

9804329
10019

THIS SHEET HAS BEEN
REPLACED

- 1 -

BYTE WRITE CAPABILITY FOR MEMORY ARRAY

Background of the Invention

1. Field of the Invention

The present invention generally relates to computer systems and more particularly to a method and device for selectively writing one or more bytes of information to a memory word in a memory array.

2. Background

Conventional computers use several pieces of interconnected hardware, including input/output devices for the user interface (such as a keyboard and display), a permanent storage device (such as a magnetic or optical disk), a temporary memory device (such as random-access memory or RAM), and a central processing unit (CPU or processor) which accesses the permanent storage and temporary memory when carrying out program instructions. The present invention relates to a method and device for accessing temporary memory arrays such as RAM.

These arrays are often organized in rows and columns and sometimes further by other groupings such as sectors, wherein a sector may include many columns. All of the cells in a given row are referred to as a memory word. A given memory cell (bit) in the array is accessed for a store or load operation by selecting the particular cell based on its address. Such selection is commonly performed by a row-decoding circuit which presents an active logic level on a conductor, known as a word line, which traverses a given row of cells. When the word line for a given row is activated,

This Sheet Has Been Replaced



all of the memory cells in the row are connected to their respective "bit lines," which are, in turn, connected to other circuitry, such as a sense amplifiers, that allow the memory cells to be accessed. For example, a word line can access an array having 32 cells in a row to provide a 32-bit value (program instruction or data). In a memory array having 2^N word lines, an address input of N bits is generally required to select one of the word lines. Each word line may have a decoder circuit, and all of the decoder circuits receive and decipher the N-bit address input. In response thereto, only one word line is selected, and all other word lines in the array are deselected.

Figure 1 illustrates an example of a memory array wherein the memory cells 2 are placed in small groups 4 of 8 bits each, and several such groups 4 are present in a given column 6. Any encoding method is used to present a unique memory address to access a particular row of memory cells. A given word line 8 (only one is shown in Figure 1) is selected by any type of decoding circuitry which may include a gate 10. Only one word line will be selected and all others will be deselected. When a word line is selected, all cells in the corresponding row are connected to their respective local bit lines 12. In the array of Figure 1, all of the local bit lines for a given group 4 of memory cells are further connected to a global bit line 14. The global bit lines are further connected to the circuitry used for reading from and writing to the cells (e.g., sense amplifiers).

One particular design (described in U.S. Patent Application Serial No. 08/717,575 (Attorney Docket No. AA9-95-143)) uses a memory cell which has two inverters

connected in a flip-flop arrangement with a word line controlling the charge on the gate of a transistor whose source is connected to one side of the flip-flop, and whose drain is connected to the bitline. A "clear" line controls the charge on the gate of a second transistor whose source is connected to the other side of the flip-flop, and whose drain is connected to ground; the clear line is thus used to clear the memory cell. While this memory-cell design possesses several advantages over prior-art designs, one disadvantage is that the cell must be cleared prior to writing to the cell, and due to the layout of cells in the array, it is further necessary that all cells in a row share the same clear line. Thus, all cells in a row must be cleared prior to writing to any cells in that row. However, arrays are also often required to support writing to individual bytes in a row, i.e., the row is composed of a plurality of bytes each having one or more bits, but only some of these bytes are to be changed. An array usually has several thousand bytes, and it is not practical to provide several thousand corresponding rows so it is convenient to group bytes together in a single row. It is also convenient to have several bytes in a single row since the read operation reads every bit in the row simultaneously so several bytes of information can be read all at once. For example, some current microprocessors have registers whose size is 8 bytes, so a memory array is more efficient if it has 8 bytes in a row (often it is desirable to read more than 8 bytes in a single read operation in order to load several registers at the same time). These processors also use instruction sets which require the processor to manipulate (write to) individual bytes. It, accordingly, has been difficult to combine this functionality with the aforementioned design since selective writing to individual

bytes in a memory row would result in a loss of the remaining information in the row due to the clearing of the entire row. It would, therefore, be desirable and advantageous to devise a method of selectively writing to individual bytes in a memory row without losing the remaining information.

