

(51) International Patent Classification:
G06F 15/78 (2006.01)(21) International Application Number:
PCT/US2010/035456(22) International Filing Date:
19 May 2010 (19.05.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/478,532 4 June 2009 (04.06.2009) US(71) Applicant (for all designated States except US): **MICRON TECHNOLOGY, INC.** [US/US]; 8000 S. Federal Way, Boise, Idaho 83707-0006 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **WALKER, Robert** [US/US]; 9000 Deerland Grove Drive, Raleigh, North Carolina 27615 (US).(74) Agents: **MANWARE, Robert A.** et al.; 7915 FM 1960 West, Suite 330, Houston, Texas 77070 (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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, 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))

[Continued on next page]

(54) Title: CONTROL OF PAGE ACCESS IN MEMORY




FIG. 3

(57) Abstract: The present techniques provide systems and methods of controlling access to more than one open page in a memory component, such as a memory bank. Several components may request access to the memory banks. A controller can receive the requests and open or close the pages in the memory bank in response to the requests. In some embodiments, the controller assigns priority to some components requesting access, and assigns a specific page in a memory bank to the priority component. Further, additional available pages in the same memory bank may also be opened by other priority components, or by components with lower priorities. The controller may conserve power, or may increase the efficiency of processing transactions between components and the memory bank by closing pages after time outs, after transactions are complete, or in response to a number of requests received by masters.



Published:

- *without international search report and to be republished
upon receipt of that report (Rule 48.2(g))*

CONTROL OF PAGE ACCESS IN MEMORY

BACKGROUND

Field of Invention

[0001] Embodiments of the invention relate generally to memory systems, and more particularly, to access of memory.

Description of Related Art

[0002] This section is intended to introduce the reader to various aspects of art that may be related to various aspects of the present invention, which are described and/or claimed below. This discussion is believed to be helpful in providing the reader with background information to facilitate a better understanding of the various aspects of the present invention. Accordingly, it should be understood that these statements are to be read in this light and not as admissions of prior art.

[0003] Electronic systems typically include one or more memory components, which may be accessed by various other components, such as other devices or applications in the electronic system. A memory component may be, for example, any data structure including a sequence or group of data elements. In some embodiments, a memory component may be a memory array, such as a dynamic random access memory (DRAM) array. A memory component may also be a bank or banks which make up a memory array. For example, a controller in the electronic system may access data in memory components, such as a memory array, to obtain data for processing. A “controller” may refer to any component, such as a processor, capable of substantially controlling access to the memory

component. For example, the controller may open or close pages in the memory component, and may receive and grant requests from other components of the system to access the memory component. The controller may be a silicon-on chip (SOC), which may be external to a memory device or memory system containing a memory component. Further, in some embodiments, the controller may also be an internal processor embedded on or packaged with the memory device containing the memory component. Further, other components of the electronic system may also store data in the memory components.

[0004] Generally, more than one component of the electronic system may request access to a memory component, which may have multiple banks, each further including a number of pages. As used herein, a “page” refers to a row, or any other sequence or group of data elements organized in a memory. However, in a typical memory component, each bank may have only one page open at a time. For example, while a first component in the system may have access to a row in a bank, a second component may have to wait for the system to close the open row before a row requested by the second component is opened. Thus, the number of transactions between system components and the memory components may be limited by the access patterns of all the components requesting access to the memory.

