

**DYNAMIC COMMAND AND/OR ADDRESS MIRRORING SYSTEM AND
METHOD FOR MEMORY MODULES**

CROSS-REFERENCE TO RELATED APPLICATIONS

[001] The present application claims the benefit of the filing date of U.S. Patent Application No. 10/773,518 entitled DYNAMIC COMMAND AND/OR ADDRESS MIRRORING SYSTEM AND METHOD FOR MEMORY MODULES, filed February 5, 2004, which is incorporated herein by reference.

TECHNICAL FIELD

[002] This invention relates to memory modules having memory devices mounted on opposite surfaces of a substrate, and more particularly to memory modules having memory devices on opposite surfaces of a substrate.

BACKGROUND OF THE INVENTION

[003] Semiconductor devices, such as memory devices, are normally in the form of a semiconductor substrate or chip mounted in a hermetically sealed package. An integrated circuit fabricated on the chip is then coupled to terminals that are accessible from outside the package. These externally accessible terminals can assume many forms, such as pins projecting outwardly and then downwardly along opposite sides of the integrated circuit package and terminal pads arranged in a group on the bottom of the integrated circuit package, which is known as a ball grid array, or "BGA" configuration.

[004] Each externally accessible terminal of the integrated circuit package is normally associated with a particular function. For example, in an integrated memory device, a first set of externally accessible terminals are input terminals for respective memory address bits A_0 - A_N , and a second set of externally accessible terminals are input terminals for respective command or status signals C_0 - C_N , such as RAS*, CAS*, and a clock signal. A third set of externally accessible terminals are input/output terminals for respective data bits D_0 - D_N . Finally, a fourth set of externally accessible terminals are reserved for power and ground.

[005] Although integrated circuits are commonly used singly in many applications, other types of integrated circuits are most commonly used in groups. For example, memory devices in general, and dynamic random access memory (“DRAM”) devices in particular, are commonly used in groups as part of memory modules. Memory modules are generally in the form of an insulative substrate, such as a printed circuit board, having several memory devices mounted on one or both surfaces of the substrate. Conductors couple the memory devices to connectors that are generally formed by terminals extending along an edge of the substrate. One common memory module is a single in-line memory module, known as a “SIMM,” which includes a single row of memory devices extending across one or both surfaces of the substrate. Another common memory module is a double in-line memory module, known as a “DIMM,” which includes two rows of memory devices extending across one or both surfaces of the substrate.

[006] A common phenomena associated with memory modules, whether SIMMs, DIMMs or some other variety, is the need for an ever increasing storage capacity. For this reason, the capacity of the memory devices mounted on the substrate, as well as the number of externally accessible terminals needed to address the memory devices, has continuously increased. The need for increased memory capacity also increases the need for memory modules having a larger number of memory devices. As a result, memory devices are now usually mounted on both sides of a memory module substrate, and the spacing between memory devices has continued to decrease. The decreased spacing between memory devices and the increased number of terminals has made it more difficult to route conductors to the externally accessible terminals of the memory devices.

[007] One technique that has been used successfully to route conductors to a large number of terminals of closely spaced memory devices is to use substrates having a large number of layers on which conductors are formed. However, it is relatively expensive to provide substrates having a large number of layers, and a large number of closely spaced layers can result in excessive cross-talk between conductors on different layers and excessive conductor capacitance.

[008] Another technique that has made it easier to route conductors to memory device terminals is mirroring in which the terminals of each memory device

mounted on one surface of the substrate are position directly opposite corresponding terminals of a memory device mounted on the opposite surface of the substrate. This mirroring can occur horizontally, in which corresponding terminals are at the same locations on opposite sides of the mirrored packages, or vertically, in which corresponding terminals are at the same locations above and below a line extending across and bisecting the mirrored packages. In either case, mirroring has the advantage of allowing conductors to extend to a single respective location on the substrate, and to then connect to a respective terminal on each surface of the substrate at that location. Significantly, there is no need to route a conductor coupled to a terminal of an integrated circuit on one surface of the substrate to a different location for coupling to the corresponding terminal of an integrated surface on the other surface of the substrate.

[009] Although memory device mirroring has the advantage of allowing more compact routing of conductors to the memory devices, it is not without some disadvantages. For memory device mirroring to occur, two different integrated circuit packages must be developed so that corresponding terminals of the two packages are mirror images of each other. The two different packages can theoretically use the same integrated circuit chip, but, in practice, this is not always feasible. In particular, it is important that the integrated circuits on one surface of the substrate respond to signals in the same manner as the integrated circuits on the opposite surface of the substrate. For the circuits to respond in the same manner, it is important for the lengths of corresponding signal paths of the two circuits be identical. Not only is it sometimes difficult to route signal lines from a circuit node to either of two different terminals, doing so creates an undesirable stub connection to the signal path between the terminal and the circuit node. This stub connection can produce signal reflections that can degrade the performance of the integrated circuit. For this reason, it can be necessary to fabricate two different integrated circuit chips, which are the mirror images of each other, for placement in the respective mirrored packages. The need to develop and stock two different integrated circuit packages, even if the same chip can be used for both packages, can significantly increase the cost of mirrored integrated circuits.

[010] To alleviate the above-described problems of mirroring integrated circuits, programmable integrated circuits have been developed. With reference to Figure 1, an integrated circuit memory device 10 includes a large number of terminals, although only terminals 12, 14 for RAS and CAS signals are shown. The RAS and CAS signals are horizontally mirrored, as explained above. The terminals 12, 14 are each coupled to a respective input of two multiplexers 16, 18. The output of the multiplexer 16 is coupled to a RAS signal node 20, and the output of the multiplexer 18 is coupled to a CAS signal node 22. The multiplexers 16, 18 are controlled by a signal line coupled to an external terminal 26 at a predetermined location. The terminal 12 is coupled to a first input of the multiplexer 16 and to a second input of the multiplexer 18. The terminal 14 is coupled to a second input of the multiplexer 16 and to a first input of the multiplexer 18. As a result, a low applied to the terminal 26 causes the terminal 12 to be coupled to the RAS signal node 20, and the terminal 14 to be coupled to the CAS signal node 22. A high applied to the terminal 26 causes the terminal 12 to be coupled to the CAS signal node 22, and the terminal 14 to be coupled to the RAS signal node 20.