Summary of the Invention

It is therefore one object of the present invention to provide an improved memory array for a computer system.

It is another object of the present invention to provide such an array that has a plurality of memory rows, wherein individual bytes may be selectively written to any of the rows.

It is yet another object of the present invention to provide such a memory array which allows clearing of the entire memory row prior to a write operation.

The foregoing objects are achieved in a method of storing information in a memory device having a memory word which includes a plurality of bytes, generally comprising the steps of storing the information from all of the bytes in the memory word in a temporary space, clearing the memory word, writing a portion of the stored information to at least one of the bytes in the memory word, and writing new information in at least one other byte in the memory word. The memory word includes a plurality of memory cells which are grouped to form the bytes, the temporary space is a cache having a plurality of storage locations, and the storing step includes the step of writing information from each of the bytes to a respective one of the storage locations in the cache. The writing steps are accomplished using at least one multiplexer which selectively writes either stored information or new information in response to a control signal. The temporary space (cache) includes a latch for each memory cell, and the multiplexer includes an enable line having on and off states, such that the multiplexer writes information stored in a latch to the

memory word if the enable line is in its off state, but writes new information to the memory word if the enable line is in its on state. The memory word can be one of a plurality of memory words in the memory device, each including a plurality of bytes, and the memory word is accessed using an addressable word line. In a specific implementation, each of the memory cells is connected to a respective transistor, each transistor has a gate connected to the word line, and the clearing step includes the step of turning on the word line to connect each memory cell to ground.

The above as well as additional objectives, features, and advantages of the present invention will become apparent in the following detailed written description.

Brief Description of the Drawings

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives, and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 is a schematic diagram of a prior-art memory array depicting selection of a memory word (row) using an addressable word line;

Figure 2 is a block diagram depicting the method of the present invention by which new data is selectively written to a portion (one or more bytes) of a given memory word;

Figure 3 is a detailed schematic diagram depicting an exemplary circuit which may be used to carry out the present invention; and

Figure 4 is a timing diagram showing the signals associated with selective writing of a portion of the memory word.

Description of the Preferred Embodiment

With reference now to the figures, and in particular with reference to Figure 2, there is depicted a block diagram of one embodiment 20 of a memory-write circuit of the present invention. Memory-write circuit 20 is generally comprised of a memory word 22; a temporary memory device such as cache 24 which reads data from the memory word; another information device 26, such as a computer processor, which provides new data to be stored in memory word 22; and a multiplexer 28 which has two inputs respectively connected to cache 24 and new data device 26. The output of the multiplexer is connected as an input to the memory word 22. Multiplexer 28 is designed to write the contents of cache 24 into memory word 22 except when an enable line 30 connected to multiplexer 28 is turned on. When enable line 30 is activated by the processor or other controller, multiplexer 28 instead writes the new data from device 26 to memory word 22. The original contents of memory word 22 may be stored in the temporary space (cache 24), and memory word 22 is then cleared, but any portion of the original contents may be rewritten to memory word 22 via multiplexer 28. Those portions of the original contents which are not re-written are selectively replaced by the new data. In this manner, individual bytes within the memory word can be selectively written even if memory word 22 is designed so as to require clearing of the entire word before any write operation.

Memory word 22 may be embodied in a conventional memory device, such as an SRAM (static random-access memory) array, which has a plurality of such memory words. In this case, a single multiplexer may be used to write the information to the memory word, or multiple multiplexers may be provided. The memory device supporting memory word 22 alternatively

can be of a novel design, as long as it receives information only after it has been cleared, i.e., the state of each cell in the word has been set to the same (low) voltage. Cache 24 similarly may be of various designs according to the particular application (hardware platform), and multiplexer 28 may be implemented in many ways, as those skilled in the art will appreciate.

