volatile memory 722 and volatile memory 724, but also may include a display subsystem 706, communication system 708, and input subsystem 710, all linked by a communications bus subsystem 702.

20 [0053] System 10 may take the form of one or more personal computers, server computers, tablet computers, home-entertainment computers such as a cable set top box, digital video recorder, video server etc., network computing devices, gaming devices such as a game console, mobile computing devices, mobile communication devices (e.g., smart phone), and/or other computing devices.

25 [0054] The processor 720 of the host system 718 is a physical device configured to execute instructions using portions of non-volatile memory 722 and the volatile memory 724. For example, the processor may be configured to execute instructions that are part of one or more applications, services, programs, routines, libraries, objects, components, data structures, or other logical constructs. Such instructions may be implemented to perform a
30 task, implement a data type, transform the state of one or more components, achieve a technical effect, or otherwise arrive at a desired result.

[0055] Non-volatile memory 722 includes one or more physical devices configured to hold instructions executable by the processor 720. Non-volatile memory 722 may include non-volatile memory such as FLASH, EPROM, EEPROM, etc., which hold

the instructions for a finite duration during processing, and which also continue to hold the instructions even when power is cut to the computing system. Volatile memory 724 may include DRAM, such as DRAM device 32 discusses above, or RAM which hold the instructions for a finite duration during processing but typically do not continue to store the instructions after the computing system is powered down.

[0056] Aspects of processor 720, permanent memory 722, and volatile memory 724, may be integrated together into one or more hardware-logic components. Such hardware-logic components may include field-programmable gate arrays (FPGAs), program- and application-specific integrated circuits (PASIC / ASICs), program- and application-specific standard products (PSSP / ASSPs), system-on-a-chip (SoC), and complex programmable logic devices (CPLDs), for example.

[0057] A display subsystem 706 may be used to present a visual representation of data held by permanent memory 722 or volatile memory 724. This visual representation may take the form of a graphical user interface (GUI). As the herein described methods and processes change the data held by non-volatile memory 722 or volatile memory 724, and thus transform the state of the non-volatile memory 722 or volatile memory 724, the state of display subsystem 706 may likewise be transformed to visually represent changes in the underlying data. Display subsystem 706 may include one or more display devices utilizing virtually any type of technology. Such display devices may be combined with processor 720 and/or non-volatile memory 722 or volatile memory 724 in a shared enclosure, or such display devices may be peripheral display devices.

[0058] Input subsystem 710 may comprise or interface with one or more user-input devices such as a keyboard, mouse, touch screen, or game controller. In some embodiments, the input subsystem may comprise or interface with selected natural user input (NUI) componentry. Such componentry may be integrated or peripheral, and the transduction and/or processing of input actions may be handled on- or off-board. Example NUI componentry may include a microphone for speech and/or voice recognition; an infrared, color, stereoscopic, and/or depth camera for machine vision and/or gesture recognition; a head tracker, eye tracker, accelerometer, and/or gyroscope for motion detection and/or intent recognition; as well as electric-field sensing componentry for assessing brain activity.

[0059] Communication subsystem 708 may be configured to communicatively couple system 10 with one or more other computing devices. Communication subsystem 708 may include wired and/or wireless communication devices compatible with one or

more different communication protocols. As non-limiting examples, the communication subsystem 708 may be configured for communication via a wireless telephone network, or a wired or wireless local- or wide-area network. In some embodiments, the communication subsystem may allow computing system 708 to send and/or receive
5 messages to and/or from other devices via a network such as the Internet.

[0060] It will be understood that the configurations and/or approaches described herein are exemplary in nature, and that these specific embodiments or examples are not to be considered in a limiting sense, because numerous variations are possible. The specific routines or methods described herein may represent one or more of any number of
10 processing strategies. As such, various acts illustrated and/or described may be performed in the sequence illustrated and/or described, in other sequences, in parallel, or omitted. Likewise, the order of the above-described processes may be changed.

[0061] The subject matter of the present disclosure includes all novel and nonobvious combinations and subcombinations of the various processes, systems and
15 configurations, and other features, functions, acts, and/or properties disclosed herein, as well as any and all equivalents thereof.