[011] In operation, two of the integrated circuit memory devices 10a,b are mounted on opposite surfaces of a substrate as shown in Figures 2A and 2B, respectively. As a result, the RAS signal is coupled to the terminal 12 of the memory devices 10a and the terminal 14 of the memory device 10b. The CAS signal is coupled to the terminal 14 of the memory devices 10a and the terminal 12 of the memory device 10b. However, the terminal 26 of the memory device 10a is coupled to ground potential, and the terminal 26 of the memory device 10b is coupled to a supply voltage. Therefore, the multiplexer 16 (Figure 1) couples the RAS signal to the RAS signal node 20 of both memory devices 10a,b, and the multiplexer 18 couples the CAS signal to the CAS signal node 22 of both memory devices 10a,b.

[012] The technique explained with reference to Figures 1 and 2A,B has the advantage of allowing mirroring to occur using a single integrated circuit mounted on opposite sides of a substrate, and avoids many of the above-mentioned disadvantages of using two different integrated circuits. However, mirroring using an internal routing circuit, such as the multiplexers 16, 18 shown in Figure 1, has the

disadvantage of requiring that a routing circuit for each terminal be fabricated on a semiconductor substrate, thereby using area that could be used for the integrated circuit itself. As a result, the use of routing circuits can significantly increase the cost of memory devices, particularly in view of the large number of terminals present in memory devices that each require a routing circuit, as well as the large number of memory devices included in many systems. The routing circuits can also introduce undesirable delays in the coupling of the RAS and CAS signals to their respective nodes 20, 22.

[013] Another problem in routing conductors to memory devices in memory modules occurs when the memory module includes a memory hub or register through which signals are routed to and from the memory devices. As shown in Figure 3, a memory module 30 includes memory hub 32 mounted on a substrate 34. The memory module 30 also includes a plurality of memory devices mounted on the substrate 34, two of which 38, 40 are shown in Figure 3. In the memory module 30 of Figure 3, each signal transmitted and received by the memory hub 32 is transmitted and received on a first set of terminals coupled to the memory device 38 on the left of the substrate 34 and a second set of terminals coupled to the memory device 40 on the right side of the substrate 34. One of the signals transmitted by the memory hub, *i.e.*, the A_0 address bit, is shown in Figure 3, and this address bit is coupled to correspondingly positioned terminals of the memory devices 38, 40. However, since the A_0 terminal is located on the left side of both memory devices 38, 40, the path to the A_0 terminal of the left memory device 38 is longer than the path to the A_0 terminal of the right memory device 40. As a result, the performance of the two memory devices 38, 40 may not be symmetrical. A similar problem exists when coupling signals from between memory devices and a register (not shown) of a registered memory module.

[014] The above-described difficulties incurred in coupling signals to and from integrated circuits, such as memory devices, creates a need for a mirroring technique that allows a single integrated circuit to be mounted on opposite surfaces of a substrate with correspondingly positioned terminals coupled together, and which does not require internal routing circuitry in each memory device.

SUMMARY OF THE INVENTION

[015] A memory module according to one aspect of the invention includes an insulative substrate; on which a plurality of identical memory devices mounted on first and second opposed surfaces of the insulative substrate. The memory devices are mounted on the substrate in a mirrored configuration. As a result, a plurality of terminals of each of the memory devices mounted on the first surface are interconnected to respective, correspondingly positioned terminals of a respective one of the memory devices mounted on the second surface. Address and command signals are coupled to the interconnected terminals. Significantly, either the address signals, the command signals or both the address and command signals are coupled to the interconnected terminals for a plurality of the memory devices in a first configuration if the memory devices mounted on the first surface of the substrate are being accessed. In contrast these signals are coupled to the interconnected terminals for a plurality of the memory devices in a second configuration that is different from the first configuration if the memory devices mounted on the second surface of the substrate are being accessed. The memory devices in one or more memory modules may be coupled to a memory controller that alters the configuration of the signals as indicated above. Alternatively, the configuration of the signals may be altered by either a memory hub or a register mounted on each of the memory modules.

BRIEF DESCRIPTION OF THE DRAWINGS

- [016] Figure 1 is a block diagram of a conventional integrated memory device that can be used in a memory module in mirrored fashion.
- [017] Figures 2A and 2B are terminal diagrams showing the routing of some of the signals for the memory device of Figure 1 when they are mounted on opposite surfaces of a memory module substrate.
- [018] Figure 3 is a plan view of a conventional memory module including a memory hub coupled to memory devices on opposite sides of the memory hub.
- [019] Figure 4 is a block diagram of a computer system including several memory modules according to one embodiment of the invention.

[020] Figure 5 is a block diagram of one embodiment of a memory controller that may be used in the computer system of Figure 4.

[021] Figure 6 is a block diagram of one embodiment of a computer system according to the present invention.

[022] Figure 7 is a block diagram of still another embodiment of a computer system according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[023] A computer system 50 according to one embodiment of the invention is shown in Figure 4. The computer system 50 includes a processor 54 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 54 includes a processor bus 56 that normally includes an address bus, a control bus, and a data bus. The processor bus 56 is typically coupled to a cache memory 58, which, is typically static random access memory ("SRAM") device. Finally, the processor bus 56 is coupled to a system controller 60, which is also sometimes referred to as a bus bridge.