One specific implementation is shown in Figure 3. A memory circuit 40 includes a memory cell 42 which is connected to a read circuit 44 and a write circuit 46. Memory circuit 40 is part of a larger circuit that includes other memory cells (not shown) which are generally identical to memory cell 42, and which can be considered as being in columns and rows of a memory array. Additional read and write circuits 44 and 46 are provided, one for each cell in a given row. However, as explained further below, it is not necessary to provide separate read and write circuits for each cell in a given column, although some components of read and write circuits 44 and 46 are duplicated for additional cells in a given column. Memory circuit 40 is specifically designed for an array having 128 cells in a column, arranged in four large groups of 32 cells, these large groups each being further broken down into eight small groups of four cells each.

A memory word is formed by a row of bits (cells), such that the bits in a given row can be collectively read or written to. A plurality of word lines are used to address each memory word; one word line may be used for each row but, in the depicted embodiment, each word has two word lines, a read word line and a write word line. This configuration allows the memory cells to have dual ports

(bit lines) so that two rows can be accessed at the same time. Two (identical) connections 48 and 50 from the read word line are provided for memory cell 42, and two connections, 52 and 54, from the write word line are provided, although the signal from connection 54 is delayed as explained further below. The read port of memory cell 42 is comprised of a first-read output 56 and a second-read output 58, while the write port is comprised of a first-write input 60 and a second-write input 62 (which is connected to the common ground as explained further below). First-read word line 48 is connected to the gate of an n-type field-effect transistor (NFET) 64 whose source is connected to first-read output 56. Second-read word line 50 is connected to the gate of another NFET 66 whose source is connected to second-read output 58. First-write word line 52 is connected to the gate of another NFET 68 whose source is connected to first-write input 60. Second-write word line 54 is connected to the gate of another NFET 70 whose source is connected to second-write input 62. The drain of NFET 64 is connected to the source of another NFET 72 whose drain is connected to ground. The drain of NFET 66 is similarly connected to the source of another NFET 74 whose drain is connected to ground. The gates of both NFETs 72 and 74 are connected to a flip-flop formed by two inverters 76 and 78. The gates of those two NFETs, the source of NFET 70 and the output of inverter 78, are all connected to the input of inverter 76. The output of inverter 76 is connected to the input of inverter 78 and to the source of NFET 68. Those skilled in the art will appreciate that this configuration of transistors and inverters provides a memory cell that is bi-stable and regenerative like a typical SRAM cell.

The operation of read circuit 44 does not pertain to the present invention since the invention is directed to a method of selectively writing to certain cells in a memory word, but the read operation is nevertheless described for the sake of completeness. Read circuit 44 includes three evaluate circuits, 80, 82, and 84, which are connected in series. The first-read output 56 is connected to the input of the first-read evaluate circuit 80, which is connected to the input of an inverter 86 and to the drains of two p-type field-effect transistors (PFETs) 88 and 90. The sources of PFETs 88 and 90 are connected to source voltage (V_{dd}). The gate of PFET 88 is connected to the system clock 89 (signal "c1"). The gate of PFET 90 is connected to the output of inverter 86. The inverter output controls the charge on the gate of another NFET 92, whose drain is connected to ground and whose source is the output of first-read evaluate circuit 80. First-read evaluate circuit 80 is actually connected to four different memory cells, the inputs from the other three non-visible cells being indicated at 94. These four cells constitute one of the small groups referred to above (eight of these small groups together form one of the four large groups in a given 128-cell column). By providing such groupings, a smaller number of components (evaluate circuits) is necessary to support all of the cells. Thirty-two such first-read evaluate circuits 80 are required to receive inputs from all 128 cells in a given column.

