SEQUENTIAL NIBBLE BURST ORDERING FOR DATA

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present invention is directed to memory devices and, more particularly, to methods and circuits for reading information into and out of the memory device.

Description of the Background

[0002] Computer designers are continually searching for faster memory devices that will permit the design of faster computers. A significant limitation on a computer's operating speed is the time required to transfer data between a processor and a memory circuit, such as a read or write data transfer. Memory devices such as dynamic random access memories (DRAMs), synchronous dynamic random access memories (SDRAMs), flash memories, etc. typically include a large number of memory cells arranged in one or more arrays, each array comprised of rows and columns. Each memory cell provides a location at which the processor can store and retrieve one bit of data, sometimes referred to as a memory bit or m-bit. The more quickly the processor can access the data within the memory cells, the more quickly it can perform a calculation or execute a program using the data.

[0003] FIG. 1 shows, in part, a typical computer system architecture. A central processing unit (CPU) or processor 10 is connected to a processor bus 12, which in turn is connected to a system or memory controller 14. The memory controller 14 may be connected to an expansion bus 16. The memory controller 14 serves as interface circuitry

between the processor 10 and a memory device 18. The processor 10 issues a command and an address which are received and translated by the memory controller 14. The memory controller 14 applies the translated command signals on a plurality of command lines 20 and the translated address on a plurality of address lines 22 to the memory device 18. These command signals are well known in the art and include, in the case of a DRAM, RAS (row address strobe), CAS (column address strobe), WE (write enable) and OE (output enable), among others. A clock signal is also provided on CLK lines 24. Corresponding to the processor-issued command and address, data is transferred between the controller 14 and the memory 18 via datapath lines 26.

[0004] Methods exist to enable memory devices, such as DRAM memory 18, to appear to external devices to be operating faster than the time it takes for the memory device to retrieve data from the array. These methods include pipeline and prefetch methods of operation. The pipeline method divides internal processing into a number of stages and sequentially processes information relating to one unit of data through each stage. Processing in each stage is carried out simultaneously in parallel, such that the rate at which data can be output from the device can be greater than the rate at which data is retrieved from the array. In the prefetch method, all internal processing is carried out in parallel, and parallel to serial conversion is performed at the input/output section.

[0005] Both the pipeline and prefetch methods can be used to support, for example, a burst mode of operation. The burst mode of operation is a mode of operation in which the starting address for a data string is provided to the memory device. The data string to be read out of the memory or written into the memory is then synchronously output or input, respectively, with a clock signal.

[0006] Historically, synchronous DRAMs have supported both an interleaved and a sequential burst mode of operation. Advance DRAM technology standards are being defined with an 8-bit external prefetch and capability to support a 4-bit or 8-bit internal prefetch. With a 4-bit internal prefetch, the sequential read or write crosses a boundary and is therefore difficult to implement as illustrated by the following table, Table 1.

Table 1

Starting Address	Internal Bits [0 1 2 3]	Internal Bits [4 5 6 7]
0	0 1 2 3	4 5 6 7
1	1 2 3 4	5 6 7 0
2	2 3 4 5	6 7 0 1
3	3 4 5 6	7 0 1 2
4	4 5 6 7	0 1 2 3
5	5 6 7 0	1 2 3 4
6	6 7 0 1	2 3 4 5
7	7 0 1 2	3 4 5 6

[0007] As seen from Table 1, except for starting addresses 0 and 4, the sequential burst cannot be executed without an 8-bit internal burst, adding cost, or a dual prefetching, which adds latency.

[0008] The existing interleave burst mode supports a 4-bit internal prefetch but some applications still use a sequential type of access burst mode. One solution is to always start the read burst at index 0 and sequence through the data. That solution is acceptable only when the word stored at index 0 is the next critical word. If the critical word is indexed at any other location, latency is introduced.

[0009] Thus, the need exists for a method and apparatus for enabling both 8-bit and 4-bit internal prefetches for new architectures without adding cost or latency to the new architecture.

SUMMARY OF THE INVENTION

[0010] The present invention is directed to a memory device comprising a plurality of arrays of memory cells and peripheral devices for reading information out of and for writing information into the memory cells. The peripheral devices include a decode circuit responsive to a first portion of address information for identifying an address and is further responsive to a second portion of the address information for identifying an order. The address may be a read address or a write address, and the order may be the order for reading data or writing data, respectively.