[024] The system controller 60 contains a memory hub controller 62 that is coupled to the processor 54. The memory hub controller 62 is also coupled to several memory modules 64a-n, and the memory modules 64a-n are coupled to each other, through a downstream bus 66 and an upstream bus 68, which couple data, address and/or control signals away from or toward, respectively, the memory hub controller 62. Each of the memory modules 64a-n includes a memory hub 76 mounted on a substrate 78. The memory hub 76 is coupled to several memory devices 80a,b on one side of the memory module 64 through a first set of command, address and data buses 82, and to several memory devices 86a,b on the other side of the memory module 64 through a second set of command, address and data buses 88. The memory devices 80a,b and 86a,b are identical to each other. The memory hub 76 efficiently routes memory requests and responses between the memory hub controller 62 and the memory devices 80, 86. Computer systems employing this architecture can have a higher bandwidth because the processor 54 can access one memory module 64a-n while another memory module 64a-n is responding to a prior memory access. For example, the processor 54 can output write data to one of the

memory modules 64a-n in the system while another memory module 64a-n in the system is preparing to provide read data to the processor 54. The operating efficiency of computer systems using a memory hub architecture can make it more practical to vastly increase data bandwidth of a memory system. A memory hub architecture can also provide greatly increased memory capacity in computer systems.

[025] The system controller 60 also serves as a communications path to the processor 54 for a variety of other components. More specifically, the system controller 60 includes a graphics port that is typically coupled to a graphics controller 90, which is, in turn, coupled to a video terminal 92. The system controller 60 is also coupled to one or more input devices 94, such as a keyboard or a mouse, to allow an operator to interface with the computer system 50. Typically, the computer system 50 also includes one or more output devices 96, such as a printer, coupled to the processor 54 through the system controller 60. One or more data storage devices 98 are also typically coupled to the processor 54 through the system controller 60 to allow the processor 54 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices 98 include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs).

[026] With further reference to Figure 4, the memory devices 80a and 86a are mounted on a first surface 100a of the substrate 78, and the memory devices 80b and 86b are mounted on a second surface 100b of the substrate 78. The memory devices 80a, 86a are preferably mounted directly opposite the memory devices 80b, 86b, respectively, and their adjacent terminals are coupled to each other and to signal lines of the buses 82, 88, respectively. As a result, a terminal located at the upper left-hand corner of one of the memory devices 80a is coupled to a correspondingly positioned terminal located at the upper right-hand corner of the opposing memory device 80b, for example.

[027] The terminals of the memory devices 80, 86 are preferably arranged so that data bus terminals of the memory devices 80a, 86a are coupled to data bus terminals of the memory devices 80b, 86b, respectively, on the opposite surface of the substrate 78. Data signals applied to the interconnected terminals of the memory

devices will therefore be written to the memory devices 80a, 86a mounted on the first surface 100a as data bits that are different from the data bits written to the memory devices 80b, 86b mounted on the second surface 100b. For example, a data signal for data bit D_1 may be written to a memory device 80a as a D_1 bit and written to a memory device 80b as a D_{15} bit. However, when the same data signals are read from the memory devices 80a and 80b, the data signal will be coupled to the signal line for the D_1 data bit even though the data signal was read from the D_{15} bit of the memory device 80b. Therefore, as long as the data bus terminals of the memory devices 80a, 86a are coupled to the data bus terminals of the memory devices 80b, 86b, respectively, data will be properly written to and read from the memory devices 80, 86, assuming that all of the data bus terminals to which the data bits are coupled are associated with the same data strobe signal.

[028] Unlike the data bus terminals of the memory devices 80, 86, the address and control bus terminals of the memory devices 80, 86 are not interchangeable. Therefore, an address signal coupled to a terminal of the memory device 80a cannot be simply coupled to a correspondingly positioned terminal of an opposing memory device 80b. Instead, each address and control signal must be coupled to a specific terminal of each of the memory devices 80, 86 regardless of the location of the memory devices. Rather than using the conventional approaches of using different memory devices having different terminal configurations or routing signals to the proper signal nodes, the address and control signals are coupled to the proper terminals of the memory devices 80, 86 by the memory hub 76 coupling different signals to the same lines of the buses 82, 88 depending upon which memory device 80, 86 is being accessed. For example, if the A_0 address bit of the memory device 80a is connected to the A_{10} address bit of the opposing memory device 80b, the memory hub 76 will couple an address signal for bit A_0 to a specific signal line of the bus 82 if the memory device 80a is being addressed, and it may couple an address signal for bit A_{10} to the same signal line of the bus 82 if the memory device 80b is being addressed. Conversely, the memory device will couple an address signal for bit A_0 to one signal line of the bus 82 if the memory device 80a is being addressed, and it will couple the address signal for bit A_0 to another signal line of the bus 82 if the memory device 80b is being addressed. Different address

and control signals are also coupled to the same lines of the bus 88 depending on whether memory device 86a or 86b is being accessed.

[029] The memory devices 80, 86 may be either of two types of memory devices each of which operates in either of two modes. Furthermore, the memory devices 80, 86 have different terminal assignments depending on which mode of operation is used by the memory devices 80, 86. For example, as is known to one skilled in the art, DRAM memory devices operating in one double data rate (“DDR”) mode known as the “DDR2” mode use one set of terminal assignments, and, DRAM memory devices operating in another DDR mode known as the “DDR3” mode, use another set of terminal assignments. For example, the same terminal to which an A_5 address bit is coupled to memory devices operating in the DDR2 mode will receive an A_{12} address bit for memory devices operating in the DDR3 mode. In each case, the memory hub 76 will couple the proper signal to each signal line of each of the buses 82, 88 depending on: (1) whether the memory devices 80, 86 are of the type that operate in the DDR2 mode or the DDR3 mode; (2) whether the memory devices 80, 86 are to the left of the memory hub 76 and are thus coupled to the hub 76 through the bus 82 or to the right of the memory hub 76 and are thus coupled to the hub 76 through the bus 88; and (3) whether the memory devices 80, 86 are mounted on the first surface 100a of the substrate 78 or the second surface 100b of the substrate 78.