CLAIMS

1. A memory controller, configured to:
 - receive memory requests from a plurality of masters via a bus fabric;
 - determine a priority level assigned with each of the plurality of masters;
 - determine an amount of memory data bus bandwidth consumed by each master on a memory data bus;
 - apply a fine-grained quality of service (QoS) logic to compute a schedule for the memory requests based on a priority class assigned to each master and an amount of memory data bus bandwidth consumed by each master, the memory data bus bandwidth consumed by each master per request being calculated based on information about a last issued request; and
 - based upon the schedule,
 - convert the memory requests to memory commands;
 - send the memory commands to at least one memory device via a memory command bus; and
 - receive a response from the memory device via the memory data bus.

2. A method for performing fine-grained bandwidth provisioning in a memory controller, comprising:
 - receiving memory requests from a plurality of masters via a bus fabric;
 - determining a priority level assigned with each of the memory requests;
 - determining an amount of memory data bus bandwidth consumed by each master on a memory data bus;
 - applying a fine-grained quality of service (QoS) logic to compute a schedule for the memory requests based on a priority class assigned to each master and an amount of memory data bus bandwidth consumed by each master, the memory data bus bandwidth consumed by each master per request being calculated based on information about a last issued request; and
 - based upon the schedule,
 - converting the memory requests to memory commands;
 - sending the memory commands to at least one memory device via a memory command bus; and

receiving a response from the memory device via a memory data bus.

3. The method of claim 2, wherein the memory device is a dynamic random access memory (DRAM) device, the memory data bus is a DRAM data bus, the memory requests are DRAM memory requests, the memory commands are DRAM memory commands, and the memory command bus is a DRAM command bus, the memory controller being, the method further comprising:

assigning each DRAM request to one of a plurality of priority queues based on the priority level.

4. The method of claim 3, further comprising:

sending the DRAM requests from the plurality of priority queues to a master scheduling multiplexer;

periodically monitoring an unused portion of bandwidth allocated to each master; and

computing the schedule so as to throttle the DRAM requests at the memory controller if the unused portion of the bandwidth drops below a predetermined level.

5. The method of claim 4, wherein periodically monitoring an unused portion of bandwidth allocated to each master includes applying a fine-grained QoS bandwidth provisioning logic to allocate a predetermined amount of bandwidth in the DRAM data bus to each of the plurality of masters.

6. The method of claim 5, periodically monitoring an unused portion of bandwidth allocated to each master includes applying a QoS credit accumulator to incrementally accumulate credits for each of the plurality of masters, the rate of accumulation corresponding to the bandwidth allocated to the master by the fine-grained QoS bandwidth provisioning logic, and the total credits for each master not exceeding a predetermined maximum QoS credit value.

7. The method of claim 6, wherein periodically monitoring an unused portion of bandwidth allocated to each master includes:

applying a fine-grained QoS charge accounting logic to calculate a bandwidth on the DRAM data bus consumed by each DRAM memory request; and
computing a credit balance for each master by subtracting the calculated bandwidth from the credits in the QoS credit accumulator corresponding to each master.

8. The method of claim 7, wherein calculating the bandwidth on the DRAM data bus consumed by each DRAM memory request includes calculating the bandwidth cost of a request based on the information about a last issued request.

9. The method of claim 8, wherein applying a fine-grained scheduling logic and throttling DRAM requests at the memory controller if an unused portion of the bandwidth drops below a predetermined level are achieved by applying a fine-grained QoS scheduling logic to schedule the DRAM requests based on the priority class of each request, and to compute the schedule so as to throttle the DRAM requests if the credit balance in the QoS credit accumulator for the master issuing the DRAM request is below a predetermined minimum QoS credit value.

10. The method of claim 9, wherein sending the DRAM requests to at least one of a plurality of DRAM devices includes:

sending the DRAM requests from the master scheduling multiplexer to a DRAM protocol engine; and

sending the DRAM requests from the DRAM protocol engine to at least one of the plurality of DRAM devices.




