MULTIPLE PROCESSOR SYSTEM AND METHOD INCLUDING MULTIPLE MEMORY HUB MODULES

TECHNICAL FIELD

[001] This invention relates to computer systems, and, more particularly, to a computer system having several processors or other memory access devices that can be coupled to several memory hub modules in a variety of configurations.

BACKGROUND OF THE INVENTION

[002] Computer systems use memory devices, such as dynamic random access memory ("DRAM") devices, to store instructions and data that are accessed by a processor. These memory devices are normally used as system memory in a computer system. In a typical computer system, the processor communicates with the system memory through a processor bus and a memory controller. The processor issues a memory request, which includes a memory command, such as a read command, and an address designating the location from which data or instructions are to be read or to which data or instructions are to be written. The memory controller uses the command and address to generate appropriate command signals as well as row and column addresses, which are applied to the system memory. In response to the commands and addresses, data are transferred between the system memory and the processor. The memory controller is often part of a system controller, which also includes bus bridge circuitry for coupling the processor bus to an expansion bus, such as a PCI bus.

[003] Although the operating speed of memory devices has continuously increased, this increase in operating speed has not kept pace with increases in the operating speed of processors. As a result, the data bandwidth between a processor and memory devices to which it is coupled is significantly lower than the data bandwidth capabilities of the processor. The data bandwidth between the processor and memory devices is limited to a greater degree by the even lower data bandwidth between the processor and the memory devices.

[004] In addition to the limited bandwidth between processors and memory devices, the performance of computer systems is also limited by latency problems that increase the time required to read data from the memory devices. More specifically, when a memory device read command is coupled to a memory device, such as a synchronous DRAM ("SDRAM") device, the read data are output from the SDRAM device only after a delay of several clock periods. Therefore, although SDRAM devices can synchronously output burst data at a high data rate, the delay in initially providing the data can significantly slow the operating speed of a computer system using such SDRAM devices.

[005] One approach to alleviating the memory latency problem is illustrated in Figure 1. As shown in Figure 1, a computer system 10 includes a processor 14 coupled to several memory modules 20a-f, although a lesser or greater number of memory modules 20 may be used. Each of the memory modules 20 includes a memory hub 24 coupled to several memory devices 28, which may be SDRAM devices. The memory modules 20 are shown in Figure 1 as being coupled to the processor 14 and to each other 20 through unidirectional input buses 30 and unidirectional output buses 38. However, it will be understood that the memory modules 20 may be coupled to the processor 14 and to each other by bi-directional buses (not shown).

[006] The memory modules 20 are shown in Figure 1 as being coupled in a point-to-point arrangement in which each bus 30, 38 is coupled only between two points. However, other bus system may alternatively be used. For example, a switched bus system as shown in Figure 2A, a shared bus system as shown in Figure 2B, or some other bus system may also be used. The switched bus system shown in Figure 2A includes a processor 40 coupled to a switching circuit 42. The switching circuit 42 is coupled to several memory modules 44a-d, a graphics processor 46 and an I/O device 48. In operation, the switching circuit 42 couples the processor 40 to either one of the memory modules 44a-d, the graphics processor 46 or the I/O device 48. The shared bus system shown in Figure 2B includes a processor 50 coupled to several memory modules 54a-c through a shared bus system 58.

[007] Any of the above-described architectures may also be used to couple multiple processors to multiple memory modules. For example, as shown in Figure 3, a pair of a processors 60, 62 are coupled through respective bi-directional bus systems 64

to respective sets of memory modules 66a-e, 68a-e. Each of the memory modules 66a-e, 68a-e includes a memory hub 24 coupled to several memory devices 28.

[008] A memory hub architecture as shown in Figures 1 and 3 can provide performance that is far superior to architectures in which a processor is coupled to several memory devices, either directly or through a system or memory controller. However, they nevertheless suffer from several limitations. For example, the architecture shown in Figure 1 does not provide a great deal of flexibility in the manner in which the processor 14 can access the memory modules 20a-f. If, for example, the buses 30-38 include a 32-bit data bus, all accesses to the memory modules 20a-f will be in 32-bit double words even if a lesser number a data bits are being read from or written to the memory modules 20a-f.