[0005] In some systems, rows are configured to be left open once they are requested, and a component targeting the open row would not have to wait for the row to be opened. This “open page policy” may be advantageous for systems where requests for certain rows in a memory component are relatively predictable (e.g., a frequently requested row). In other systems, rows may be configured to close once a transaction is completed. This “closed page policy” may be implemented in systems where page access patterns are more

random. However, the open page policy may require more power to operate, and the closed page policy may require more time to process each access request. Furthermore, in typical electronic systems, once a system is configured to have an open page policy or a closed page policy, the policy is set. The limitations of opening only one page at a time, and the set policy on page access in a memory bank may limit the efficiency of transactions in an electronic system

BRIEF DESCRIPTION OF DRAWINGS

[0006] Certain embodiments are described in the following detailed description and in reference to the drawings in which:

[0007] FIG. 1 depicts a block diagram of a processor-based system in accordance with an embodiment of the present technique;

[0008] FIG. 2 illustrates a memory array having multiple banks with multiple pages, in accordance with an embodiment of the present technique;

[0009] FIG. 3 depicts a flow chart illustrating a process for controlling a transaction between a master and an assigned page in the memory array, in accordance with one or more embodiments of the present technique; and

[0010] FIG. 4 depicts a flow chart illustrating a process for controlling a transaction between multiple masters and a requested page in the memory array, in accordance with one or more embodiments of the present technique.

DETAILED DESCRIPTION

[0011] In a typical electronic system, processing or operating speed may be partially limited by the access of memory by different components of the system. For example, more than one device may request access to different pages in a memory bank. As memory banks may typically have only one open page at a time, one or more devices may delay access to a requested page when the memory bank already has a different page open. Such delays may limit the number of transactions between components of the system and the memory, which may also limit the efficiency of the electronic system.

[0012] Some systems implement different methods of controlling access to pages in a memory component to address the limitations in accessing a memory component. For example, as discussed, some systems may implement an “open page policy” when accesses to a certain page in a bank are relatively predictable (e.g., a certain page is frequently accessed). For example, a graphics or video system may benefit from the open page policy. One disadvantage to this policy may be that more power is used in keeping the page constantly open. The “closed page policy” may conserve some power, and may be advantageous to systems where accesses to pages in the memory component are less predictable. For example, a modem may access a memory component which operates under a closed page policy, as the access patterns of a typical modem may be comparatively random. However, the closed page policy may cause unnecessary delays when a device requests access to a particular page more than once in immediate succession, as the page would still have to be closed by the controller each time one transaction is complete. Typically, once a system sets a memory bank for either the open page or closed page policies, access to that memory bank may always operate under the set policy. This

limitation, in addition to the limitations resulting from opening only one page in a bank at one time, may reduce the efficiency of transactions between system components and the memory component.

[0013] One or more embodiments of the present techniques involve systems and methods of controlling access to more than one page in a memory bank at one time. A controller may substantially control access to more than one concurrently open page in the memory bank, and may organize access patterns to more efficiently process transactions between a master component (a component in the system requesting access to the memory component, hereinafter “master”) and each memory bank. In one embodiment, the controller may prioritize a master and assign a page in the memory component to the prioritized master component. As more than one page may be open in the memory, additional requests from other masters may also be processed based on a number of factors including, for example, an order of received requests, a priority of the other masters, or some combination of factors. In another embodiment, the controller may assign different priorities to multiple masters. When a master requests access to a page, the page may be opened by the controller, and may remain open until a request from a higher priority master is received, or until the transaction times out. Some embodiments of the present techniques may assign a higher priority to masters requiring low latency, and may assign certain pages to such low latency masters.

[0014] Further, in some embodiments, methods of controlling page access may be dynamic, as a controller may change page access control methods based on a condition of the system (e.g., based on a change in the condition of the system). Page access control methods may also be a hybrid of other methods. For example, any memory bank may

include different combinations of open pages, closed pages, or assigned pages, and page accesses may be based on a priority of masters, a chronological order of received requests, or any combination of the above.

[0015] Different components requesting access to the memory component may be referred to as a “master,” and may request a transaction from the controller of the memory component. A “transaction” may refer to a write or a read to a page in the memory component, or any type of access to any part of the memory component.