FIG. 1

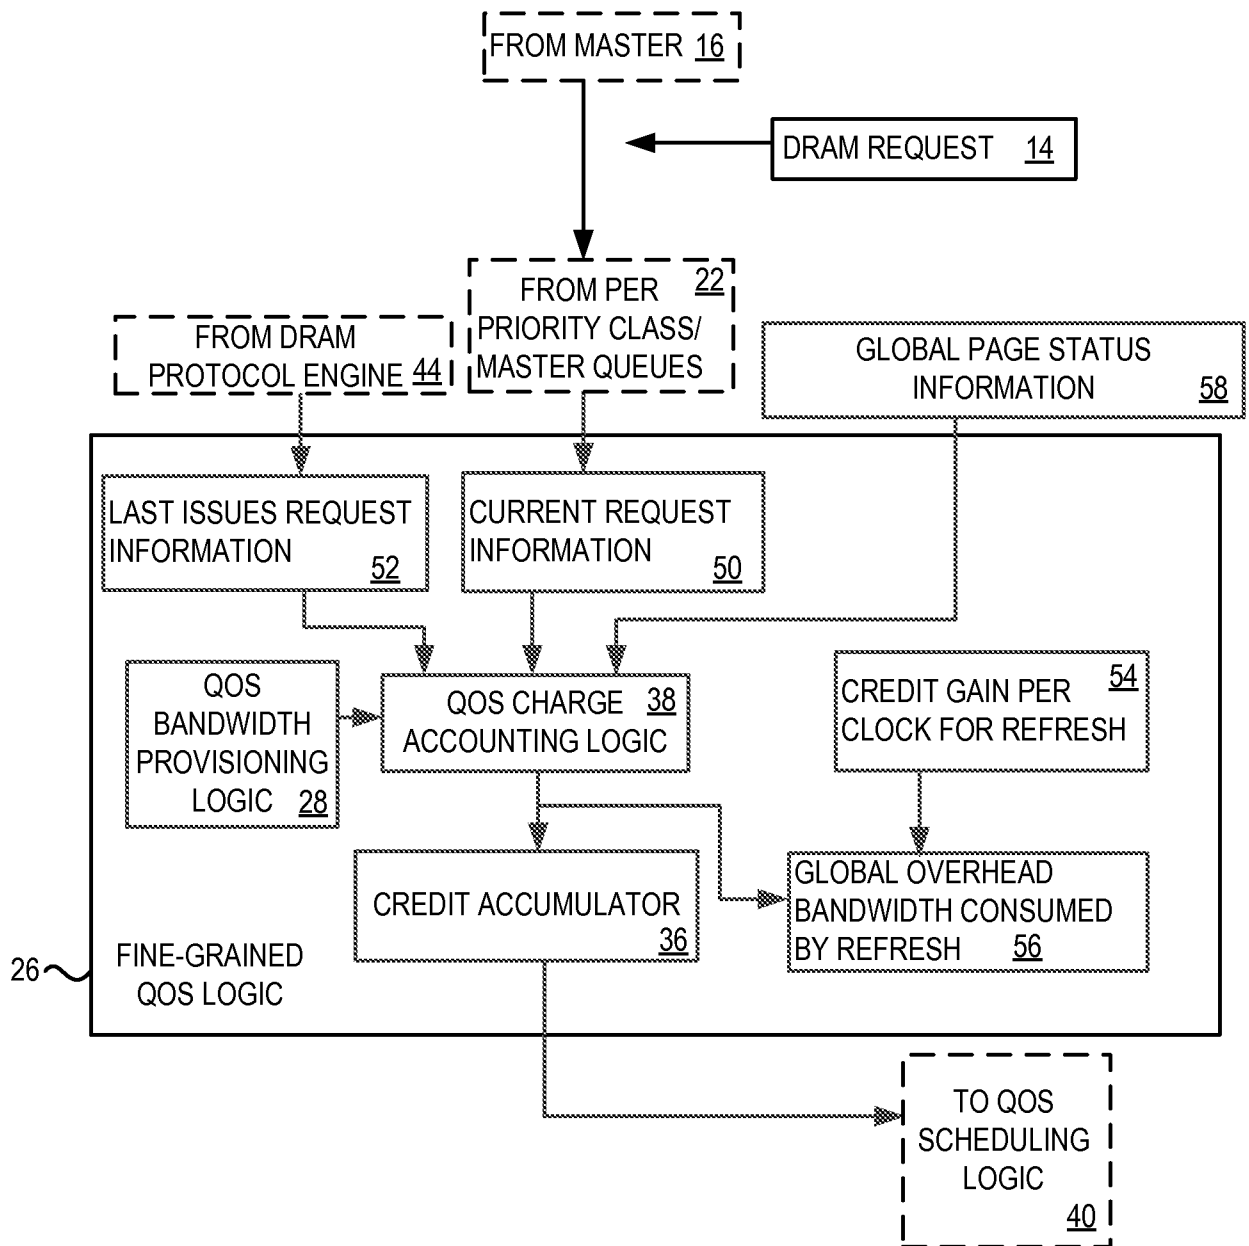


FIG. 2

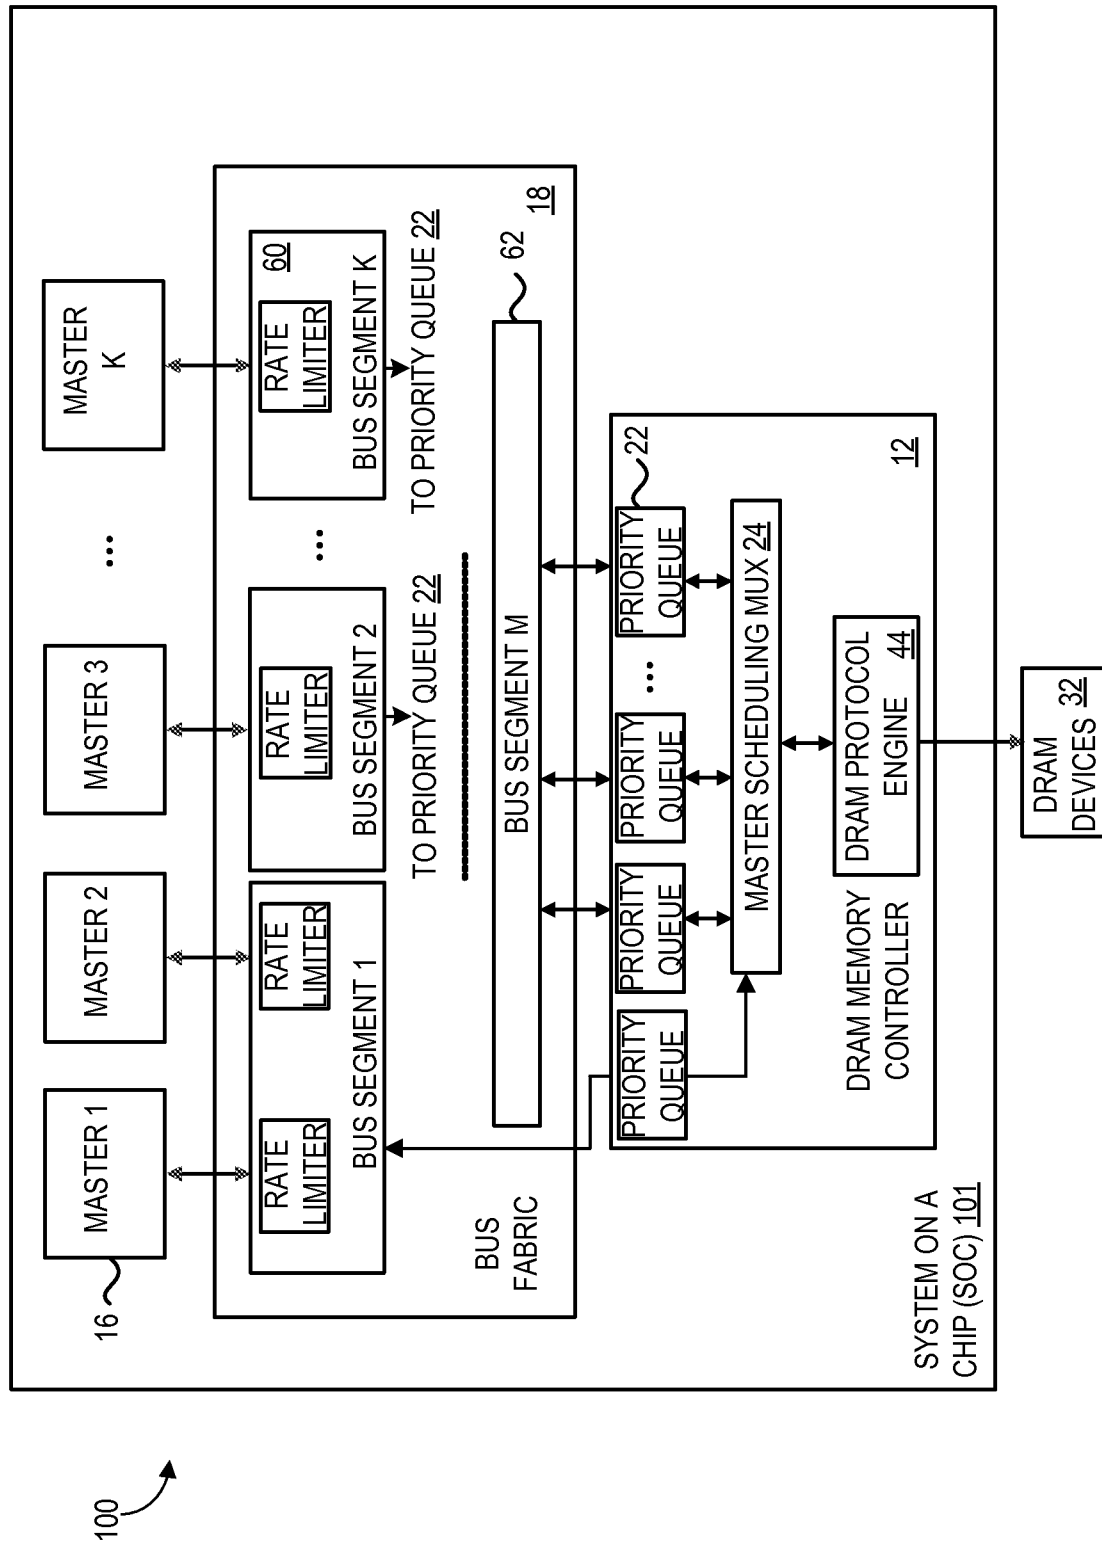


FIG. 3

4/7

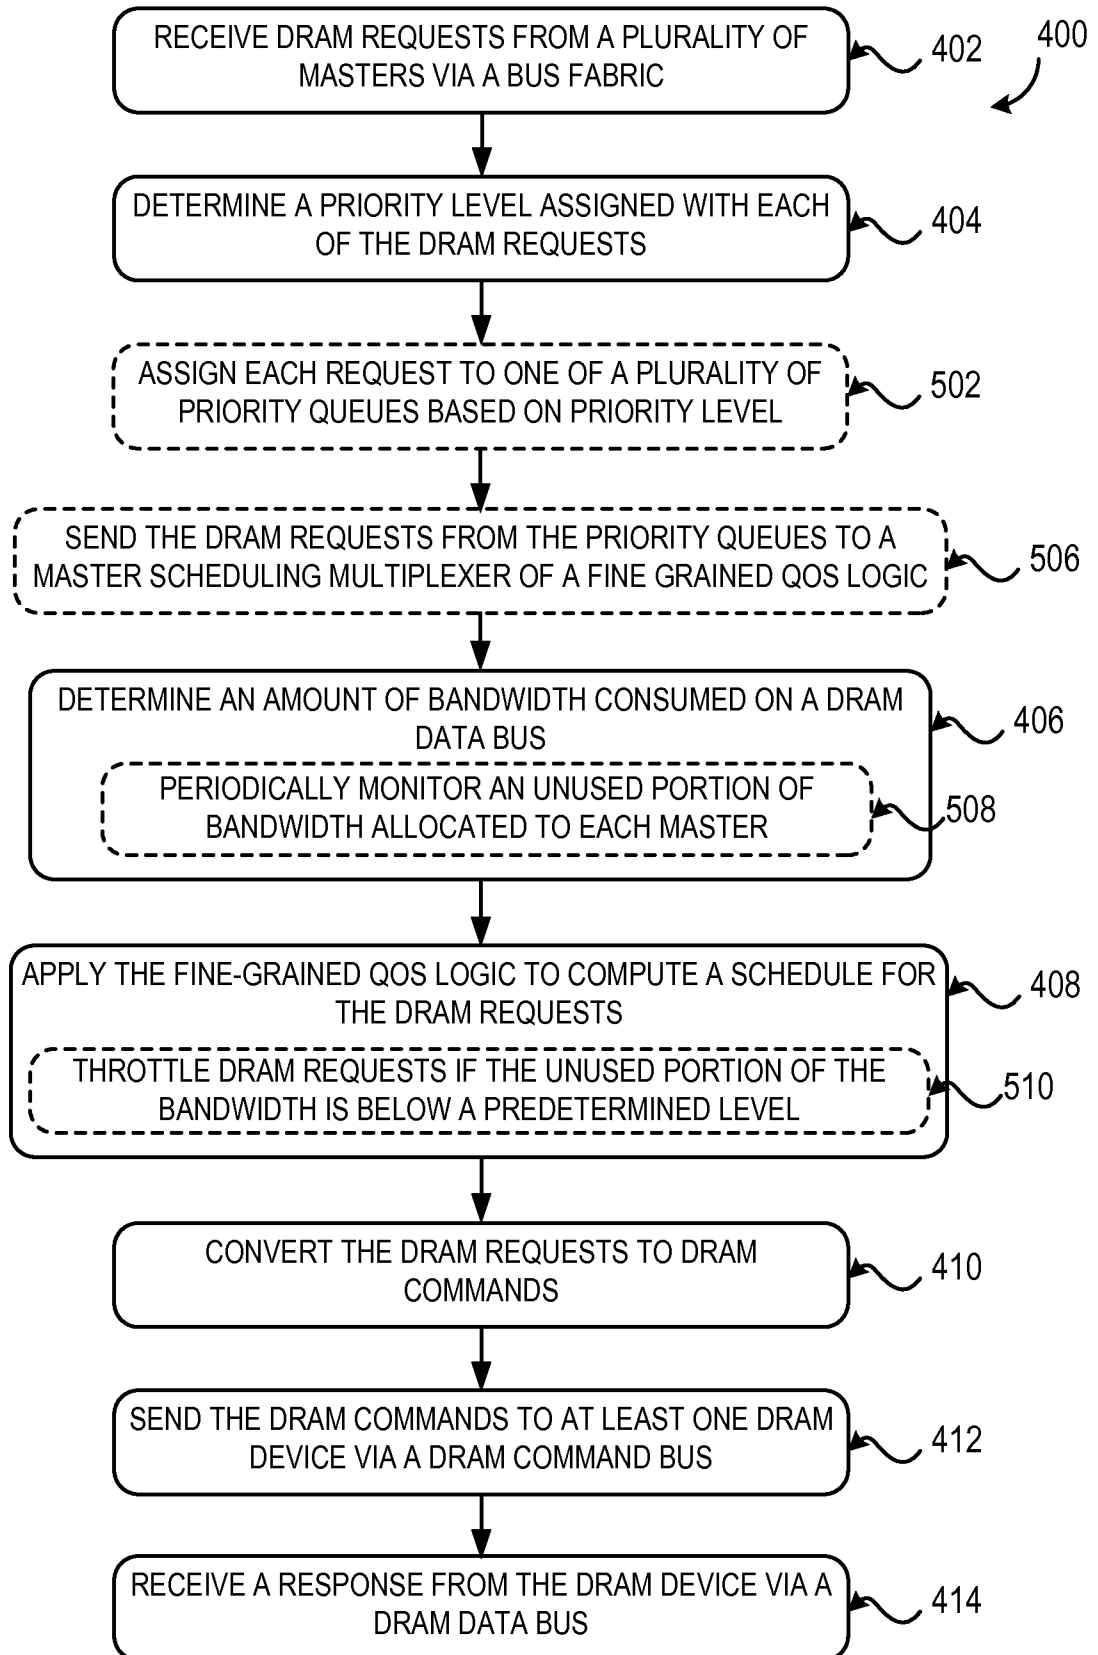


FIG. 4

5/7

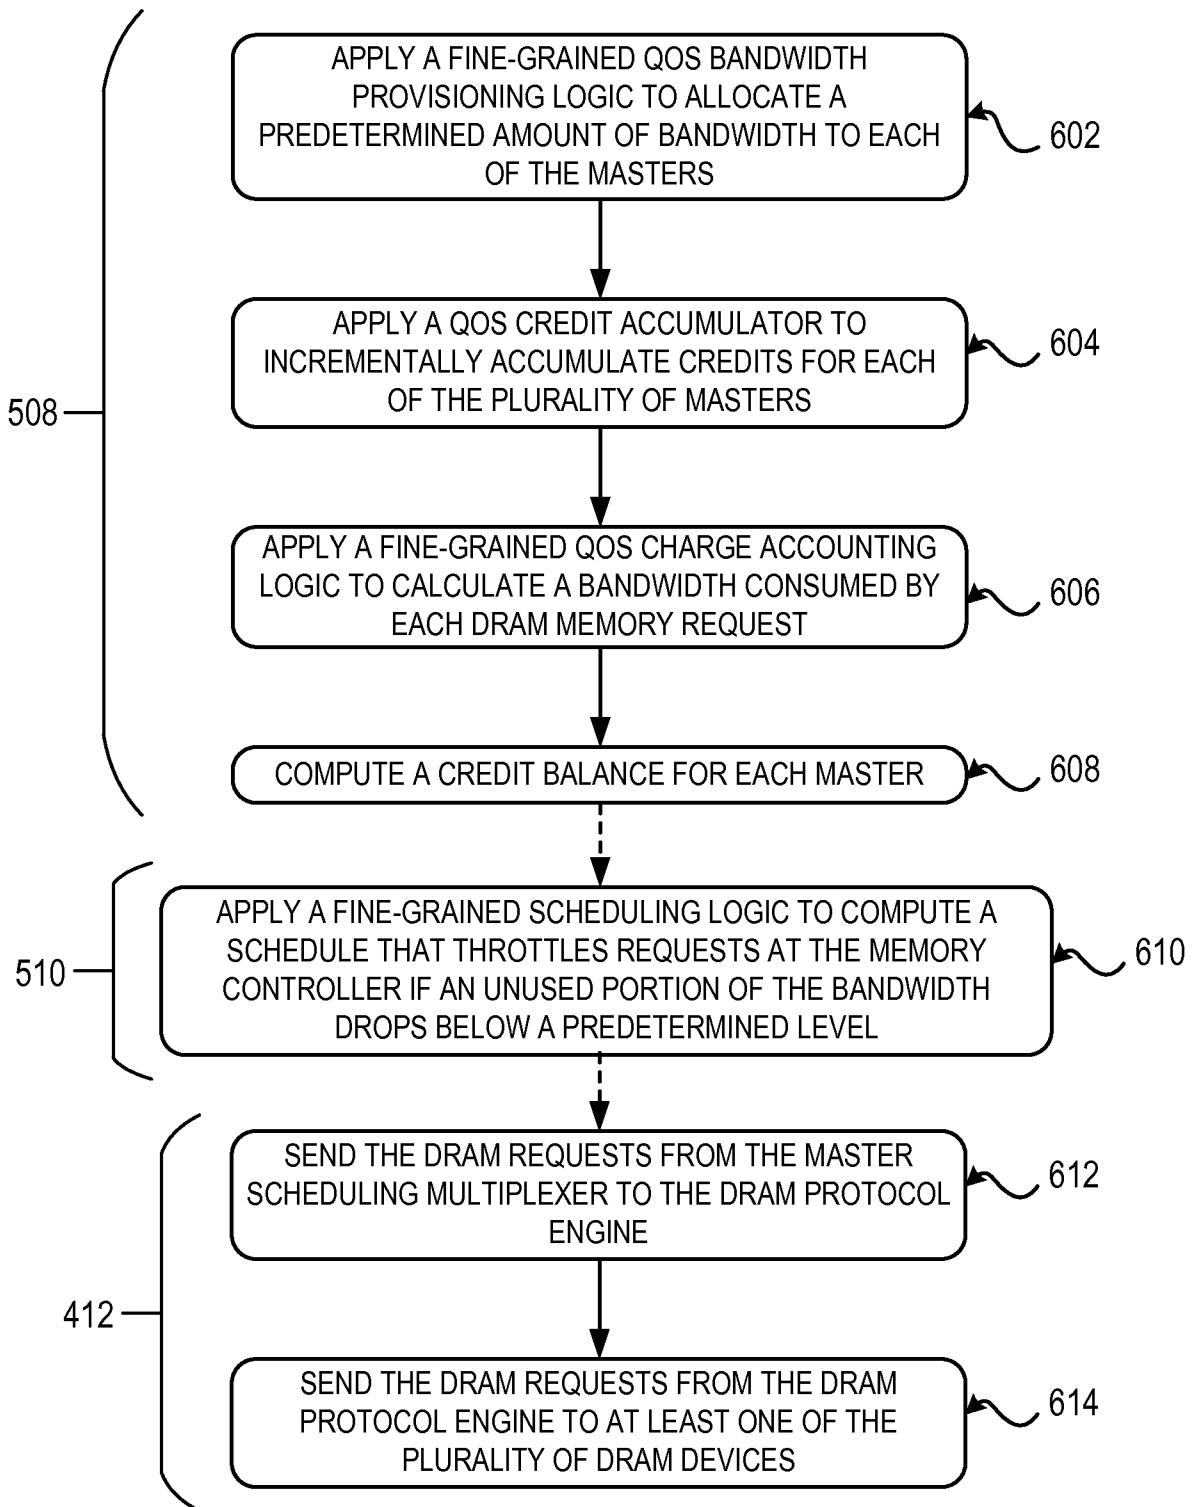


FIG. 5

6/7

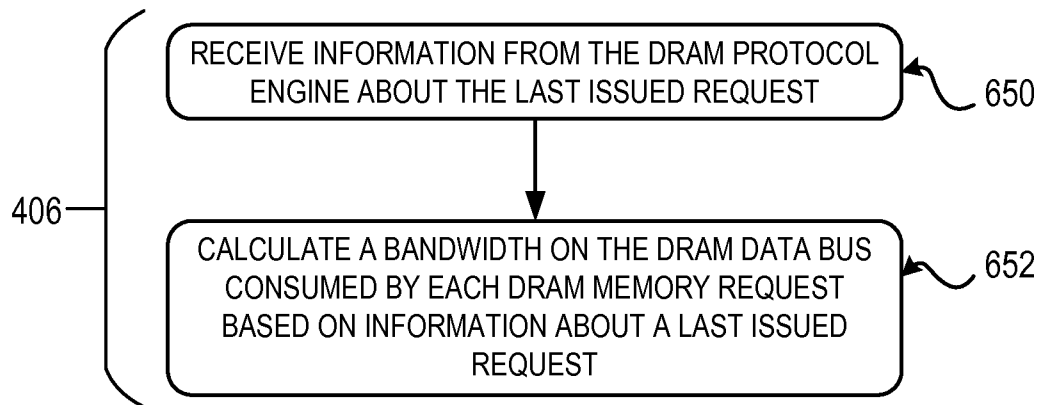


FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2015/024414

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 2007/113023 A1 (ARULAMBALAM [US] ET AL) 17 May 2007 (2007-05-17)	1-3