[009] The flexibility of the architectures shown in Figures 1 and 3 are also limited in other respects. For example, the architecture shown in Figure 3 does not provide a great deal of flexibility in the manner in which the processors 60, 62 can access the memory modules 66a-e, 68a-e, respectively. Although the processor 60 can access any of the memory modules 66a-f, and the processor 62 can access any of the memory modules 68a-e, the processor 60 cannot access any of the memory modules 68a-e nor can the processor 62 access any of the memory modules 66a-e. As a result, if the processor 60 writes sufficient data to the memory modules 66a-e to reach the storage capacity of the modules 66a-e, the processor 60 will be unable to store any further data even though there may be substantial unused capacity in the memory modules 68a-e. Finally, the memory modules 66, 68 cannot be used to allow the processors 60, 62 to communicate with each other.

[010] Another limitation of the memory architectures shown in Figures 1 and 3 is the relatively high latency that the processors 14, 60, 62 incur in accessing their respective memory modules 20, 66, 68. Insofar as each memory module is accessed through any memory module that is between it and the processor, substantial delays may be incurred in coupling address, data and control signals through the intervening memory modules. Further, if any of the memory modules 20, 66, 68 becomes defective, the memory modules that must be accessed through the defective memory module become unusable.

[011] There is therefore a need for a memory system architecture that is relatively fault-intolerant, that provides relatively low latency memory accesses, an that allows multiple processor to have a great deal of flexibility in the manner in which they access hub-based memory modules.

SUMMARY OF THE INVENTION

[012] A memory system includes a plurality of memory requestors coupled to a first rank of memory modules. The memory modules in the first rank each include a first set of memory ports corresponding in number to the number of memory requestors. Each of the memory ports in the first rank is coupled to a respective one of the memory requestors. The memory modules in the first rank further include a second set of memory ports. The memory system also includes a second rank of memory modules each of which has at least one memory port coupled to at least one memory module in the first rank through a memory port in the second set. Each of the memory modules in the first and second ranks include a plurality of memory devices and a memory hub coupled to the memory devices and to the memory ports in the first set and any second set. The memory hub preferably includes a plurality of memory controllers coupled to respective memory devices in the module, a plurality of link interfaces each of which is coupled to either one of the memory requestors or another module, and a cross bar switch having a first plurality of switch ports coupled to respective link interfaces and a plurality of memory ports coupled to respective memory controllers. The cross bar switch is operable to selectively couple each of the link interfaces to any one of the memory controllers.

BRIEF DESCRIPTION OF THE DRAWINGS

[013] Figure 1 is a block diagram of a conventional processor-based electronic system including several memory modules each of which include a memory hub coupled to several memory devices.

[014] Figures 2A and 2B are block diagrams showing various conventional architectures for coupling memory modules to multiple processors.

[015] Figure 3 is a block diagram of a conventional processor-based electronic system including multiple processors coupled to respective sets of memory modules each of which include a memory hub coupled to several memory devices.

[016] Figure 4 is a block diagram of a processor-based system coupled to several memory modules according to one embodiment of the invention.

[017] Figure 5 is a block diagram of a memory module that can be used in the processor-based systems of Figures 4 and 6.

[018] Figure 6 is a block diagram of a processor-based system coupled to several memory modules according to another embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

[019] A processor-based electronic system 100 according to one example of the invention is shown in Figure 4. The system 100 includes three processors 104, 106, 108 and a direct memory access ("DMA") device 110, such as a graphics controller. The DMA device 110 and each of the processors 104-108 includes four memory access ports 112, 114, 116, 118. The ports 112-118 preferably include a data port as well as either individual or a shared control and address ports. However, it will be understood that some other memory port configuration may be used, such as a port for receiving and transmitting packets. The system 100 also includes a first rank 130 of four memory modules 132, 134, 136, 138 each of which includes a first set of four memory access ports 142, 144, 146, 148. As explained below, each of the memory modules 132-138 includes a memory hub coupled to eight memory devices, which are preferably dynamic random access memory ("DRAM") devices, and, more preferably, synchronous DRAM ("SDRAM") devices. However, it will be understood that a greater or lesser number of memory devices may be coupled to the memory hub in each of the memory modules 132-138.