[030] In one embodiment of the invention, the memory hub 76 will couple signal to the memory devices 80, 86 according to Table 1 below in which A_0 - A_{16} are row and column addresses, BA_0 - BA_3 are bank addresses, CASZ is an active low column address strobe signal, CSZ0 and CSZ1 are active low chip select signals to select memory devices 80, 86 on the first surface 100a and second surface 100b, respectively, ODT0 and ODT1 are signals, RASZ is an active low row address strobe signal and WEZ is an active low write enable signal.

Table 1

Hub Left Terminal Location	Hub Right Terminal Location	82a, 86a DDR2 CS ₀	82b, 86b DDR2 CS ₁	82a, 86a DDR3 CS ₀	82b, 86b DDR3 CS ₁
----------------------------------	-----------------------------------	-------------------------------------	-------------------------------------	-------------------------------------	-------------------------------------

C ₅	C ₂₇	A ₀	A ₁₀	A ₃	A ₀
C ₁₁	B ₂₁	A ₁	A ₂	A ₆	A ₅
C ₆	C ₂₆	A ₂	A ₁	BA ₃	BA ₂
C ₁₂	B ₂₀	A ₃	A ₄	A ₂	A ₁
B ₅	B ₂₇	A ₄	A ₃	A ₁	A ₂
B ₁₁	A ₂₁	A ₅	A ₆	A ₁₂	A ₁₃
B ₆	B ₂₆	A ₆	A ₅	A ₅	A ₆
B ₁₂	C ₁₉	A ₇	A ₈	A ₇	A ₈
A ₅	B ₂₉	A ₈	A ₇	A ₈	A ₇
A ₁₃	A ₁₉	A ₉	A ₁₁	A ₉	A ₁₁
D ₁₂	C ₂₁	A ₁₀	A ₀	BA ₂	BA ₃
A ₄	A ₂₉	A ₁₁	A ₉	A ₁₁	A ₉
A ₁₂	A ₂₀	A ₁₂	A ₁₃	A ₄	A ₁₄
A ₆	A ₂₇	A ₁₃	A ₁₂	A ₁₃	A ₁₂
A ₁₁	A ₂₂	A ₁₄	A ₁₅	A ₁₅	A ₁₅
A ₇	A ₂₈	A ₁₅	A ₁₄	A ₁₄	A ₄
A ₁₀	A ₂₃	A ₁₆	A ₁₆	A ₁₆	A ₁₆
D ₇	D ₂₇	BA ₀	BA ₀	BA ₁	BA ₀
E ₁₃	D ₂₀	BA ₁	CASZ	A ₁₀	WEZ
E ₇	E ₂₅	BA ₂	BA ₂	CSZ ₁	CSZ ₁
F ₁₂	F ₂₂	BA ₃	BA ₃	ODT ₁	ODT ₁
D ₈	D ₂₆	CASZ	BA ₁	CASZ	RASZ
D ₆	D ₂₈	CSZ ₀	CSZ ₀	WEZ	A ₁₀
D ₁₁	D ₂₁	CSZ ₁	CSZ ₁	BA ₀	BA ₁
C ₉	E ₂₈	ODT ₀	ODT ₀	ODT ₀	ODT ₀
D ₁₃	C ₂₀	ODT ₁	ODT ₁	A ₀	A ₃
F ₉	E ₂₂	RASZ	WEZ	CSZ ₀	CSZ ₀
E ₁₀	D ₂₂	WEZ	RASZ	RASZ	CASZ

[031] In one embodiment, the memory hub 76 has externally accessible terminals arranged in a grid, and the location of each terminal is designated by a

letter designating the row of the terminal and a number designating the column of the terminal. For example, A1 is a terminal located at the upper left hand corner of the memory hub 76, and F29 is a terminal located at the lower right hand corner of the memory hub 76. The locations of these terminals that are coupled to the memory devices 80 through the bus 82 are listed in the first column of Table 1, and the locations of these terminals that are coupled to the memory devices 86 through the bus 88 are listed in the second column of Table 1. The signals applied to the terminals of the memory hub 76 when the memory devices 82a, 86a on the first surface 100a of the substrate 78 are being accessed and the devices 82a, 86a are of the type operating in the DDR2 mode are listed in the third column of Table 1. The signals applied to the terminals of the memory hub 76 when the memory devices 82b, 86b on the second surface 100b are being accessed and the devices 82b, 86b are of the type operating in the DDR2 mode are listed in the fourth column of Table 1. The signals applied to the terminals of the memory hub 76 when the memory devices 82a, 86a are of the type operating in the DDR3 mode and are being accessed are listed in the fifth column of Table 1. Finally, the signals applied to the terminals of the memory hub 76 when the memory devices 82b, 86b are of the type operating in the DDR3 mode and are being accessed are listed in the sixth column of Table 1.

[032] For example, the C5 terminal of the memory hub 76, which is coupled to a signal line of the bus 82, receives the A_0 address bit when the memory devices 80a are being accessed and they are of the type operating in the DDR2 mode, and it receives the A_3 address bit when the memory devices 80a are being accessed and they are of the type operating in the DDR3 mode. If the memory devices 86a are being accessed, the C27 terminal of the memory hub 76, which is coupled to a signal line of the bus 88, receives the A_0 address bit when the memory devices 86a are being accessed and they are of the type operating in the DDR2 mode, and it receives the A_3 address bit when the memory devices 80a are being accessed and they are of the type operating in the DDR3 mode. If the memory devices 80b are being accessed, the C5 terminal of the memory hub 76, which is coupled to a signal line of the bus 82, receives the A_{10} address bit when the memory devices 80b are of the type operating in the DDR2 mode, and it receives the A_0 address bit when the memory devices 80b are being accessed and they are of the

type operating in the DDR3 mode. If the memory devices 86b are being accessed, the C27 terminal of the memory hub 76, which is coupled to a signal line of the bus 88, receives the A_{10} address bit when the memory devices 80b are of the type operating in the DDR2 mode, and it receives the A_0 address bit when the memory devices 80b are being accessed and they are of the type operating in the DDR3 mode.