The output of first-read evaluate circuit 80 is connected to the input of second-read evaluate circuit 82, which is very similar to first-read evaluate circuit 80. The input of second-read evaluate circuit 82 is connected to the input of another inverter 96 and to the drains of two

more PFETs 98 and 100. The sources of PFETs 98 and 100 are again connected to V_{dd} . The gate of PFET 98 is connected to a delayed clock 99 (signal "c1+"). The gate of PFET 100 is connected to the output of inverter 96. The inverter output controls the charge on the gate of another NFET 102, whose drain is connected to ground and whose source is the output of second-read evaluate circuit 82. Second-read evaluate circuit 82 is actually connected to eight different first-read evaluate circuits, the inputs from the other seven non-visible circuits being indicated at 104. The memory cells corresponding to these eight connections together constitute one of the four large groups in a given 128-cell column. Four such second-read evaluate circuits 82 are required to receive inputs from all 128 cells in a given column.

The output of second-read evaluate circuit 82 is connected to the input of third-read evaluate circuit 84, which is also similar to first- and second-read evaluate circuits 80 and 82. The input of third-read evaluate circuit 84 is connected to the input of another inverter 106 and to the drains of two more PFETs 108 and 110. The sources of PFETs 108 and 110 are again connected to V_{dd} . The gate of PFET 108 is connected to a further-delayed clock 109 (signal "c1++"). The gate of PFET 110 is connected to the output of inverter 106. The inverter output is the output of third-read evaluate circuit 84, and is also the output of read circuit 44. This circuit (including circuits 80, 82, and 84) replaces a conventional sense amplifier. Third-read evaluate circuit 84 is actually connected to four different second-read evaluate circuits, the inputs from the other three non-visible circuits being indicated at 112. Only one third-read evaluate circuit 84 is required to receive inputs from all 128 cells in a given column. The

second-read output 58 is used to similarly generate a second (independent) output for the read operation.

When the memory array is idle and the circuits are precharged, the nodes between the read evaluate circuits are high, and the output of third-read evaluate circuit 84 is low. When a read word line is turned on, if memory cell 42 is in a low state (zero), then the read evaluate circuits will remain in the same state, with the output of the third-read evaluate circuit remaining low. If memory cell 42 is in a high state (one) when the read word line is turned on, then the read evaluate circuits will flip and the output of third-read evaluate circuit will change to high.

The write operation for the memory word is now explained, and particularly how this method allows for writing to selected bytes in the memory word, even though the memory word is completely cleared before the new data is written (the term "bytes" refers to a group of any number of bits, including just one). Write circuit 46 includes three evaluate circuits 114, 116, and 118 which, similarly to read evaluate circuits 80, 82 and 84, are connected in series. First-write input 60 is connected to the first-write evaluate circuit 114, that is, to the input of an inverter 120 and to the drains of two PFETs 122 and 124. First-write input 60 is also connected to a writing NFET 126 discussed further below. The sources of PFETs 122 and 124 are connected to V_{dd} . The gate of PFET 122 is connected to a restore signal 125 (signal "restore bitline"). The gate of PFET 124 is connected to the output of inverter 120. The inverter output controls the charge on the gate of another NFET 128, whose drain is connected to ground and whose source is the output of first-write evaluate circuit 114.

First-write evaluate circuit 114 is actually connected to four different memory cells, the inputs from the other three non-visible cells being indicated at 130 (and corresponding to the same three cells indicated by 94). Thirty-two such first-write evaluate circuits 114 are required to receive inputs from all 128 cells in a given column.

The output of first-write evaluate circuit 114 is connected to the input of second-write evaluate circuit 116, which is similar to first-write evaluate circuit 114. The input of second-write evaluate circuit 116 is connected to the drains of an NFET 132 and a PFET 134. The source of PFET 134 is connected to V_{dd} and its gate is connected to another restore signal 135 (signal "restore group4"). The gate of NFET 132 is connected to a select line 137 (signal "enable read path"), and the source of NFET 132 is connected to the input of another inverter 136 and to the drains of two more PFETs 138 and 140. The sources of PFETs 138 and 140 are again connected to V_{dd} . The gate of PFET 138 is connected to another restore signal 139 (signal "restore group32"). The gate of PFET 140 is connected to the output of inverter 136. The inverter output controls the charge on the gate of another NFET 142, whose drain is connected to ground and whose source is the output of second-write evaluate circuit 116; the output of inverter 136 is also connected to a latch 144 described further below. Second-write evaluate circuit 116 is actually connected to eight different first-write evaluate circuits, the inputs from the other seven non-visible circuits being indicated at 146. Four such second-write evaluate circuits 116 are required to receive inputs from all 128 cells in a given column.