[020] The memory access ports 112, 114, 116, 118 of the first processor 104 are coupled to the memory access port 142 of each of the memory modules 132, 134, 136, 138, respectively, through respective buses 162, 164, 166, 168. Similarly, the memory access ports 112, 114, 116, 118 of the second processor 106 are coupled to the memory access port 144 of each of the memory modules 132, 134, 136, 138, respectively, through respective buses 172, 174, 176, 178, and the memory access ports 112, 114,

116, 118 of the third processor 108 are coupled to the memory access port 146 of each of the memory modules 132, 134, 136, 138, respectively, through respective buses 182, 184, 186, 188. As a result, any of the processors 102-106 can access any of the memory modules 132-138. In a like manner, the memory access ports 112, 114, 116, 118 of the DMA device 110 are coupled to the memory access port 148 of each of the memory modules 132, 134, 136, 138, respectively, through respective buses 192, 194, 196, 198. Thus, the DMA device 108 can also access each of the memory modules 132, 134, 136, 138.

[021] Each of the memory modules 132, 134, 136, 138 also includes a second set of four memory access ports 202, 204, 206, 208 that are coupled to a second rank 210 of four memory modules 212, 214, 216, 218. More specifically, the memory access ports 202, 204, 206, 208 of the memory module 132 are coupled to a respective memory access port 222 of the memory modules 212, 214, 216, 218, respectively, through respective buses 232, 234, 236, 238. Similarly, the memory access ports 202, 204, 206, 208 of the memory module 134 are coupled to the memory access port 224 of each of the memory modules 212, 214, 216, 218, respectively, through respective buses 242, 244, 246, 248, and the memory access ports 202, 204, 206, 208 of the memory module 136 are coupled to the memory access port 226 of each of the memory modules 212, 214, 216, 218, respectively, through respective buses 252, 254, 256, 258. Finally, the memory access ports 202, 204, 206, 208 of the memory module 138 are coupled to the memory access port 228 of each of the memory modules 212, 214, 216, 218, respectively, through respective buses 262, 264, 266, 268.

[022] Each of the memory modules 212-218 in the second rank 210, like the memory modules 132-138 in the first rank 130, includes a memory hub coupled to eight memory devices. As explained in greater detail below, each of the memory hubs in the first rank 130 of memory modules 132-138 includes a crossbar switch (not shown in Figure 4) that can couple any of the memory access ports 112-118 to any of the memory access ports 202-208. In the same manner, a memory hub in each of the memory modules 212-218 in the second rank 210 can couple any of the memory access ports 202-208 to any of the memory access ports 222-228. As a result, any of the processors 102-106 and the DMA device 108 can access any of the memory modules 132-138 directly and any of the memory modules 212-218 through the memory modules 132-

138. This close proximity between the processors 102-106 and the memory modules 132-138, 212-218 and between the DMA device 108 and the memory modules 132-138, 212-218 results in a relatively low latency compared to latencies that are achievable with memory architectures of the type illustrated in Figures 1 and 3.

[023] An additional advantage of the memory topography shown in Figure 4 is that the data bandwidth between any of the processors 102-106 or the DMA device 108 and any of the memory modules 212-218 in the second rank 210 can be varied by simply varying the number of interconnections to the memory modules 212-218. For example, if each of the buses are 16-bits wide, the processor 106 can be coupled to any of the memory modules 212-218 through a 16-bit data bus by using only a single one of the buses extending from the processor 106 to one of the modules 132-138 and a single one of the buses extending from one of the modules 132-138 to one of the modules 212-218. The processor 106 can be coupled to any of the memory modules 212-218 through a 32-bit data bus by being coupled to two of the modules 132-138, and from each of these two modules 132-138 to one of the modules 212-218 though a respective bus. The processor 106 can be coupled to any of the memory modules 212-218 through a 48-bit data bus by being coupled to three of the modules 132-138, and from each of these three modules 132-138 to one of the modules 212-218 though a respective bus. Finally, the processor 106 can be coupled to any of the memory modules 212-218 through a 64-bit data bus by being coupled to all four of the modules 132-138, and from each of these four modules 132-138 to one of the modules 212-218 though a respective bus.