[033] The routing of signals to the terminals of the memory devices 80, 86 based on whether they are of the type operating in the DDR2 or the DDR3 mode is static since it does not change during the operation of the memory devices 80, 86. However, the routing of signals from either the left set of terminals A1-F13 of the memory hub 76 to the memory devices 80 or from right set of terminals A19-29 of the memory hub 76 and to the memory devices 86 is dynamic, and therefore changes rapidly while the memory devices 80, 86 are being accessed.

[034] Although the address and most of the control signals must be dynamically mirrored, in one embodiment some control signals can be applied to multiple symmetrically positioned terminals of the memory device 80, 86 so that dynamic mirroring is not required. For example, a first clock enable CKE signal is applied to a terminal of the memory device 80, and a second clock enable CKE signal is applied to a terminal of the memory device 86 regardless of whether the memory devices 80, 86 are of the type operating in the DDR2 or the DDR3 mode. Regardless of which memory device 80, 86 is being addressed, the addressed memory device will receive a CKE signal at the appropriate terminal. The signals for which dynamic mirroring is not required, as well as the locations of the terminals of the memory hub 76 to which those signals are applied in one embodiment, are listed in Table 2, below.

Table 2

Hub Left Terminal Location	Hub Right Terminal Location	82a, 86a DDR2 CS_0	82b, 86b DDR2 CS_1	82a, 86a DDR3 CS_0	82b, 86b DDR3 CS_1
D ₉	E ₁₉	CKE ₀	CKE ₀	CKE ₀	CKE ₀
F ₆	E ₂₉	CKE ₁	CKE ₁	CKE ₁	CKE ₁
F ₈	F ₂₆	CLK ₀	CLK ₀	CLK ₀	CLK ₀

G ₁₁	G ₂₄	CLK ₁	CLK ₁	CLK ₁	CLK ₁
F ₇	F ₂₅	CLKZ ₀	CLKZ ₀	CLKZ ₀	CLKZ ₀
G ₁₂	G ₂₃	CLKZ ₁	CLKZ ₁	CLKZ ₁	CLKZ ₁
F ₁₃				RESET	RESET

[035] The memory hub 76 used in the memory modules 64a-n of Figure 4 includes a memory controller, and one embodiment of a memory controller 200 that can be used is shown in Figure 5. Briefly, the memory controller 200 receives high-level macro commands, such as ACTIVATE ROW, COLUMN, and PRECHARGE, converts these commands to DRAM commands, schedules the DRAM commands for outputting at the proper time, and routes the scheduled DRAM commands to the correct terminals of the memory hub 76, as previously explained. With reference to Figure 5, the memory controller 200 includes a command queue 204 that receives the high-level macro commands from the memory hub controller 62 (Figure 4). The Command Queue 204 translates the received macro commands to DRAM command signals, such as RASZ, CASZ, WEZ, etc., places the command signals in a queue in the order that their corresponding macro commands were received, and then outputs the commands in the proper order. The Command Queue 204 also receives address signals, which are placed in a queue in the order that they were received, and they are subsequently output in the proper order. The command and address signals that are output from the Command Queue 204 are the command and address signals listed in Tables 1 and 2.

[036] The DRAM command signals and address signals from the Command Queue 204 are applied to a Command Scheduler 210, which spaces the command and address signals apart from each other with the proper delay. The delay is measured in periods of a clock CLK signal, which is also applied to the Command Scheduler 210. For example, the Command Scheduler 210 might schedule the CASZ to be output three clock periods after the RASZ signals was output from the Command Scheduler 210. When the command and address signals are output from the Command Scheduler 210, they are stored in a micro Command Shifter 214, which is basically a shift register that is driven by the CLK signal. The command and address signals are stored in the Command Shifter 214 in the proper

order and with the proper spacing because they were shifted into the Command Shifter 214 from the Command Scheduler 210 in that manner.

[037] When the properly timed and ordered command and address signals are shifted out of the Command Shifter 214 responsive to the CLK signal, they are applied to one input of a Multiplexer 220. The Multiplexer 220 also receives the DRAM command and address signals at a second input directly from the Command Queue 204. In operation, the Multiplexer 220 couples the command and address signals to its output directly from the Command Shifter 214 as these signals are received when the memory devices 80, 86 have been idle. As a result, the latency penalty that would be incurred in being coupled through the Command Scheduler 210 and the Micro Command Shifter 214 is avoided. After the initial command and address signals have been coupled to the output of the Multiplexer 220, the Multiplexer 220 selects the output of the Micro Command Shifter 214 for coupling to its output.

[038] The command and address signals at the output of the Multiplexer 220 are applied to Signal Swap Multiplexers 230 that are controlled by a CONFIG command coupled to the Multiplexers 230 through a bus 232. The CONFIG command indicates whether the memory devices 80, 86 are of the type operating in the DDR2 or DDR3 modes. The Multiplexers 230 are also controlled by CSZ₀ and CSZ₁ signals, which indicate whether the memory devices 80a, 86a on the first surface 100a or the memory devices 80b, 86b on the second surface 100b are being accessed. The Signal Swap Multiplexers 230 can be implemented using a large number of individual multiplexers arranged in a matrix in a manner that will be apparent to one skilled in the art. The Signal Swap Multiplexers 230 route the command and address signals to the terminals of the memory hub 76 as shown in Tables 1 and 2.