The output of second-write evaluate circuit 116 is connected to the input of third-write evaluate circuit 118. The input of third-write evaluate circuit 118 is connected to the input of another inverter 148 and to the drains of two more PFETs 150 and 152. The sources of PFETs 150 and 152 are again connected to V_{dd} . The gate of PFET 150 is connected to another restore signal 151 (signal "restore group128"). The gate of PFET 152 is connected to the output of inverter 148. The inverter output is the output 149 ("cast out") of third-write evaluate circuit 118. This output, or bit line, is connected to conventional circuitry (such as a sense amplifier) to finish the write evaluation/restore operation. Third-read evaluate circuit 118 is actually connected to four different second-write evaluate circuits, the inputs from the other three non-visible circuits being indicated at 154. Only one third-write evaluate circuit 118 is required to receive inputs from all 128 cells in a given column.

The data value in memory cell 42 is actually written by NFET 126. When the write word line is turned on, first-write word line 52 allows NFET 126 to drive the memory cell via transistor 68. Before this occurs, however, several other steps occur, due to timing of the various clock signals, as can be understood with reference to Figure 4. Two primary signals, 89 and 155 ($c1$ and ci), are used to generate a four-phase clock system. Signal ci is identical in duration to signal $c1$, but it is 90° out of phase. Signal 157 ($c2$) is the complement of signal $c1$, and signal 159 ($ci_$) is the complement of signal ci . These four signals, in turn, allow generation of other signals having a duration which is any multiple of a quarter-cycle.

During the first quarter-cycle, all cell values in the memory word are read and stored to a respective cache element, e.g., latch 144. Latch 144 is essentially an SRAM element (a memory cell), but its output 156 is multiplexed with other data via an NFET 158. The drain of NFET 158 is connected to the device 165 that provides the new data signal, e.g., a register or a cell in another memory array, and the gate of NFET 158 is connected to the enabling device, i.e., the computer processor (signal "byte write" 161). During the first quarter-cycle, NFETs 68 and 132 are turned on (due to the word line signal 52 ("first word line" on Figure 4)), and the enable read path signal 137), and PFETs 122, 134, 138, and 150 are turned off (due to the restore bitline signal 125, restore group4 signal 135, restore group32 signal 139, and restore group128 signal 151). Writing NFET 126 is off during the read/store cycle due to the "drive bitline" signal 163.

During the second quarter-cycle, all of the memory cells in the row are cleared, such as when NFET 70 pulls down the cell to ground due to the signal on the second-write word line 54. That delayed signal is generated using two AND gates 160 and 162. Gate 160 has as its inputs clock signals 155 and 157 (ci and c2); gate 162 has as its inputs the (first) write word line and the output of gate 160. The output of gate 162 drives NFET 70.

During the third quarter-cycle, data is rewritten to the memory cell in response to the drive bitline signal 163 becoming active. The restore group32 signal 139 has turned off this quarter-cycle to avoid race conditions. If the byte write signal 161 is turned on during this quarter-cycle (by the processor), then new data will be written to cell 42

instead of rewriting the previously stored data, thereby achieving the goal of the present invention, namely, selectively overwriting portions of an existing memory word. During the fourth quarter-cycle, the "rst_dsb" signal 167 clears latch 144 so that it can accept new data during the next write operation.

Although the invention has been described with reference to specific embodiments, this description is not meant to be construed in a limiting sense. Various modifications of the disclosed embodiment, as well as alternative embodiments of the invention, will become apparent to persons skilled in the art upon reference to the description of the invention. It is therefore contemplated that such modifications can be made without departing from the spirit or scope of the present invention as defined in the appended claims.