Y	paragraph [0056] - paragraph [0090] -----	4,5
Y	US 2011/154352 A1 (DESOTA [US] ET AL) 23 June 2011 (2011-06-23)	4,5
	paragraph [0033] - paragraph [0039] -----	
A	EP 2 431 884 A1 (APPLE INC [US]) 21 March 2012 (2012-03-21)	6-10
	paragraph [0066] - paragraph [0068] paragraph [0093] - paragraph [0099] -----	
A	US 2013/246727 A1 (HIKONO [JP]) 19 September 2013 (2013-09-19)	1-10
	paragraph [0009] - paragraph [0012] paragraph [0034] - paragraph [0052] paragraph [0067] - paragraph [0071] -----	
	-/-	



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

27 July 2015

Date of mailing of the international search report

04/08/2015

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

McDonagh, Fintan

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/024414

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 2008/209093 A1 (HUANG [TW] ET AL) 28 August 2008 (2008-08-28) paragraph [0020] - paragraph [0033] -----	1-10
A	US 2003/074519 A1 (WEBER [US]) 17 April 2003 (2003-04-17) paragraph [0014] - paragraph [0035] -----	1-10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2015/024414

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007113023	A1	17-05-2007	NONE
US 2011154352	A1	23-06-2011	US 2011154352 A1 23-06-2011 US 2012191939 A1 26-07-2012 US 2015032977 A1 29-01-2015
EP 2431884	A1	21-03-2012	AU 2011302452 A1 28-03-2013 CN 102402490 A 04-04-2012 EP 2431884 A1 21-03-2012 GB 2483763 A 21-03-2012 HK 1168159 A1 16-05-2013 JP 5610636 B2 22-10-2014 JP 2012074042 A 12-04-2012 KR 20120029366 A 26-03-2012 NL 2007411 C 09-05-2012 RU 2013117127 A 27-10-2014 TW 201216056 A 16-04-2012 WO 2012036905 A1 22-03-2012
US 2013246727	A1	19-09-2013	CN 103324590 A 25-09-2013 JP 2013196321 A 30-09-2013 US 2013246727 A1 19-09-2013
US 2008209093	A1	28-08-2008	NONE
US 2003074519	A1	17-04-2003	AT 377793 T 15-11-2007 DE 60223394 T2 28-08-2008 EP 1435039 A1 07-07-2004 JP 4095032 B2 04-06-2008 JP 2005508550 A 31-03-2005 US 2003074519 A1 17-04-2003 WO 03040927 A1 15-05-2003