[0011] The present invention also includes a read sequencer circuit or both a write sequencer circuit and a read sequencer circuit for reordering bits to be written to or read from, as the case may be, the memory in response to another portion of the

address information. The necessary address information is routed to the sequencer circuits by an address sequencer.

[0012] The present invention is also directed to a method of reading a word from a memory array in at least two prefetch operations, wherein the order of the prefetch operations is controlled by an address bit, or writing a word in two n-bit bytes under the control of the address bit.

[0013] In one implementation of the present invention, the new burst sequence splits, for example, an 8-bit burst into two 4-bit bursts with a sequential interleave within each burst sequence. That enables each of the 4-bit bursts to be output from a memory array before the 8-bit burst is required to be output from the memory device. To implement that operation, the most significant column address bits (for example CA3-CAi) identify which 8-bit burst is selected. Those address bits may be referred to as a first portion of the address information. Address bit CA2, referred to as a second portion of the address information, identifies which of the two 4-bit bursts are fetched first from the memory array. CA0 and CA1 may then be used to identify which of the prefetched 4-bits are to be asserted first, with the remaining 3 bits output in sequential order from the first bit.

[0014] The present invention allows sequential type of interleaves for applications requiring them and provides access to the most critical word first without adding any latency to the system. Those, and other advantages and benefits, will become apparent from the detailed description of the preferred embodiments hereinbelow.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] For the present invention to be easily understood and readily practiced, the present invention will now be described, for purposes of illustration and not limitation, in conjunction with the following figures, wherein:

[0016] Fig. 1 is a functional block diagram of a computer system architecture;

[0017] Fig. 2 is simplified block diagram of an architecture for implementing the present invention;

[0018] Figs. 3A and 3B are timing diagrams comparing a 4-bit prefetch to an 8-bit prefetch, respectively; and

[0019] Fig. 4 is a simplified block diagram of a computer system in which the present invention may be used.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0020] Advanced DRAM technology (ADT) specifies an 8-bit external prefetch and supports either a 4 or 8-bit internal prefetch. Typical DRAMs support a sequential and an interleaved burst mode of operation. However, a sequential interleave is not compatible with a DRAM having a double pumped 4-bit internal prefetch DRAM architecture. The present invention allows for a new burst ordering sequence to support a multiple internal prefetch architecture for applications requiring a sequential-like burst sequence. The present invention allows sequential types of interleaves for applications requiring them, and provides access to the most critical word first.

[0021] Turning now to FIG. 2, FIG. 2 illustrates a simplified block diagram of an architecture for a DRAM capable of implementing the present invention. The DRAM memory device 29 is comprised of a command/address input buffer 30 responsive to a command bus or command lines and an address bus or address lines. A command decoder and sequencer 32 and an address sequencer 34 are each responsive to the command/address input buffer 30.

[0022] A bank address decoder 36 is responsive to the address sequencer 34 while bank control logic 38 is responsive to the bank address decoder 36. A series of row latch/decoder/drivers 40 is responsive to the bank control logic 38 and the address sequencer 34. One row latch/decoder/driver 40 is provided for each memory array 42. Illustrated in FIG. 2 are four memory arrays labeled bank 0 through bank 3. Accordingly, there are four row latch/decoder/driver circuits 40 each responsive to one of bank 0 through bank 3.

[0023] A column latch/decode circuit 44 is responsive to the address sequencer 34. The column latch/decode circuit 44 receives the most significant bits of the column address CA3-CAi, where "i" in the present example equals 9. The most significant bits CA3-CAi may be thought of as a first portion of the address and is used for identifying a word to be read. The column latch/decode circuit 44 also receives one of the least significant column address bits CA0-CA2; in the present example, the column latch/decode circuit 44 receives the column address bit CA2 which may be referred to as a second portion of the address. The identified word which is to be read may be, for example, an 8-bit word. The word will be read in two

4-bit bytes, and the second portion of the address identifies which of the first or the second n-bit bytes is to be read first.

[0024] An I/O gating circuit 46 is responsive to the column latch/decode circuit 44 for controlling sense amplifiers within each of the memory arrays 42.

[0025] The DRAM 29 may be accessed through a plurality of data pads 48 for either a write operation or a read operation. For a write operation, data on data pads 48 is received by receivers 50 and passed to input registers 52. A write sequencer circuit 54 orders the two 4-bit bytes comprising each 8-bit byte in response to, for example, column address bits col 0-1. The ordered bytes are then input to a write latch and driver circuit 56 for input to the memory arrays 42 through the I/O gating circuit 46. Data which is to be read from the memory arrays 42 is output through the I/O gating circuit 46 to a read latch 58. From the read latch 58, the information is input to a read sequencer circuit 60 which orders the read data in response to, for example, column address bits col 0-1. The ordered data is then output to an output mux 62 and then onto the data pads 48 through drivers 64.