[039] The properly routed and timed command and address signals are coupled from the Signal Swap Multiplexers 230 to Ring Buffers 240, 242. The Ring Buffer 240 is coupled through the bus 82 (Figure 4) to the memory devices 80 on the left hand side of the memory modules 64, and the Ring Buffer 242 is coupled through the bus 88 to the memory devices 86 on the right hand side of the memory modules 64. Each of the Ring Buffers 240, 242 is basically a first-in, first-out buffer

that is driven by one of the clock signals CLK_0 , CLK_1 , $CLKZ_0$ or $CLKZ_1$ that is applied to the memory devices 80, 86. The command and address signals are therefore shifted from the Ring Buffers 240, 242 and applied to the terminals of the memory hub 76 in synchronism with the operation of the memory devices 80, 86.

[040] Although the memory controller 200 can be used in the memory hub 76 in the computer architecture shown in Figure 4, it can also be used as a stand-alone memory controller in the computer system 300 shown in Figure 6. The computer system 300 uses many of the same components that are used in the computer system 50 of Figure 4. Therefore, in the interests of brevity, these components have been provided with the same reference numerals, and an explanation of their operation will not be repeated. The computer system 300 differs from the computer system 50 by using memory modules 310 that do not contain a memory hub. Instead, the memory devices 80, 86 in each memory module 310 are coupled directly to the memory controller 200 in the system controller 60 through a memory bus 320. The memory controller 200 applies the command and address signals shown in Tables 1 and 2 to the signal lines of the bus 320 depending on which memory devices 80a, 86a, 80b or 86b are being accessed and whether the memory devices 80, 86 are of the type operating in either the DDR2 mode or the DDR3 mode, as explained above and as shown in Table 1.

[041] Another embodiment of a computer system 400 according to the invention is shown in Figure 7. The computer system 400 also uses many of the same components that are used in the computer systems 50 and 300 of Figure 4 and 6, respectively, and an explanation of the operation of these components will not be repeated. The computer system 400 differs from the computer system 300 by using memory modules 410 having registers 420 coupling the command and address signals to the memory devices 80, 86. Registered memory modules are well known in the art. The computer system 400 uses a memory controller 430 in the system controller 60 that may operate in the same manner as the memory controller 200 to apply the command and address signals to the registers 420 depending on which memory devices 80a, 86a, 80b or 86b are being accessed and whether the memory devices 80, 86 are of the type operating in either the DDR2 mode or the DDR3 mode. Alternatively, the memory controller 420 may operate in a conventional

manner, and circuitry like the Signal Swap Multiplexers 230 shown in Figure 5 may be included in each register 420 to apply the command and address signals to the memory devices 80, 86 depending on which memory devices 80a, 86a, 80b or 86b are being accessed and whether the memory devices 80, 86 are of the type operating in either the DDR2 mode or the DDR3 mode. Other variations will be apparent to one skilled in the art.

[042] The various embodiments of the inventions have the advantage of routing the proper command and address signals to the memory devices 80, 86 at the memory controller level rather than at each of the memory device level. As a result, the number of circuits needed to perform this function is markedly less than if such circuitry was in each of the memory devices 80, 86. The various embodiments thus allow a mirroring with single memory device used on both surfaces of the substrate and on both the left and right sides of the substrate. From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, it will be understood by one skilled in the art that various modifications may be made without deviating from the spirit and scope of the invention. For example, instead of using registered memory modules, circuitry like the Signal Swap Multiplexers 230 could simply be used on each memory module. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

What is claimed is:

1. A method of accessing a plurality of memory devices in which a plurality of terminals of a first of the memory devices are interconnected with a corresponding plurality of terminals of a second of the memory devices in a manner that causes the first and second memory devices to function differently responsive to respective address or control signals applied to the interconnected terminals, the method comprising:

if the first memory device is being accessed, applying control or address signals to the interconnected terminals according to a first set of terminal assignments; and

if the second memory device is being accessed, applying control or address signals to the interconnected terminals according to a second set of terminal assignments that is at least in part different from the first set of terminal assignments.

2. The method of claim 1 wherein the parts of the first and second terminal assignments that are different comprise address terminal assignments.

3. The method of claim 1 wherein the parts of the first and second terminal assignments that are different comprise control terminal assignments.

4. The method of claim 1, further comprising applying data signals to the interconnected terminals, and wherein the data signals are applied to the interconnected terminals according to a common set of terminal assignments regardless of whether the first memory device or the second memory device is being accessed

5. The method of claim 1 wherein the memory devices comprise dynamic random access memory devices.

6. A method of applying address and control signals to a plurality of identical memory devices in which a plurality of terminals of a first of the memory devices are interconnected with a corresponding plurality of terminals of a second of the memory devices in mirrored configuration, the method comprising:

applying a set of control signals or a set of address signals to the interconnected terminals in a first arrangement if the first memory device is being accessed; and

applying a set of control signals or a set of address signals to the interconnected terminals in a second arrangement if the second memory device is being accessed, the second arrangement being different from the first arrangement.

7. The method of claim 6 wherein the act of applying a set of control signals or a set of address signals to the interconnected terminals in first and second arrangements that are different from each other comprise applying a set of control signals to the interconnected terminals in first and second arrangements that are different from each other.

8. The method of claim 6 wherein the act of applying a set of control signals or a set of address signals to the interconnected terminals in first and second arrangements that are different from each other comprise applying a set of address signals to the interconnected terminals in first and second arrangements that are different from each other.

9. The method of claim 6 wherein the act of applying a set of control signals or a set of address signals to the interconnected terminals in first and second arrangements that are different from each other comprise applying a set of

both address signals and control signals to the interconnected terminals in first and second arrangements that are different from each other.