[0016] Now turning to the figures, FIG. 1 depicts a processor-based system, generally designated by reference numeral 10. As is explained below, the system 10 may include various electronic components. The system 10 may be any of a variety of types such as a computer, pager, cellular phone, personal organizer, control circuit, etc. In a typical processor-based system, one or more processors 12, such as a microprocessor, control the processing of system functions and requests in the system 10. As is explained below, the processor 12 and other subcomponents of the system 10 may include memory devices manufactured in accordance with one or more embodiments of the present technique.

[0017] The system 10 typically includes a power supply 14. For instance, if the system 10 is a portable system, the power supply 14 may advantageously include a fuel cell, a power scavenging device, permanent batteries, replaceable batteries, and/or rechargeable batteries. The power supply 14 may also include an AC adapter, so the system 10 may be plugged into a wall outlet, for instance. The power supply 14 may also include a DC adapter such that the system 10 may be plugged into a vehicle cigarette lighter, for instance.

[0018] Various other components may be coupled to the processor 12 depending on the functions that the system 10 performs. For instance, an input device 16 may be coupled to the processor 12. The input device 16 may include buttons, switches, a keyboard, a light pen, a mouse, a digitizer and stylus, and/or a voice recognition system, for instance. A display 18 may also be coupled to the processor 12. The display 18 may include an LCD, an SED display, a CRT display, a DLP display, a plasma display, an OLED display, LEDs, and/or an audio display, for example. Furthermore, an RF sub-system/baseband processor 20 may also be coupled to the processor 12. The RF sub-system/baseband processor 20 may include an antenna that is coupled to an RF receiver and to an RF transmitter (not shown). One or more communication ports 22 may also be coupled to the processor 12. The communication port 22 may be adapted to be coupled to one or more peripheral devices 24 such as a modem, a printer, a computer, or to a network, such as a local area network, remote area network, intranet, or the Internet, for instance.

[0019] The processor 12 generally controls the system 10 by implementing software programs whose implementing instructions are stored in the memory. The software programs may include an operating system, database software, drafting software, word processing software, and/or video, photo, or sound editing software, for example. The memory is operably coupled to the processor 12 to store and facilitate execution of instructions to implement various programs. For instance, the processor 12 may be coupled to the system memory 26, which may include dynamic random access memory (DRAM), and/or synchronous dynamic random access memory (SDRAM). The system memory 26 may include volatile memory, non-volatile memory, or a combination thereof.

The system memory 26 is typically large so that it can store dynamically loaded instructions and data.

[0020] The processor 12 may also be coupled to non-volatile memory 28, which is not to suggest that system memory 26 is necessarily volatile. The non-volatile memory 28 may include read-only memory (ROM), such as an EPROM, resistive read-only memory (RROM), and/or flash memory to be used in conjunction with the system memory 26. The size of the ROM is typically selected to be just large enough to store instructions for any necessary operating system or application programs, and fixed data. Additionally, the non-volatile memory 28 may include a high capacity memory such as a tape or disk drive memory, such as a hybrid-drive including resistive memory or other types of non-volatile solid-state memory, for instance.

[0021] The processor 12 may communicate with and/or control elements of the system memory 26. In some embodiments, the processor 12 may include a controller capable of controlling access to memory components in the system memory 26, in accordance with one or more of the present techniques. Though a processor 12 may include many different processor components capable of performing various tasks in the system 10, a processor capable of controlling access to memory components may also be referred to as a controller 12. Further, the system memory 26 may include memory components which may be controlled by the controller 12, including, for example, a memory array 30, illustrated in FIG. 2. The memory array 30 may have multiple banks 32 (as labeled by “Bank 0” to “Bank N”), and each bank 32 may have multiple pages 34 (as labeled by “Page 0” to “Page N”). In embodiments, more than one page 34a-34d may be open concurrently, such that a

controller 12 may more efficiently allow transactions between different masters and different pages 34 in the same bank 32.