[0026] The command/address input buffer 30, command decoder and sequencer 32, address sequencer 34, bank address decoder 36, bank control logic 38, the row latch/decoder/drivers 40, column latch decode circuit 44, I/O gating circuit 46 , the receivers 50, input registers 52, write sequence circuit 54, write latch and driver circuit 56, read latch 58, read sequence circuit 60 output mux 62 and drivers 64 are considered to be a plurality of peripheral devices for reading information out of and writing information into the memory cells of the arrays. The description of the forgoing elements as a plurality of peripheral devices is intended to provide a description of the presently preferred embodiment, and is not intended to limit the scope of the invention to only the recited devices. Those of ordinary skill in the art will recognize that other combinations of devices may be used to implement the plurality of peripheral devices, particularly where other memory architectures are used.

[0027] In general terms, the purpose of the read sequencer circuit 60 is to reorder the prefetched portions of the read word in response to certain the least significant address bits CA0 – CA2; in this example CA0 and CA1 are used.

[0028] The first n-bit prefetch (in this example, the first 4-bit prefetch identified by CA2) is reordered according to the start address identified by CA0 and CA1 as follows:

Starting Address	Internal Bits [0 1 2 3]	Internal Bits [4 5 6 7]
0	0 1 2 3	4 5 6 7
1	1 2 3 0	5 6 7 4
2	2 3 0 1	6 7 4 5
3	3 0 1 2	7 4 5 6
4	4 5 6 7	6 1 2 3
5	5 6 7 4	1 2 3 0
6	6 7 4 5	2 3 0 1
7	7 4 5 6	3 0 1 2

[0029] In operation, when a read command is received, the value on the bank address inputs BA0 and BA1 (not shown) selects one of the memory arrays 42. Address information is then received which identifies a row or rows within each array 42. The address provided on inputs CA3 through CAi (where "i" in the present example equals 9) selects the starting column location. Referring to FIG. 2, CA3 – CA9 are input to the column latch/decode circuit 44 to identify a word to be read. CA2 is also input to the column latch/decode circuit 44 for the purpose of identifying which portion of the word is to be read first. The bits CA0 and CA1 are input to the read sequencer circuit 60. That information identifies the start address such that the bits can be reordered thereby enabling the most critical word to be output first by the mux 62.

[0030] For a write operation, the bank is identified in the same manner as for a read operation. Similarly, the starting column address is identified in the same manner. The signals available at inputs CA0-CA2 are input to write sequencer 54 which reorders the bits as described. Although FIG. 2 shows both a write sequencer circuit 54 and a read sequencer circuit 60, the memory can operate with just the read sequencer circuit 60.

[0031] FIG. 3A illustrates a timing diagram for an 8-bit external prefetch using a 4-bit internal prefetch. As can be seen, after the read latency period, the data available at the output pads appears as an 8-bit byte, although the word was constructed from two 4-bit bytes. While the first 8-bit byte is made available at the data pads, a next 8-bit byte can be processed internally in two 4-bit prefetches as shown in the figure. In contrast, in FIG. 3B, the 8-bit byte is prefetched from the memory in one step.

[[0032] The timing diagram illustrated in FIG. 3A is the timing diagram for a 4-bit double pumped array. The array runs at a frequency of 1/4 that of the IO

frequency. Because not all 8 bits of data may be available for data scramble prior to the memory device outputting data to the external data pads, a data scramble must be performed on the 4 bit boundaries. That places a limit on the maximum data frequency that can be supported.

[0033] The timing diagram illustrated in FIG. 3B illustrates an 8-bit single pumped array. That array runs at a frequency of 1/8 of that of the data frequency. All 8 bits are available for data scramble prior to outputting data to the data pads, such that the output scramble may be completed on an 8 bit byte. The maximum data frequency is scaleable (the core is not a limiting factor) at the expense of die size.

[0034] Advantages of the present invention include the ability to support 4-bit internal prefetches at low cost with no addition to device latency, the critical word needed by the system is output first, and a sequential type burst for applications not supporting interleaved bursts is possible.