CLAIMS:

1. A method of storing information in a memory device having a memory word which includes a plurality of bytes, comprising the steps of:

storing the information from all of the bytes in the memory word in a temporary space;

clearing the memory word;

writing a portion of the stored information to at least one of the bytes in the memory word; and

writing new information in at least one other byte in the memory word.

2. The method of Claim 1 wherein:

the memory device includes a plurality of memory cells which are grouped to form the bytes;

said temporary space is a cache having a plurality of storage locations; and

said storing step includes the step of writing information from each of said bytes to a respective one of said storage locations in said cache.

3. The method of Claim 1 wherein the memory word is one of a plurality of memory words in the memory device, each including a plurality of bytes, and further comprising the step of accessing the memory word using an addressable word line.

4. The method of Claim 1 wherein said writing steps are accomplished using at least one multiplexer which selectively writes either stored information or new information in response to a control signal.



5. The method of Claim 1 further comprising the step of reading the information from all of the bytes in the memory word prior to said storing step.

6. The method of Claim 1 wherein the memory device has a plurality of cells providing bits for the memory word, and said writing steps are performed simultaneously on all of the bits.

7. The method of Claim 2 wherein said writing steps include the steps of writing stored information into a first plurality of the memory cells, and writing new information into a second plurality of the memory cells.

8. The method of Claim 2 wherein said writing steps are accomplished using at least one multiplexer which selectively writes either stored information or new information in response to a control signal.

9. The method of Claim 3 wherein:
the memory word has a plurality of bi-stable memory cells;

each of the memory cells is connected to a respective transistor, and each transistor has a gate connected to the word line; and

said clearing step includes the step of turning on the word line to connect each memory cell to ground.



10. The method of Claim 4 wherein:

the temporary space includes at least one latch;

the multiplexer includes an enable line having on and off states, such that the multiplexer writes information stored in the latch to the memory word if the enable line is in its off state, but writes new information to the memory word if the enable line is in its on state.

11. The method of Claim 9 wherein all memory cells in the memory word are cleared simultaneously.





- 21 -

12. A device for storing and retrieving data, comprising:

a memory device having a memory word formed of a plurality of memory cells;

a cache for temporarily storing data contained in the memory word;

multiplexer means for selectively writing stored data from said cache or new data to different cells in the memory word; and

processor means for controlling said multiplexer means.

13. The device of Claim 12 wherein said memory device is a RAM array and the memory word is one of a plurality of memory words in said RAM array, the memory word being accessed by an addressable word line connected to said processor means.

14. The device of Claim 12 further comprising means for clearing the memory word after data from the memory word has been stored in said cache.

15. The device of Claim 12 wherein said cache includes a plurality of latches, one for each memory cell,

16. The device of Claim 12 wherein said processor means is connected to said multiplexer means by an enable line having on and off states, such that said multiplexer means writes data stored in said cache to said memory word if said enable line is in said off state, but writes new data to said memory word if said enable line is in said on state.



* C L M S *

17. The device of Claim 13 wherein:

each of said memory cells is connected to a first transistor providing a write port;

each of said memory cells is further connected to a second transistor providing a ground connection; and

said word line is connected to first and second word lines, said first word line being connected to the gates of each of said first transistors, and said second word line being connected to the gates of each of said second transistors such that said memory cells are cleared when said second word line is turned on.

18. The device of Claim 13 wherein each of said plurality of memory words is connected to a respective word line such that each memory word can be addressed by said processor means, and further comprising a plurality of caches, respectively connected to said plurality of memory words, and a plurality of multiplexers respectively connected to said plurality of caches such that stored data or new data can be selectively written to portions of each of said memory words.

19. The device of Claim 17 wherein said second word line is connected to said second transistor via a delay circuit such that data contained in a given memory cell is stored in said cache before the memory cell is cleared.