[0022] One embodiment of the present techniques, controlling access to multiple open pages in a memory bank, is a page assignment process 50, depicted in the flow chart of FIG. 3. The process starts (block 52) when the controller 12 assigns a priority to a component which may request access to the memory array 30 (a “master”) (FIG. 2) (block 54). The controller 12 may assign a priority to a master based on an application that is running in the system 10 (FIG. 1) or based on applications running or requests made by other components in the system 10. The controller 12 may then assign a page 34a to the priority master (block 56). For example, if the system 10 is initiating a graphics program, the controller 12 may assign a higher priority to the graphics master, and the assigned page 34a may be based on predicted access patterns (e.g., the controller may assign pages 34 that the graphics master is likely to access).

[0023] The controller may receive a request from a master for a transaction with a page 34a-34d (block 58). As the controller 12 has already assigned a priority master, and assigned page 34a to the priority master (block 54 and 56), the controller 12 may determine whether the master making the request is a priority master to which page 34a has been assigned (block 60). If the requesting master is not a priority master, the requesting master may still be able to access a page 34a-34d in the bank 32, as in one or more embodiments, more than one page 34 may be open concurrently. Pages 34 in the same bank 32 as the assigned page 34a may be available to the requesting master based on other factors, such as limitations on the number of open pages 34 per bank 32, or requests from other masters in the system 10. The controller 12 may determine whether the page requested by the

requesting master is available, based on these factors (block 62). If the requested page is not available, the controller 12 may queue the request (block 64) until the controller 12 determines that the requested page is available (block 62). If the requested page is available, the controller 12 may open the page (block 66) even if the requesting master does not have priority, as more than one page 34 may be open concurrently in one bank 32.

[0024] If the controller 12 determines that the requesting master is the priority master, the controller 12 may open the assigned page 34a for the priority master (block 66). The assigned page 34a may operate on a closed page policy, meaning the page 34a may be closed once the transaction is complete. Thus, the controller 12 may determine whether the transaction is complete (block 68). If the transaction is still running (i.e., the priority master is still accessing the assigned page 34a), the assigned page 34a may remain open (block 70). If the controller 12 determines that the transaction has completed, the controller 12 may close the assigned page 34a (block 72). In one embodiment, the controller 12 may be configured to receive a request for a transaction (block 58) after closing the assigned page (block 72), or at any other point in the process 50.

[0025] Another embodiment of the present techniques, controlling access to more than one open page in a memory bank, may be depicted in the open page control process 100 of FIG. 4. The process 100 may start (block 102) when the controller 12 assigns priority for more than one (N) masters (block 104). For example, in applications where more than one device in a system 10 may randomly access different pages 34 in a memory array 30, the controller 12 may increase the number of transactions between masters and pages 34 in the memory array 30 by increasing the likelihood that a page 34a-34d will be available for a priority master. The controller 12 may receive a request for a transaction from a priority

master (block 106), and open the requested page 34b (block 108). Thus, the priority master may access the opened page 34b even if another page 34a-34d in the bank 32 may also be open, as more than one page 34 may be open in the same bank 32. The controller 12 may receive a request for a different page 34c from a different priority master (block 110), and the controller 12 may also open the different requested page 34c. Though more than one page 34 may be open concurrently in the same bank 32, the process 100 illustrates an example of how a controller 12 may control the access of pages 34 when additional requests exceed a number of available pages 34 in a bank 32. In embodiments, not every request for a page 34 may result in an opening of the requested page 34, and the controller 12 may determine whether the requesting master has higher priority than the priority master which is accessing the opened page 34b (block 112). If the controller 12 determines that the new requesting master has a higher priority, the controller may then determine whether available pages remain in the memory bank 32, or whether another page 34c may be concurrently open with other open pages 34 in the bank 32 (block 114). For example, a bank 32 may have eight pages 34 (though only four are illustrated in FIG. 2), the bank 32 may be limited to have only four pages 34 open concurrently. If the limit has been reached, the controller may close the opened page 34b (block 120) and open the page 34c requested by the master with the higher priority (block 108).