[024] One embodiment of a memory hub 300 that may be used in the memory modules 132-138, 212-218 of Figure 4 is shown in Figure 5. The memory hub 300 includes four link interfaces 304a-d that are coupled to respective buses, such as the buses in the first set shown in Figure 4. Similarly, four additional link interfaces 308a-d are included that are also coupled to respective buses, such as the buses in the second set shown in Figure 4. All of the link interfaces 304, 308 are coupled to a cross bar switch 310, which may be of a conventional or hereinafter developed design. The cross bar switch 310 can couple any of the link interfaces 304a-d to any of the link interfaces 308-a-d, as previously explained with respect to Figure 4. The link interfaces 304a-d, 308a-d may be either unidirectional or duplex interfaces, and the nature of the memory accesses coupled to or from the link interfaces 304a-d, 308a-d may vary as desired, including

conventional DRAM address, control and data signals, shared address and control signals and packetized memory access signals.

[025] The cross bar switch 310 can also couple any of the link interfaces 304a-d, 308a-d to four DRAM controllers 314a-d, each of which is coupled to a plurality of DRAM devices (not shown in Figure 5). The DRAM controllers 314a-d may be conventional DRAM controllers or some hereinafter developed design for a DRAM controller. The specific structure and operation of the DRAM controllers 314a-d will, of course, depend on the nature of the DRAM devices used in a memory module. The cross bar switch 310 couples the link interfaces 304a-d to the DRAM controllers 314a-d to allow any of a plurality of memory access devices to write data to or read data from DRAM devices coupled to the controllers 314a-d, as also explained above with reference to Figure 5. The cross bar switch 310 couples the link interfaces 308a-d to the DRAM controllers 314a-d to allow any data to be transferred to or from DRAM devices coupled to the DRAM controllers 314a-d from or to, respectively, other memory modules containing a memory hub 300.

[026] The memory hub 300 also includes a cache memory 320a-d and a write buffer 324a-d for each of the DRAM devices serviced by a respective DRAM controller 314a-d. As is well known in the art, each of the cache memories 320a-d, which may be a static random access memory ("SRAM") device, stores recently or frequently accessed data stored in the DRAM devices serviced by the respective DRAM controller 314a-d. The write buffers 324a-d accumulate write addresses and data directed to DRAM devices serviced by a respective one of the DRAM controllers 314a-d if the DRAM devices are busy servicing a read memory request or there are other read requests pending. By accumulating the write memory requests in this manner, they can be processed more efficiently in a pipelined manner since there is no need to incurs delays associated with alternating write and read requests.

[027] As mentioned above, data can be transferred from one memory module containing a memory hub 300 to another memory module containing a memory hub 300. These inter-module data transfers are controlled by a direct memory access ("DMA") engine 330, which may be of a conventional or hereinafter developed design. The DMA engine 330 may also be used to transfer data from a partially defective

memory module to a properly functioning memory module prior to disabling the operation of the partially defective memory module.

[028] The memory hub 300 will generally include components in addition to those shown in Figure 5. However, these have been omitted in the interest of brevity or clarity. Also, in some applications, components shown in Figure 5 may be omitted. For example, the write buffers 324a-d may be omitted if write/read access turnovers are acceptable. Also, although the memory hub 300 shown in Figure 5 includes two sets of four link interfaces 304a-d, 308a-d and four DRAM controllers 314a-d, the number of sets of link interfaces, the number of link interfaces in each set and the number of DRAM controllers may vary as desired.

[029] An alternative embodiment of a processor-based electronic system 350 is shown in Figure 6. The system 350 includes four memory requestors 352a-d, such as processors or direct memory access devices, each of which is coupled to a first rank 354 of four memory modules 356a-d through buses, generally indicated at 358. The memory modules 356a-d are thus configured in the same manner as in the embodiment of Figure 4. However, rather than coupling each of the memory modules 356 in the first rank 354 to each of the memory modules in a second rank of memory modules as in the embodiment of Figure 4, the memory modules 356a-d are each coupled to a respective set of four memory modules in a second rank 360. Thus, the first memory module 356a is coupled to four memory modules 362a-d, the second memory module 356b is coupled to four memory modules 362e-h, the third memory module 356c is coupled to four memory modules 362i-l, and the fourth memory module 356d is coupled to four memory modules 362m-p. The advantage of the topography shown in Figure 6 over the topography shown in Figure 4 is that the second rank 360 of memory modules 362 provide a larger memory capacity than the memory modules in the second rank 210 using the topography shown in Figure 4. However, a disadvantage of the memory topography shown in Figure 6 is that it provides less bandwidth and flexibility in accessing the memory modules 362 in the second rank 360.

