

MEMORY HUB WITH INTERNAL CACHE AND/OR MEMORY ACCESS PREDICTION

TECHNICAL FIELD

This invention relates to computer systems, and, more particularly, to a
5 computer system having a memory hub coupling several memory devices to a processor
or other memory access device.

BACKGROUND OF THE INVENTION

Computer systems use memory devices, such as dynamic random access
memory ("SDRAM") devices, to store instructions and data that are accessed by a
10 processor. In a typical computer system, the processor communicates with the system
memory through a processor bus and a memory controller. The processor issues a
command, such as a read command, and an address designating the location from which
data or instructions are to be read. The memory controller uses the command and
address to generate appropriate command signals as well as row and column addresses,
15 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.

Although the operating speed of memory devices has continuously
20 increased, this increase in operating speed has not kept pace with increases in the
operating speed of processors. Even slower has been the increase in operating speed of
memory controllers coupling processors to memory devices. The relatively low speed
of memory controllers and memory devices limits the communication bandwidth
between the processor and the memory devices.

25 In addition to the limited bandwidth between processors and memory
devices, the performance of computer systems and is also limited by latency problems
that increase the time required to read data from system memory devices. More
specifically, when a memory device read command is coupled to a system memory

device, such as an asynchronous DRAM ("SDRAM ") device, the read data is 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.

One approach to alleviating at the memory latency problem is to use multiple memory devices coupled to the processor through a memory hub. Computer systems employing this architecture can have a higher bandwidth because a processor can access one memory device while another memory device is responding to a prior memory access. For example, the processor can output write data to one of the memory devices in the system while another memory device in the system is preparing to provide read data to the processor. However, although computer systems using memory hubs may provide superior performance, they nevertheless often fail to operate at optimum speed. One of the reasons such computer systems fail to operate at optimum speed is that conventional memory hubs are essentially single channel systems since all control, address and data signals must pass through common memory hub circuitry. As a result, when the memory hub circuitry is busy communicating with one memory device, it is not free to communicate with another memory device. Furthermore, although computer systems using memory hubs can provide a greater memory bandwidth, they still suffer from latency problems of the type described above. More specifically, although the processor may communicate with one memory device while another memory device is preparing to transfer data, is sometimes necessary to receive data from one memory device before the data from another memory device can be used. In the event data must be received from one memory device before data received from another memory device can be used, the latency problem continues to slow the operating speed of such computer systems.

There is therefore a need for a computer architecture that provides the advantages of a memory hub architecture and also minimize this latency problems common in such systems, thereby providing a memory devices with high bandwidth and low latency.

SUMMARY OF THE INVENTION

A memory hub that may be used in a computer system includes a memory access device interface coupled to a processor or other memory access device, and a plurality of memory interfaces each of which is coupled to a respective memory device. Each of the memory interfaces includes a memory controller and, according to one aspect of the invention, a memory cache. Each of the memory interfaces is coupled to the memory access device interface by a switch. In operation, data read from or written to a memory device coupled to one of the memory interfaces are stored in the cache memory for the memory interface. In response to a subsequent memory read request, the cache memory is checked to determine whether the data corresponding to the memory read request are stored in the cache memory. In the event of a cache hit, the requested data are provided from the cache memory. Otherwise, the requested data are provided by the memory device. According to another aspect of the invention, each memory interface includes a memory controller and a prediction unit. The prediction unit predicts an address from which data are likely to be read based on an address from a prior memory access. The prediction unit then causes the memory controller in the respective memory interface to read data from the predicted address. The memory hub may be physically included in a system controller, a memory module, or some other component of a computer system or other electronic system using memory devices.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of a computer system according to one embodiment of the invention in which the memory hub is included in a system controller.

Figure 2 is a block diagram of a computer system according to another embodiment of the invention in which the memory hub is included in a memory module.

Figure 3 is a block diagram of a memory hub used in the computer systems of Figures 1 and 2.

DETAILED DESCRIPTION OF THE INVENTION

A computer system 100 according to one embodiment of the invention is shown in Figure 1. The computer system 100 includes a processor 104 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 104 includes a processor bus 108 that normally includes an address bus, a control bus, and a data bus. In addition, the computer system 100 includes one or more input devices 108, such as a keyboard or a mouse, coupled to the processor 104 through a system controller 110 to allow an operator to interface with the computer system 100. Typically, the computer system 100 also includes one or more output devices 114 coupled to the processor 104 through the system controller 110, such output devices typically being a printer or a video terminal. One or more data storage devices 120 are also typically coupled to the processor 104 through the system controller 110 to allow the processor 104 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices 120 include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs). The processor 104 is also typically coupled to cache memory 124, which is usually static random access memory ("SRAM").