[0026] If the controller determines that the master making the new request does not have higher priority (block 112), the controller 12 may determine if a timeout has occurred (block 116). A “timeout” may refer to situations where, for example, the priority master does not return results in some number of clock cycles, or some number of transactions. A timeout may also be based on the number of other requests from lower priority masters that

enter the queue between the highest priority master transactions. In one or more embodiments, the process 100 may dynamically control access to pages in memory by altering access patterns based on a condition of the system 10. For example, although a controller 12 may determine that a requesting master has lower priority, the controller 12 may track the number of requests from lower priority masters, or track the number of lower priority masters in the queue. Even though a higher priority master may be accessing a page 34, timing out and closing the page 34 used by the higher priority master when a certain number of requests are made by lower priority masters may increase the utilization of the memory and the overall performance of the system 10. If the controller 12 determines that a timeout has occurred (block 116), the controller may close the opened page 34b (block 120) and open a different requested page (block 108) based on the requests of different priority masters. If the controller 12 determines that a timeout has not occurred (block 116), the controller 12 may keep the page 34b open (block 118). At any point in the process 100, the controller 12 may receive requests from different masters in the system 10.

[0027] In some embodiments, the priority of masters may be based on latency requirements of the masters. Low latency masters may refer to components in the system 10 making memory requests where speed of accessing memory is significant. It may be advantageous to reduce any latency in the memory access time for a low latency master. For example, and referring back to FIG. 2, one or more pages 34 in a memory bank 32 may be assigned to low latency masters. The pages 34 assigned to low latency masters may be accessed under a close page policy, such that once a low latency master requests a transaction and enters the queue, the controller 12 may determine that the low latency master has high priority, and may open the assigned page 34. In one embodiment, all low

latency masters may use the same assigned page 34, or each low latency master may have a different assigned page. In some embodiments, hybrid control methods may be implemented, and the pages 34 in a bank 32 may be assigned or accessed through a combination of methods. For example, in a memory array 30 with four available pages 34 per bank 32, a master A may be a low latency master, and may be assigned to a page 34a. Masters B and C may also be low latency masters, and may share a page 34b. Two other pages 34c and 34d may be used for all other masters. In embodiments, the four available pages 34 may be open concurrently, and the opening and closing of the pages 34 may be controlled by the controller 12, and may be based on the access patterns of the masters in the system 10.

[0028] Furthermore, the present techniques may also be dynamic, and may include modifying a page access method depending on the operation of the system 10. The methods of controlling open pages in a memory component are not necessarily permanent, and may be dynamically altered to improve the efficiency of processing transactions between the masters and the memory component. For example, if the queue of pending transaction requests reaches a certain threshold, then the processing patterns of a current method (e.g., process 50 or process 100) may be modified to allow any master to use any available pages. The methods which assign a page to a priority master, or the methods which prioritize transaction requests based on the priority of the master may be temporarily changed to allow the controller 12 to process the pending transactions more efficiently. In some embodiments, once the queue length falls below a lower threshold, the controller 12 may return to an original method of controlling open pages.

[0029] Any of the systems and methods previously described of controlling open pages in a memory component may be combined. Additionally, throughout the operation of the system 10, any methods may be changed based on various conditions in the system 10. Methods of page access and changes in methods based on system conditions may be dynamic and/or hybrid, and may all be implemented to increase the efficiency of processing transactions between masters and memory components in the system 10.

[0030] While the invention may be susceptible to various modifications and alternative forms, specific embodiments have been shown by way of example in the drawings and have been described in detail herein. However, it should be understood that the invention is not intended to be limited to the particular forms disclosed. Rather, the invention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the following appended claims.