10. The method of claim 6, further comprising applying data signals to the interconnected terminals in a common arrangement regardless of whether the first memory device is being accessed or the second memory device is being accessed.

11. The method of claim 6 wherein the memory devices comprise dynamic random access memory devices.

12. A method of applying address or control signals to a plurality of identical memory devices mounted on first and second surfaces of memory module substrate in a mirrored configuration so that a plurality of terminals of each of the memory devices mounted on the first surface are interconnected to respective, correspondingly positioned terminals of a respective one of the memory devices mounted on the second surface, the method comprising:

coupling address or control signals to the interconnected terminals for a plurality of the memory devices in a first configuration if the memory devices mounted on the first surface of the substrate are being accessed; and

coupling address or control signals to the interconnected terminals for a plurality of the memory devices in a second configuration that is different from the first configuration if the memory devices mounted on the second surface of the substrate are being accessed.

13. The method of claim 12 wherein the acts of coupling address or control signals to the interconnected terminals in the first or second configuration comprises:

coupling address or control signals to the memory module; and

within the memory module, re-arranging the signals coupled to the memory module to either the first configuration or the second configuration prior to coupling the address or control signals to the interconnected terminals.

14. The method of claim 13 wherein the act of re-arranging the signals coupled to the memory module to either the first configuration or the second configuration comprises re-arranging the signals in a memory hub that is structured to independently access the memory devices.

15. The method of claim 12 wherein the act of coupling address or control signals to the interconnected terminals for a plurality of the memory devices in a first or second configuration comprises coupling address signals to the interconnected terminals in a first or second configuration.

16. The method of claim 12 wherein the act of coupling address or control signals to the interconnected terminals for a plurality of the memory devices in a first or second configuration comprises coupling control signals to the interconnected terminals in a first or second configuration.

17. The method of claim 12, further comprising applying data signals to the interconnected terminals in a common configuration regardless of whether the memory devices being accessed are mounted on the first surface of the substrate or the second surface of the substrate.

18. The method of claim 12 wherein the memory devices comprise dynamic random access memory devices.

19. A memory module, comprising:
an insulative substrate;

a plurality of identical memory devices mounted on first and second opposed surfaces of the insulative substrate, the memory devices being mounted on the substrate in a mirrored configuration so that a plurality of terminals of each of the memory devices mounted on the first surface are interconnected to respective, correspondingly positioned terminals of a respective one of the memory devices mounted on the second surface; and

a memory access device mounted on the substrate, the memory access device having a plurality of terminals that are coupled through the substrate conductors to respective ones of the interconnected terminals, the memory access device being operable to receive a memory request and, in response, to couple address and control signals to the interconnected terminals for a plurality of the memory devices, the address or control signals being coupled to the interconnected terminals in a first configuration if the memory devices mounted on the first surface of the substrate are being accessed, and the address or control signals being coupled to the interconnected terminals in a second configuration that is different from the first configuration if the memory devices mounted on the second surface of the substrate are being accessed.

20. The memory module of claim 19 wherein the memory access device is centrally position on the insulative substrate, and wherein the memory devices are positioned to both sides of the memory access device, the memory access device being operable to couple respective sets of address or control signals in the first or second configuration to the memory devices on each side of the memory access device.

21. The memory module of claim 19 wherein the memory access device comprises a memory hub that is structured to generate the address and control signals to access the memory devices responsive to the memory requests.

22. The memory module of claim 19 wherein the memory hub comprises:

a command queue that is operable to receive the memory requests, the memory queue further being operable to convert the memory requests into respective sets of command and address signals and to output the command and address signals in the order that the respective memory requests were received;

a command scheduler coupled to the command queue to receive the command and address signals from the command queue, the command scheduler arranging the timing of the command and address signals;

a micro command shifter coupled to receive the command and address signals from the command scheduler after the timing of the command and address signals have been arranged, the micro command shifter being operable to output the command and address signals in synchronism with the operation of the memory devices; and

a multiplexer coupled to the micro command shifter to receive the command signals or the address signals from the micro command shifter, the multiplexer being operable to arrange the command or address signals in either the first configuration or the second configuration depending on whether the memory devices on the first surface or the memory devices on the second surface are being accessed, the multiplexer being operable to couple the command or address signals in either the first configuration or the second configuration to the interconnected terminals.

23. The memory module of claim 22, further comprising a ring buffer coupled between the multiplexer and the memory devices.

24. The memory module of claim 19 wherein the memory access device comprises a register that is structured to receive and store address and control signals forming each of the memory requests, and couple the stored address and control signals to the memory devices.

25. The memory module of claim 19 wherein the memory devices each comprise a dynamic random access memory device.

26. A processor-based system, comprising:
a processor having a processor bus;
a system controller coupled to the processor bus, the system controller having a peripheral device port, the system controller further comprising a controller coupled to a system memory port;
at least one input device coupled to the peripheral device port of the system controller;
at least one output device coupled to the peripheral device port of the system controller;
at least one data storage device coupled to the peripheral device port of the system controller; and
a memory module coupled to the system memory port of the system controller, the memory module comprising:
an insulative substrate;
a plurality of identical memory devices mounted on first and second opposed surfaces of the insulative substrate, the memory devices being mounted on the substrate in a mirrored configuration so that a plurality of terminals of each of the memory devices mounted on the first surface are interconnected to respective, correspondingly positioned terminals of a respective one of the memory devices mounted on the second surface; and
a memory access device mounted on the substrate, the memory access device having a plurality of terminals that are coupled through the substrate conductors to respective ones of the interconnected terminals, the memory access device being coupled to the controller to receive a memory request from the controller and, in response, to couple address and control signals to the interconnected terminals for a plurality of the memory devices, the address or control signals being coupled to the

interconnected terminals in a first configuration if the memory devices mounted on the first surface of the substrate are being accessed, and the address or control signals being coupled to the interconnected terminals in a second configuration that is different from the first configuration if the memory devices mounted on the second surface of the substrate are being accessed.