[0035] The present invention is also directed to a method of reading a word from a memory array in at least two prefetch operations, wherein the order of the prefetch operations is controlled by at least one address bit. The present invention is also directed to a method of outputting an n-bit word in two ½ n-bit prefetch steps from a plurality of memory arrays in response to an address bit. The present invention is also directed to a method comprised of prefetching the first portion of a word from a memory array and prefetching a second portion of the word from the memory array in an order determined by an address bit.

[0036] FIG. 4 is a block diagram of one example of a computer system 110 in which the present invention may be implemented. The computer system 110 includes a processor 112, a memory subsystem 114, and an expansion bus controller 116. The memory subsystem 114 and the expansion bus controller 116 are coupled to the processor 112 via a local bus 118. The expansion bus controller 116 is also coupled to at least one expansion bus 120, to which various peripheral devices 121 – 123 such as mass storage devices, keyboard, mouse, graphic adapters, and multimedia adapters may be attached. Processor 112 and memory subsystem 114 may be integrated on a single chip.

[0037] The memory subsystem 114 includes a memory controller 124 which is coupled to a plurality of memory modules 125, 126 via a plurality of signal lines 129, 130, 129a, 130a, 129b, 130b, 129c and 130c. The plurality of data signal lines 129, 129a, 129b, 129c are used by the memory controller 124 and the memory modules

125, 126 to exchange data DATA. Addresses ADDR are signaled over a plurality of address signal lines 132, clock signals CLK are applied on a clock line 133, and commands CMD are signaled over a plurality of command signal lines 134. The memory modules 125, 126 include a plurality of memory devices 136 – 139, 136' – 139' and a register 141, 141', respectively. Each memory device 136 – 139, 136' – 139' may be a high speed synchronous memory device. Although only two memory modules 125, 126 and associated signal lines 129 – 129c, 130 – 130c are shown in FIG. 5, it should be noted that any number of memory modules can be used.

[0038] The plurality of signal lines 129 – 129c, 130 –130c, 132, 133, 134 which couple the memory modules 125, 126 to the memory controller 124 are known as the memory bus 143. The memory bus 143 may have additional signal lines which are well known in the art, for example chip select lines, which are not illustrated for simplicity. Each column of memory devices 136 – 139, 136' - 139' spanning the memory bus 143 is known as a rank of memory. Generally, single side memory modules, such as the ones illustrated in FIG. 4, contain a single rank of memory. However, double sided memory modules containing two ranks of memory may also be used.

[0039] Read data is output serially synchronized to the clock signal CLK, which is driven across a plurality of clock signal lines, 130, 130a, 130b, 130c. Write data is input serially synchronized to the clock signal CLK, which is driven across the plurality of clock signal lines 130, 130a, 130b, 130c by the memory controller 124. Commands and addresses are also clocked using the clock signal CLK which is driven by the memory controller 124 across the registers 141, 141' of the memory modules 125, 126, respectively, to a terminator 148. The command, address, and clock signal lines 134, 132, 133, respectively, are directly coupled to the registers 141, 141' of the memory modules 125, 126, respectively. The registers 141, 141' buffer those signals before they are distributed to the memory devices 136 – 139, 136' – 139' of the memory modules 125, 126, respectively.

[0040] While the present invention has been described in conjunction with preferred embodiments thereof, those of ordinary skill in the art will recognize that many modifications and variations are possible. Such modifications and variations fall within the scope of the present invention which is limited only by the following claims.

What is claimed is:

1. A memory device, comprising:

a plurality of arrays of memory cells; and

peripheral devices for reading information out of and for writing information into said memory cells, said peripheral devices comprising:

a decode circuit responsive to a first portion of address information for identifying an address, said decode circuit further responsive to a second portion of the address information for identifying an order;

an address sequencer for routing at least one bit of the address information; and

a sequencer circuit response to said address sequencer.

2. The memory of claim 1 wherein said sequencer circuit comprises a read sequencer circuit.

3. The memory of claim 2 additionally comprising a write sequencer circuit responsive to said address sequencer.

4. The memory of claim 1 wherein said address sequencer routes at least two of the CA0, CA1 and CA2 column address bits.

5. The memory of claim 1 wherein said decode circuit is responsive to column address bits CA3 through CAi for identifying an address, and is responsive to at least one of column address bits CA0 to CA2 for identifying an order.