CLAIMS

1. A system comprising:

a memory component comprising a plurality of pages, wherein more than one of the plurality of pages can be opened concurrently; and

a controller coupled to the memory component and configured to substantially control access to the plurality of pages.
2. The system, as set forth in claim 1, wherein the memory component comprises a plurality of memory banks, each comprising a portion of the plurality of pages.
3. The system, as set forth in claim 1, wherein the processor is configured to substantially control when each of the plurality of pages is open or closed.
4. The system, as set forth in claim 2, wherein the controller is configured to open more than one page concurrently in one of the plurality of memory banks, wherein the open page allows a transaction between a component of the memory system and the memory component.
5. The system, as set forth in claim 4, wherein the transaction comprises one or more of reading data from the open page or writing data to the open page.

6. The system, as set forth in claim 4, wherein the controller is configured to keep the more than one page open after the transaction is complete.
7. The system, as set forth in claim 4, wherein the controller is configured to close the more than one page after each transaction is complete.
8. The system, as set forth in claim 1, comprising a plurality of components configured to request access to one or more of the plurality of pages.
9. The system, as set forth in claim 8, wherein the controller is configured to selectively open or close one or more of the plurality of pages in response to a request for access from the plurality of components.
10. The system, as set forth in claim 1, wherein the controller is further configured to organize access patterns.
11. The system, as set forth in claim 1, wherein the controller is further configured to prioritize a master component and assign a page in the memory component to the prioritized master component.
12. The system, as set forth in claim 1, wherein the controller is further configured to assign different priorities to components of the system.

13. The system, as set forth in claim 1, wherein the controller is further configured to open a page when a component in the system requests access to the page and to keep the page open until either a request from another component with a higher priority is received or a timeout occurs.

14. The system, as set forth in claim 1, wherein the controller is further configured to change a page access control method based on a condition of the system.

15. The system, as set forth in claim 1, wherein the controller is further configured to substantially control access based on a priority of a component requesting access.

16. The system, as set forth in claim 1, wherein the controller is further configured to substantially control access based on an order of received requests for access.

17. The system, as set forth in claim 1, wherein the controller is further configured to dynamically control access to the pages by altering access patterns based on a condition of the system.

18. A method of controlling a memory component in a system, comprising:
receiving a first request for accessing the memory component from a master component in the system;
opening a first page in the memory component in response to the first request;
receiving a second request for accessing a second page in the memory component from another master component; and

opening the second page in response to the second request while the first page is open.

19. The method, as set forth in claim 18, wherein the first page is assigned to one or more masters.

20. The method, as set forth in claim 18, wherein the first page is opened if the first request is received from a master component to which the first page has been assigned.

21. The method, as set forth in claim 18, further comprising determining whether the first page is assigned to the master component from which the first request was received.

22. The method, as set forth in claim 21, wherein opening the first page comprises opening the first page only if the first page is available to be opened, if the first page is not assigned to the master component from which the first request was received.

23. The method, as set forth in claim 18, further comprising determining if the first page is available to be opened.

24. The method, as set forth in claim 23, wherein determining if the first page is available to be opened comprises determining whether opening the first page will exceed a limitation on the number of open pages that can be open in the memory component.

25. The method, as set forth in claim 23, wherein determining if the first page is available to be opened is based on whether there are other requests from components in the system.

26. The method, as set forth in claim 18, further comprising operating the first page on a closed page policy if the first page is assigned to the component from which the first request was received.

27. The method, as set forth in claim 18, further comprising closing the first page if the transaction between the component from which the first request was received and the first page is complete.

28. The method, as set forth in claim 18, further comprising closing the first page when a transaction between the master requesting access to the first page and the first page times out.

29. The method, as set forth in claim 18, further comprising closing the first page when a number of lower priority masters request access to the memory component.

30. The method, as set forth in claim 18, wherein the second request is received after the first page is closed.

31. The method, as set forth in claim 18, further comprising determining whether the second page is available before opening the second page in response to the second request.

32. The method, as set forth in claim 18, comprising closing the second page if the transaction between the component from which the second request was received and the second page is complete.