27. The processor-based system of claim 26 wherein the memory access device is centrally position on the insulative substrate, and wherein the memory devices are positioned to both sides of the memory access device, the memory access device being operable to couple respective sets of address or control signals in the first or second configuration to the memory devices on each side of the memory hub.

28. The processor-based system of claim 26 wherein the memory access device comprises a memory hub that is operable to generate the address and command signals in response to a higher level memory request from the controller.

29. The processor-based system of claim 28 wherein the memory hub comprises:

- a command queue that is operable to receive the memory requests, the memory queue further being operable to convert the memory requests into respective sets of command and address signals and to output the command and address signals in the order that the respective memory requests were received;

- a command scheduler coupled to the command queue to receive the command and address signals from the command queue, the command scheduler arranging the timing of the command and address signals;

- a micro command shifter coupled to receive the command and address signals from the command scheduler after the timing of the command and address signals have been arranged, the micro command shifter being operable to

output the command and address signals in synchronism with the operation of the memory devices; and

a multiplexer coupled to the micro command shifter to receive the command signals or the address signals from the micro command shifter, the multiplexer being operable to arrange the command or address signals in either the first configuration or the second configuration depending on whether the memory devices on the first surface or the memory devices on the second surface are being accessed, the multiplexer being operable to couple the command or address signals in either the first configuration or the second configuration to the interconnected terminals.

30. The processor-based system of claim 29, further comprising a ring buffer coupled between the multiplexer and the memory devices.

31. The processor-based system of claim 26 wherein the memory devices each comprise a dynamic random access memory device.

32. The processor-based system of claim 26 wherein the memory access device comprises a register that is operable to store the address and command signals received from the controller, and to subsequently couple the stored address and command signals to the memory devices.

33. A processor-based system, comprising:
a processor having a processor bus;
a system controller coupled to the processor bus, the system controller having a peripheral device port and a system memory port;
at least one input device coupled to the peripheral device port of the system controller;
at least one output device coupled to the peripheral device port of the system controller;

at least one data storage device coupled to the peripheral device port of the system controller;

at least one memory module coupled to the system memory port of the system controller, the memory module comprising:

an insulative substrate; and

a plurality of identical memory devices mounted on first and second opposed surfaces of the insulative substrate, the memory devices being mounted on the substrate in a mirrored configuration so that a plurality of terminals of each of the memory devices mounted on the first surface are interconnected to respective, correspondingly positioned terminals of a respective one of the memory devices mounted on the second surface, the interconnected terminals of the at least one memory module being coupled to the to the system memory port of the system controller; and

a memory controller coupled to the system memory port of the system controller, the memory controller being operable to couple address and control signals to the interconnected terminals of the at least one memory module in a first configuration if the memory devices mounted on the first surface of the substrate of the at least one memory module are being accessed, and the address or control signals being coupled to the interconnected terminals of the at least one memory module in a second configuration that is different from the first configuration if the memory devices mounted on the second surface of the substrate of the at least one memory module are being accessed.

34. The processor-based system of claim 33 wherein the memory controller comprises:

a command queue that is operable to receive the memory requests, the memory queue further being operable to convert the memory requests into respective sets of command and address signals and to output the command and address signals in the order that the respective memory requests were received;

a command scheduler coupled to the command queue to receive the command and address signals from the command queue, the command scheduler arranging the timing of the command and address signals;

a micro command shifter coupled to receive the command and address signals from the command scheduler after the timing of the command and address signals have been arranged, the micro command shifter being operable to output the command and address signals in synchronism with the operation of the memory devices; and

a multiplexer coupled to the micro command shifter to receive the command signals or the address signals from the micro command shifter, the multiplexer being operable to arrange the command or address signals in either the first configuration or the second configuration depending on whether the memory devices on the first surface or the memory devices on the second surface are being accessed, the multiplexer being operable to couple the command or address signals in either the first configuration or the second configuration to the interconnected terminals of the at least one memory module.

35. The processor-based system of claim 34, further comprising a ring buffer coupled between the multiplexer and the interconnected terminals of the memory devices.

36. The processor-based system of claim 33 wherein the memory devices each comprise a dynamic random access memory device.

37. A memory hub having an input port and a plurality of output terminals, the memory hub being responsive to a memory request received at the input port to couple address and control signals to the output terminals, the address or control signals being coupled to the output terminals in a first configuration if the memory request is directed to a first memory device, the address or control signals being coupled to the output terminals in a second configuration if the memory

request is directed to a second memory device, the second configuration being different from the first configuration.

38. The memory hub of claim 37 wherein the memory hub comprises:

- a command queue that is operable to receive the memory requests, the memory queue further being operable to convert the memory requests into respective sets of command and address signals and to output the command and address signals in the order that the respective memory requests were received;

- a command scheduler coupled to the command queue to receive the command and address signals from the command queue, the command scheduler arranging the timing of the command and address signals;

- a micro command shifter coupled to receive the command and address signals from the command scheduler after the timing of the command and address signals have been arranged, the micro command shifter being operable to output the command and address signals in synchronism with the operation of the memory devices; and

- a multiplexer coupled to the micro command shifter to receive the command signals or the address signals from the micro command shifter, the multiplexer being operable to arrange the command or address signals in either the first configuration or the second configuration depending on whether the memory devices on the first surface or the memory devices on the second surface are being accessed, the multiplexer being operable to couple the command or address signals in either the first configuration or the second configuration to the interconnected terminals.

39. The memory hub of claim 38, further comprising a ring buffer coupled between the multiplexer and the memory devices.

40. The memory hub of claim 37 wherein the command and address signals generated by the memory hub comprise dynamic random access memory command and address signals.