[030] From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. For example, although the processors 104-108 and the DAM device 110 are

shown as being coupled directly to the memory modules 132-138, it will be understood that they may be coupled through other devices, such as bus bridges. Also, the systems 100, 350 shown in Figures 4 and 6, respectively, would normally include components in addition to those shown. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

What is claimed is:

1. A memory hub for use in coupling each of a plurality of memory requestors to each of a plurality of memory devices, the memory hub, comprising:

a plurality of memory controllers, each of the memory controllers being coupled to at least one of the memory devices;

a first plurality of link interfaces; and

a cross bar switch having a first plurality of switch ports and a plurality of memory ports, each of the switch ports being coupled to a respective one of the link interfaces, and each of the memory ports being coupled to a respective one of the memory controllers, the cross bar switch being operable to selectively couple each of the link interfaces to any one of the memory controllers.

2. The memory hub of claim 1, further comprising a second plurality of link interfaces, and wherein the cross bar switch further includes a second plurality of switch ports each of which is coupled to a respective one of the link interfaces in the second plurality of link interfaces.

3. The memory hub of claim 1 wherein each of the memory controllers comprises a dynamic random access memory controller.

4. The memory hub of claim 1, further comprising a respective cache memory device associated with each of the memory controllers, each of the cache memory devices being operable to store data that is accessed through a respective one of the memory ports of the cross bar switch.

5. The memory hub of claim 1 wherein each of the link interfaces comprises a bi-directional link interface.

6. The memory hub of claim 1 wherein each of the link interfaces comprises a pair of uni-directional links one of which includes an input port and the other of which includes an output port.

7. The memory hub of claim 1, further comprising a respective write buffer associated with each of the memory controllers, each of the write buffers being operable to receive write data and associated write addresses through a respective one of the memory ports of the cross bar switch and to store the write data and addresses for subsequent coupling to a respective one of the memory controllers.

8. The memory hub of claim 7 wherein each of the write buffers accumulate write data and addresses for a plurality of write memory accesses and then sequentially couple the plurality of write data and addresses to a respective one of the memory controllers without any intervening read memory accesses.

9. The memory hub of claim 1, further comprising a direct memory access device coupled to the cross bar switch and the memory controllers, the direct memory access device being operable to cause the cross bar switch and each of the memory controllers to perform memory write and read accesses.

10. A memory module, comprising:

a substrate;

a plurality of memory devices carried by the substrate; and

a memory hub carried by the substrate, the memory hub comprising:

a plurality of memory controllers, each of the memory controllers being coupled to at least one of the memory devices;

a first plurality of link interfaces, each of the link interfaces being for coupling to each of a plurality of memory requestors; and

a cross bar switch having a first plurality of switch ports and a plurality of memory ports, each of the switch ports being coupled to a

respective one of the link interfaces, and each of the memory ports being coupled to a respective one of the memory controllers, the cross bar switch being operable to selectively couple each of the link interfaces to any one of the memory controllers.

11. The memory module of claim 10 wherein the memory hub further comprising a second plurality of link interfaces, and wherein the cross bar switch further includes a second plurality of switch ports each of which is coupled to a respective one of the link interfaces in the second plurality of link interfaces.

12. The memory module of claim 10 wherein each of the memory devices comprises a dynamic random access memory device.

13. The memory module of claim 10 wherein the memory hub further comprises a respective cache memory device associated with each of the memory controllers, each of the cache memory devices being operable to store data that is accessed through a respective one of the memory ports of the cross bar switch.

14. The memory module of claim 10 wherein the substrate comprises a semiconductor substrate, and wherein the memory devices, the link interfaces, the memory controllers and the cross bar switch are fabricated as a common integrated circuit in the semiconductor substrate.