The system controller 110 also includes a memory hub 130 for controlling access to several system memory devices 140a-d, each of which may be a synchronous dynamic random access memory ("SDRAM"). The memory hub 130 allows the processor 104 to write data to and read data from each of the system memory devices 140a-d. The memory hub 130 is coupled to each of the system memory devices 140a-d through a bus system 142, which normally includes a control bus, an address bus and a data bus.

Although the memory hub 130 is shown in Figure 1 coupled to the processor 104, it will be understood that the memory hub 130 may also be coupled to other components in a computer system chipset (not shown) and may also allow other devices (not shown) to write data to and read data from the system memory devices 140a-d in a direct memory operation, as is well known in the art. Also, the memory hub 130 may be physically included as a part of components of an electronic system other

than the system controller 110. For example, a computer system 144 shown in Figure 2 uses most of the same components that are used in the computer system 100 of Figure 1. In the interest of brevity, such common components have been provided with the same reference numerals, and an explanation of their operation will not be repeated. The
5 computer system 144 differs from the computer system 100 shown in Figure 1 in that the memory hub 130 is not included in the system controller 110. Instead, the system controller 110 is coupled to a plurality of memory modules 146, such a double in-line memory modules ("DIMMs"). Each of the memory modules 146 includes the memory
10 hub 130 and a plurality of memory devices 148, which may be SDRAM or some other type of memory device. The memory hub 130 operates in essentially the same manner explained above with reference to Figure 1 to cache data stored in the memory modules 146.

Although Figures 1 and 2 show the memory hub 130 included in the system controller 110 and the memory modules 146, respectively, it will be understood
15 that the memory hub 130 may be a stand-alone unit or it may be included in other components of a computer system or other system using memory devices.

One embodiment of the memory hub 130 is shown in Figure 3 in which the memory hub 130 is coupled to the processor 104 and three memory devices 140a-c, which, in the example illustrated in Figure 3, are SDRAM devices. The memory hub
20 130 is shown coupled to the processor 104 in a point-to-point arrangement in which there are no other devices coupled to the connection between the processor 104 and the memory hub 130. This type of interconnection provides better signal coupling between the processor 104 and the memory hub 130 for several reasons, including relatively low capacitance, relatively few line discontinuities to reflect signals and relatively short
25 signal paths. However, a multi-drop interconnection may alternatively be used in which other devices (not shown) are coupled to the interconnection between the processor 104 and the memory hub 130.

The memory hub 130 includes a processor interface 150 that is coupled to the processor 104 through a plurality of bus and signal lines, as is well known in the
30 art. The processor interface 150 is, in turn, coupled to a switch 160 through a plurality

of bus and signal lines, including a write data bus 154 and a read data bus 156, although a single bi-directional data bus may alternatively be provided to couple data in both directions between the processor interface 150 in the switch 160. The processor interface 150 is also coupled to switch 160 through a request line 164 and a snoop line 5 168. A snoop signal coupled from the switch 160 to the processor interface 150 through the snoop line 168 is used to maintain cache consistency, as will be described in greater detail below. A request signal coupled from the processor interface 150 to the switch 160 through the request line 164 provides the switch 160 with information corresponding to a request to transfer data through the switch 160. It will be 10 understood, however, that the processor interface 150 maybe coupled to the switch 160 with a greater or lesser number of buses and signal lines or buses and signal lines different from those illustrated in Figure 3.

The switch 160 is also coupled to three memory interfaces 170a-c which are, in turn, coupled to the system memory devices 140a-c, respectively. By providing a 15 separate and independent memory interface 170a-c for each system memory device 140a-c, respectively, the memory hub 130 avoids bus or memory bank conflicts that typically occur with single channel memory architectures. The switch 160 is coupled to each memory interface through a plurality of bus and signal lines, including a write data bus 174, read data bus 176 and a request line 178. However, it will be understood that a 20 single bi-directional data bus may alternatively be used instead of a separate write data bus 174 and read data bus 176. Significantly, each memory interface 170a-c is specially adapted to the system memory devices 140a-c to which it is coupled. More specifically, each memory interface 170a-c is specially adapted to provide and receive the specific signals received and generated, respectively, by the system memory device 140a-c to 25 which it is coupled. Also, the memory interfaces 170a-c are capable of operating with system memory devices 140a-c operating at different clock frequencies. As a result, the memory interfaces 170a-c isolate the processor 104 from changes that may occur at the interface between the memory hub 130 and memory devices 140a-c coupled to the hub 130, and it provides a more controlled environment to which the memory devices 30 140a-c may interface.