6. The memory of claim 1 wherein said memory device includes a DRAM.

7. A memory device, comprising:

a plurality of arrays of memory cells; and

a plurality of peripheral devices for reading information out of and for writing information into said memory cells, said peripheral devices comprising:

a decode circuit responsive to a first portion of address information for identifying a word to be read, said decode circuit further responsive to a second portion of the address information for identifying the order in which portions of said word are read;

an address sequencer for routing at least one bit of the address information; and

a read sequencer for ordering bits received from said plurality of arrays in response to said address sequencer.

8. The memory of claim 7 wherein said address sequencer routes at least two of the CA0, CA1 and CA2 column address bits.

9. The memory of claim 7 wherein said decode circuit is responsive to column address bits CA3 through CAi for identifying the word to be read, and is responsive to at least one of column address bits CA0 to CA2 for identifying the order in which portions of said word are prefetched.

10. The memory of claim 9 wherein said word is prefetched in n-bit bytes where n is equal to 4-bits.

11. The memory of claim 7 additionally comprising a write sequencer circuit responsive to said address sequencer.

12. The memory of claim 7 wherein said memory device includes a DRAM.

13. A system, comprising:

a processor;

a memory controller responsive to said processor;

a first bus interconnecting said processor and said memory controller;

a plurality of memory devices; and

a second bus for interconnecting said memory controller and said plurality of memory devices, each memory device comprising:

a plurality of arrays of memory cells; and

peripheral devices for reading information out of and for writing information into said memory cells, said peripheral devices comprising:

a decode circuit responsive to a first portion of address information for identifying an address, said decode circuit further responsive to a second portion of the address information for identifying an order;

an address sequencer for routing at least one bit of the address information; and

a sequencer circuit response to said address sequencer.

14. The system of claim 13 wherein said sequencer circuit comprises a read sequencer circuit.

15. The system of claim 14 additionally comprising a write sequencer circuit responsive to said address sequencer.

16. The system of claim 13 wherein said address sequencer routes at least two of the CA0, CA1 and CA2 column address bits.

17. The system of claim 13 wherein said decode circuit is responsive to column address bits CA3 through CAi for identifying an address, and is responsive to at least one of column address bits CA0 to CA2 for identifying an order.

18. The system of claim 13 wherein said plurality of memory devices includes a plurality of DRAMs.

19. A system, comprising:

a processor;

a memory controller responsive to said processor;

a first bus interconnecting said processor and said memory controller;

a plurality of memory devices; and

a second bus for interconnecting said memory controller and said plurality of memory devices, each memory device comprising:

a plurality of arrays of memory cells; and

a plurality of peripheral devices for reading information out of and for writing information into said memory cells, said peripheral devices comprising:

a decode circuit responsive to a first portion of address information for identifying a word to be read, said decode circuit further responsive to a second portion of the address information for identifying the order in which portions of said word are read;

an address sequencer for routing at least one bit of the address information; and

a read sequencer for ordering bits received from said plurality of arrays in response to said address sequencer.

20. The system of claim 19 wherein said address sequencer routes at least two of the CA0, CA1 and CA2 column address bits.

21. The system of claim 19 wherein said decode circuit is responsive to column address bits CA3 through CAi for identifying the word to be read, and is responsive to at least

one of column address bits CA0 to CA2 for identifying the order in which portions of the word are prefetched.

22. The system of claim 21 wherein said word is prefetched in n-bit bytes where n is equal to 4-bits.

23. The system of claim 19 additionally comprising a write sequencer circuit responsive to said address sequencer.

24. The system of claim 19 wherein said memory plurality of memory devices includes a plurality of DRAMs.

25. A method, comprising:

outputting an n-bit word in two 1/2n bit prefetch steps from a plurality of memory arrays in response to an address bit.

26. The method of claim 25 additionally comprising the step of reordering the bits of each 1/2n bit prefetch in response to other address bits.

27. The method of claim 26 wherein said address bit and said other address bits include any of the least significant column address bits.

28. The method of claim 25 additionally comprising the step of writing an n-bit word into one of the plurality of memory arrays in response to said address bit.

29. A method, comprising:

prefetching a first portion of a word from a memory array; and

prefetching a second portion of said word from said memory array, said first and second portions being determined by an address bit.

30. The method of claim 29 additionally comprising the step of reordering the bits of each of said first and second portions in response to other address bits.

31. The method of claim 30 wherein said address bit and said other address bits include any one of the column address bits CA0 to CA2.

32. The method of claim 29 additionally comprising the step of writing an n-bit word into the memory array in response to said address bit.

33. A method, comprising:

reading a word from a memory array in at least two prefetch operations, wherein the order of said prefetch operations is controlled by an address bit.