15. The memory module of claim 10 wherein each of the link interfaces in the memory hub comprises a bi-directional link interface.

16. The memory module of claim 10 wherein each of the link interfaces in the memory hub comprises a pair of uni-directional links one of which includes an input port and the other of which includes an output port.

17. The memory module of claim 10 wherein the memory hub further comprises a respective write buffer associated with each of the memory controllers, each of the write buffers being operable to receive write data and associated write addresses through a respective one of the memory ports of the cross bar switch and to store the write data and addresses for subsequent coupling to a respective one of the memory controllers.

18. The memory module of claim 17 wherein each of the write buffers accumulate write data and addresses for a plurality of write memory accesses and then sequentially couple the plurality of write data and addresses to a respective one of the memory controllers without any intervening read memory accesses.

19. The memory module of claim 10 wherein the memory hub further comprises a direct memory access device coupled to the cross bar switch and the memory controllers, the direct memory access device being operable to cause the cross bar switch and each of the memory controllers to perform memory write and read accesses.

20. A memory system, comprising:

a plurality of memory requestors;

a first rank of memory modules coupled to the memory requestors, each of the memory modules in the first rank being coupled to a plurality of the memory requestors, each of the memory modules comprising:

a plurality of memory devices; and

a memory hub comprising:

a plurality of memory controllers, each of the memory controllers being coupled to at least one of the memory devices;

a first plurality of link interfaces, each of the link interfaces being coupled to a respective one of the memory requestors; and

a cross bar switch having a first plurality of switch ports and a plurality of memory ports, each of the switch ports being coupled to a respective one of the link interfaces, and each of the memory ports being coupled to a respective one of the memory controllers, the cross bar switch being operable to selectively couple each of the link interfaces to any one of the memory controllers.

21. The memory system of claim 20 wherein the memory hub further comprising a second plurality of link interfaces, and wherein the cross bar switch further includes a second plurality of switch ports each of which is coupled to a respective one of the link interfaces in the second plurality of link interfaces.

22. The memory system of claim 20 wherein each of the memory devices comprises a dynamic random access memory device.

23. The memory system of claim 20 wherein the memory hub further comprises a respective cache memory device associated with each of the memory controllers, each of the cache memory devices being operable to store data that is accessed through a respective one of the memory ports of the cross bar switch.

24. The memory system of claim 20 wherein the substrate comprises a semiconductor substrate, and wherein the memory devices, the link interfaces, the memory controllers and the cross bar switch are fabricated as a common integrated circuit in the semiconductor substrate.

25. The memory system of claim 20 wherein each of the link interfaces in the memory hub comprises a bi-directional link interface.

26. The memory system of claim 20 wherein each of the link interfaces in the memory hub comprises a pair of uni-directional links one of which includes an input port and the other of which includes an output port.

27. The memory system of claim 20 wherein the memory hub further comprises a respective write buffer associated with each of the memory controllers, each of the write buffers being operable to receive write data and associated write addresses through a respective one of the memory ports of the cross bar switch and to store the write data and addresses for subsequent coupling to a respective one of the memory controllers.

28. The memory system of claim 27 wherein each of the write buffers accumulate write data and addresses for a plurality of write memory accesses and then sequentially couple the plurality of write data and addresses to a respective one of the memory controllers without any intervening read memory accesses.

29. The memory system of claim 20 wherein the memory hub further comprises a direct memory access device coupled to the cross bar switch and the memory controllers, the direct memory access device being operable to cause the cross bar switch and each of the memory controllers to perform memory write and read accesses.

30. The memory system of claim 20 wherein at least some of the memory requesters comprise a processor.

31. The memory system of claim 20 wherein at least some of the memory requesters comprise a direct memory access device.

32. The memory system of claim 20, further comprising a second rank of memory modules coupled to a plurality of the memory modules in the first rank, each of the memory modules in the second rank comprising:

a plurality of memory devices; and

a memory hub comprising:

a plurality of memory controllers, each of the memory controllers being coupled to at least one of the memory devices;

a first plurality of link interfaces, each of the link interfaces being coupled to a respective one of the memory modules in the first rank; and

a cross bar switch having a first plurality of switch ports and a plurality of memory ports, each of the switch ports being coupled to a respective one of the link interfaces, and each of the memory ports being coupled to a respective one of the memory controllers, the cross bar switch being operable to selectively couple each of the link interfaces to any one of the memory controllers.