33. A method of controlling open pages in a memory bank, comprising:
receiving a first request for a first transaction from one of a number of priority masters;
opening a first page in a memory bank in response to the first request;
receiving a second request for a second transaction from a different master; and
opening a second page in the memory bank in response to the second request.

34. The method, as set forth in claim 33, comprising closing the first page if a timeout occurs.

35. The method, as set forth in claim 34, wherein a timeout occurs if the component from which the first request was received does not return results in a certain number of clock cycles.

36. The method, as set forth in claim 34, wherein a timeout occurs if the component from which the first request was received does not return results in a certain number of transactions.

37. The method, as set forth in claim 34, wherein a timeout occurs if a certain number of requests are received from components that have a lower priority than the component from which the first request was received.

38. The method, as set forth in claim 34, wherein a timeout occurs if requests are received from a certain number of components that have a lower priority than the component from which the first request was received.

39. The method, as set forth in claim 33, further comprising closing the first page when the different master has a higher priority than the one of the number of priority masters from which the first request was received and the second page can not be opened unless the first page is closed.

40. The method as set forth in claim 39, wherein the second page can not be opened without closing the first page because doing so would exceed a limit on the number of pages in the bank which can be concurrently open.

41. The method, as set forth in claim 33, further comprising closing the first page when the first transaction does not produce a result in a segment of time.

42. The method, as set forth in claim 33, wherein the second page is opened when the different master has a higher priority than the one of the number of priority masters.

43. The method, as set forth in claim 33, wherein the second request is received from one of the number of components that have a lower priority.

44. The method, as set forth in claim 33, further comprising assigning a page in the memory bank to one or more of the number of priority masters.

45. The method, as set forth in claim 33, further comprising keeping the first page open after the first transaction is complete.

46. The method, as set forth in claim 33, further comprising closing the first page after the first transaction is complete.

47. The method, as set forth in claim 33, comprising:
keeping the first page open after the first transaction is complete; and
closing the second page after the second transaction is complete.

48. The method, as set forth in claim 33, comprising:
closing the first page after the first transaction is complete; and
keeping the second page open after the second transaction is complete.

49. A method of controlling more than one open pages in a memory bank, comprising:
assigning a first page to a first low latency master;
opening the first page for a first transaction between the first page and the first low latency master;

opening a second page for a second transaction between the second page and a non-priority master.

50. The method, as set forth in claim 49, further comprising assigning the first page to the first low latency master and one or more additional low latency masters.

51. The method, as set forth in claim 49, further comprising opening one or more additional pages in the memory bank for one or more additional transactions between the one or more additional pages and one or more masters in the electronic system.

52. The method, as set forth in claim 49, comprising one or more of:
closing the first page when the first transaction is complete;
closing the first page when the first transaction times out;
closing the second page when the second transaction is complete; and
closing the second page when the second transaction times out.

53. The method, as set forth in claim 49, wherein the memory bank comprises a plurality of pages configured to be open concurrently.

54. A method of dynamically controlling access to pages in a memory component, comprising:

receiving a request from a component coupled to the memory component for a transaction with a page in the memory component;

opening the page in response to the request, wherein a controller is configured to close the page when the transaction is complete; and
reconfiguring the controller to leave the page open when the transaction is complete.

55. The method, as set forth in claim 54, wherein receiving the request comprises receiving more than one request from more than one component coupled to the memory component.

56. The method, as set forth in claim 55, wherein opening the page comprises opening more than one page concurrently in response to receiving more than one request.

57. The method, as set forth in claim 54, wherein the controller is configured to close the page when the transaction is complete based on a condition in a system comprising the memory component.

58. The method, as set forth in claim 54, wherein the controller is configured to leave the page open when the transaction is complete based on an access pattern of the component.

1 / 3




FIG. 1

2 / 3

FIG. 2





FIG. 3

3 / 3




FIG. 4