The switch 160 coupling the processor interface 150 to the memory interfaces 170a-c can be any of a variety of conventional or hereinafter developed switches. For example, the switch 160 may be a cross-bar switch that can simultaneously couple at the processor interface 150 and the memory interfaces 170a-c to each other. The switch 160 can also be a set of multiplexers that do not provide the same level of connectivity as a cross-bar switch but nevertheless can couple the processor interface 150 to each of the memory interfaces 170a-c. The switch 160 may also include arbitration logic (not shown) to determine which memory accesses should receive priority over other memory accesses. Bus arbitration performing this function is well known to one skilled in the art.

With further reference to Figure 3, each of the memory interfaces 170a-c includes a respective memory controller 180 and a respective cache memory unit 184. The memory controller 180 performs the same functions as a conventional memory controller by providing control, address and data signals to the system memory device 140a-c to which it is coupled and receiving data signals from the system memory device 140a-c to which it is coupled. The cache memory unit 184 includes the normal components of a cache memory including a tag memory, a data memory and a comparator, as is well known in the art. The memory devices used in the cache memory unit 184 may be either DRAM devices, static random access memory ("SRAM") devices, other types of memory devices, or a combination of all three. Furthermore, any or all of these memory devices as well as the other components used in the cache memory unit 184 may be either embedded or stand-alone devices.

The use of the cache memory unit 184 in each memory interface 170a-c allows the processor 104 to receive data responsive to a read command directed to a respective system memory device 140a-c without waiting for the memory device 140a-c to provide such data in the event that the data was recently read from or written to that memory device 140a-c. The cache memory unit 184 thus reduces the read latency of the system memory devices 140a-c to maximize the memory bandwidth of the computer system. Similarly, the processor 104 can store write data in the cache memory unit 184 and then perform other functions while the memory controller 180 in the same memory

interface 170a-c transfers the write data from the cache memory unit 184 to the system memory device 140a-c to which it is coupled.

To further reduce the memory access latency provided by the memory hub 130, each memory interface 170a-c may be provided with a prefetch unit 190. The
5 prefetch unit 190 is able to predict the likely address of a subsequent memory read request using conventional algorithms. The memory controller 180 in the same memory interface 170a-c can then perform the memory access in the background while the processor 104 is either accessing a different system memory device 140 or performing other functions. When the processor 104 subsequently provides a command to the
10 memory hub 130 to read data from the predicted address, the read data will already be present in the cache memory unit 180 and can thus be quickly provided to the processor 104.

From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration,
15 various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

1. A memory hub, comprising:
a memory access device interface structured to interface with a memory access device;
a plurality of memory interfaces structured to interface with respective memory devices, each of the memory interfaces including a memory controller and a memory cache; and
a switch coupling the memory access device interface to each of the memory interfaces.
2. The memory hub of claim 1 wherein the memory access device interface comprises a processor interface structured to interface with a processor.
3. The memory hub of claim 1 wherein each of the memory interfaces further comprises a prediction unit structured to predict an address from which data are likely to be read based on an address from a prior memory access and to cause the memory controller in the respective memory interface to output signals indicative of a memory read operation from the predicted address.
4. The memory hub of claim 3 wherein the prediction unit is further structured to cause the memory interface to store in the cache memory read data received responsive to the signals indicative of a memory read operation.
5. The memory hub of claim 1 wherein each of the memory interfaces operates at the same clock speed.
6. The memory hub of claim 1 wherein the switch comprises a cross-bar switch.

7. The memory hub of claim 1 wherein the switch comprises a multiplexer switch.

8. The memory hub of claim 1 wherein the cache memory comprises dynamic random access memory.

9. A memory hub, comprising:
a memory access device interface structured to interface with a memory access device;

a plurality of memory interfaces structured to interface with respective memory devices, each of the memory interfaces including a memory controller and a prediction unit structured to predict an address from which data are likely to be read based on an address from a prior memory access and to cause the memory controller in the respective memory interface to output signals indicative of a memory read operation from the predicted address; and

a switch coupling the memory access device interface to the memory interfaces.

10. The memory hub of claim 9 wherein the memory access device interface comprises a processor interface.

11. The memory hub of claim 9 wherein each of the memory interfaces operates at the same clock speed.

12. The memory hub of claim 9 wherein the switch comprises a cross-bar switch.

13. The memory hub of claim 9 wherein the switch comprises a multiplexer switch.

14. A computer system, comprising:
a processing unit operable to perform computing functions;
a system controller coupled to the processing unit;
at least one input device coupled to the processing unit through the system controller;
at least one output device coupled to the processing unit through the system controller;
at least one data storage devices coupled to the processing unit through the system controller;
a plurality of memory devices; and
a memory hub comprising:
a processor interface coupled to the processor;
a plurality of memory interfaces coupled to respective ones of the memory devices, each of the memory interfaces including a memory controller and a memory cache; and
a switch coupling the processor interface to each of the memory interfaces.