33. The memory system of claim 20, further comprising a plurality of sets of memory modules in a second rank, the number of sets of memory modules in the second rank corresponding in number to the number of memory modules in the first rank, the memory modules in each set in the second rank being coupled to a respective one of the memory modules in the first rank, each of the memory modules in the second rank comprising:

a plurality of memory devices; and

a memory hub comprising:

a plurality of memory controllers, each of the memory controllers being coupled to at least one of the memory devices;

a link interface coupled to the memory module in the first rank to which the respective set of memory modules in the second rank corresponds; and

a cross bar switch having a switch port and a plurality of memory ports, the switch port being coupled to the link interface, and each of the memory ports being coupled to a respective one of the memory controllers, the cross bar switch being operable to selectively couple the link interface to any one of the memory controllers.

34. The memory system of claim 20 wherein the memory devices, the link interfaces, the memory controllers and the cross bar switch are fabricated as a common integrated circuit in a semiconductor substrate.

35. A memory system, comprising:

a plurality of memory requestors;

a first rank of memory modules coupled to the memory requestors, the memory modules in the first rank each including a first set of memory ports corresponding in number to the number of memory requestors, each of the memory ports in the first rank being coupled to a respective one of the memory requestors, the memory modules in the first rank further including a second set of memory ports, each of the memory modules in the first rank including a plurality of memory devices and a memory hub coupled to the memory devices and to the memory ports in the first and second sets;

a second rank of memory modules each including at least one memory port coupled to a memory module in the first rank through a memory port in the second set, each of the memory modules in the second rank including a plurality of memory devices and a memory hub coupled to the memory devices and to the at least one memory port.

36. The memory system of claim 35 wherein each of the memory modules in the second rank includes a plurality of memory ports corresponding in number to the number of memory modules in the first rank, the memory ports in

each of the memory modules in the second rank being coupled to the memory modules in the first rank through respective memory ports in the second set.

37. The memory system of claim 35 wherein each of a plurality of sets of memory modules in the second rank are coupled to a respective one of the memory modules in the first rank, each of the memory modules in the first rank including a memory port coupled to the memory module in the first rank to which the respective set of memory modules in the second rank corresponds.

38. The memory system of claim 35 wherein the memory hub in each of the memory modules in the first rank comprise a cross bar switch operable to couple any of the memory ports in the first set to any of the memory ports in the second set and to any of the memory devices.

39. The memory system of claim 38 wherein the cross bar switch is further operable to couple any of the memory ports in the second set to any of the memory devices.

40. A method of accessing memory devices from a plurality of memory request ports, the method comprising:

arranging a first plurality of memory modules in a first rank, each of the memory modules in the first rank including a plurality of memory devices;

arranging a second plurality of memory modules in a second rank, each of the memory modules in the second rank including a plurality of memory devices;

accessing each of the memory modules in the first rank from any of the memory request ports; and

accessing each of the memory modules in the second rank from any of the memory request ports through at least one of the memory modules in the first rank.

41. The method of claim 40 wherein the act of accessing each of the memory modules in the second rank through at least one of the memory modules in the first rank comprises accessing each of the memory modules in the second through each of the memory modules in the first rank.

42. The method of claim 40 wherein the act of accessing each of the memory modules in the second rank through at least one of the memory modules in the first rank comprises accessing a set of memory modules in the second rank through a memory module in the first rank corresponding to the set of memory modules in the second rank.

43. The method of claim 40, further comprising accessing each of the memory modules in the first rank from at least some of the memory modules in the second rank.

44. The method of claim 40 wherein the act of accessing each of the memory modules in the second rank from any of the memory request ports through at least one of the memory modules in the first rank comprises varying the number of memory modules in the first rank through which a memory module in the second rank is accessed from one of the memory request ports to vary the data bandwidth between the memory request port and the memory module in the second rank that is being accessed.

45. The method of claim 40, further comprising accumulating write memory accesses in at least some of the memory modules and sequentially executing the stored write memory accesses.