15. The computer system of claim 14 wherein the memory hub is physically included in the system controller.

16. The computer system of claim 14 wherein the plurality of memory devices are physically packaged in a memory module, and wherein the memory hub is physically included in the memory module.

17. The computer system of claim 14 wherein each of the memory interfaces further comprises a prediction unit structured to predict an address from which data are likely to be read based on an address from a prior memory access and to cause the memory controller in the respective memory interface to apply to the memory device to which

the memory interface is coupled output signals indicative of a memory read operation from the predicted address.

18. The computer system of claim 15 wherein the prediction unit is further structured to cause the memory interface to store in the cache memory read data received from the respective memory device responsive to the signals indicative of a memory read operation.

19. The computer system of claim 14 wherein each of the memory interfaces operates at the same clock speed.

20. The computer system of claim 14 wherein the switch comprises a cross-bar switch.

21. The computer system of claim 14 wherein the switch comprises a multiplexer switch.

22. The computer system of claim 14 wherein the cache memory comprises dynamic random access memory.

23. The computer system of claim 14 wherein each of the memory devices comprises a dynamic random access memory device.

24. The computer system of claim 21 wherein each of the dynamic random access memory device comprises a synchronous dynamic random access memory device.

25. A computer system, comprising:
a processing unit operable to perform computing functions;
a system controller coupled to the processing unit;

at least one input device coupled to the processing unit through the system controller;

at least one output device coupled to the processing unit through the system controller;

at least one data storage devices coupled to the processing unit through the system controller;

a plurality of memory devices; and

a memory hub comprising:

a processor interface coupled to the processor;

a plurality of memory interfaces coupled to respective ones of the memory devices, each of the memory interfaces including a memory controller and a prediction unit structured to predict an address from which data are likely to be read based on an address from a prior memory access and to cause the memory controller in the respective memory interface to output to the memory device to which the memory interface is coupled signals indicative of a memory read operation from the predicted address; and

a switch coupling the processor interface to each of the memory interfaces.

26. The computer system of claim 25 wherein the memory hub is physically included in the system controller.

27. The computer system of claim 25 wherein the plurality of memory devices are physically packaged in a memory module, and wherein the memory hub is physically included in the memory module.

28. The computer system of claim 25 wherein each of the memory interfaces operates at the same clock speed.

29. The computer system of claim 25 wherein the switch comprises a cross-bar switch.

30. The computer system of claim 25 wherein the switch comprises a multiplexer switch.

31. The computer system of claim 25 wherein each of the memory devices comprises a dynamic random access memory device.

32. A method of accessing a plurality of memory devices, comprising:
directing a memory access request to a first of a plurality of memory devices coupled to memory hub;

storing data read from or written to the first memory device in a cache memory located in the memory hub;

subsequently directing a memory read request to the first memory device;

in response to the memory read request, detecting whether the data corresponding to the memory read request are stored in the cache memory located in the memory hub;

if the data corresponding to the memory read request are determined to be stored in the cache memory located in the memory hub, providing the read data from the cache memory; and

if the data corresponding to the memory read request are determined to be not stored in the cache memory located in the memory hub, providing the read data from the first memory device.

33. The method of claim 32 further comprising:

predicting an address from which data are likely to be read from the first memory device based on an address from a prior memory access to the first memory device;

providing read data from the predicted address in the first memory device; and

storing the read data from the predicted address in the cache memory in the memory hub.

34. The method of claim 32 wherein the act of storing data read from or written to the first memory device in a cache memory in the memory hub comprises storing the data read from or written to the first memory device in a cache memory dedicated to the first memory device.

35. The method of claim 32 wherein the memory access request on which the prediction was based comprises a read memory access.

36. The method of claim 32 wherein the memory access request on which the prediction was based comprises a write memory access.

37. A method of accessing a plurality of memory devices, comprising:
directing memory access requests to respective addresses in a plurality of memory devices coupled to memory hub;

within the memory hub, predicting at least one address from which data are likely to be read from the first memory device based on the addresses to which the memory access requests were directed; and

providing respective read data from the predicted addresses in the memory devices prior to receiving memory read requests directed to the predicted addresses.

38. The method of claim 37 wherein the memory access requests on which the predictions were based comprise read memory requests.

39. The method of claim 37 wherein the memory access requests on which the predictions were based comprise write memory requests.