

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
28 January 2010 (28.01.2010)

PCT

(10) International Publication Number
WO 2010/011503 A2

(51) International Patent Classification:

G11C 5/02 (2006.01) G11C 7/00 (2006.01)
G11C 7/22 (2006.01) H01L 23/12 (2006.01)

(21) International Application Number:

PCT/US2009/050155

(22) International Filing Date:

9 July 2009 (09.07.2009)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

12/176,951 21 July 2008 (21.07.2008) US

(71) Applicant (for all designated States except US): **MICRON TECHNOLOGY, INC.** [US/US]; 8000 South Federal Way, Boise, ID 83716-9632 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **LABERGE, Paul, A.** [US/US]; 1335 Meadow Avenue, Shoreview, MN 55126 (US). **JEDDELOH, Joseph, M.** [US/US]; 4302 Reiland Lane, Shoreview, MN 55126 (US). **JOHNSON, James, B.** [US/US]; 5977 East Gateway Drive, Boise, ID 83716 (US).

(74) Agents: **BULCHIS, Edward, W.** et al.; Dorsey & Whitney LLP, US Bank Centre, 1420 5th Ave, Suite 3400, Seattle, WA 98101 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: MEMORY SYSTEM AND METHOD USING STACKED MEMORY DEVICE DICE, AND SYSTEM USING THE MEMORY SYSTEM




FIGURE 2

(57) Abstract: A memory system and method uses stacked memory device dice coupled to each other and to a logic die. The logic die may include a timing correction system that is operable to control the timing at which the logic die receives signals, such as read data signals, from each of the memory device dice. The timing correction controls the timing of the read data or other signals by adjusting the timing of respective strobe signals, such as read strobe signals, that are applied to each of the memory device dice. The memory device dice may transmit read data to the memory device at a time determined by when it receives the respective strobe signals. The timing of each of the strobe signals is adjusted so that the read data or other signals from all of the memory device dice are received at the same time.

WO 2010/011503 A2

**MEMORY SYSTEM AND METHOD USING STACKED MEMORY
DEVICE DICE, AND SYSTEM USING THE MEMORY SYSTEM**

TECHNICAL FIELD

[001] This invention relates to memory devices, and, more particularly, in one or more embodiments to a memory system having a plurality of stacked memory device dice connected to a logic die.

BACKGROUND OF THE INVENTION

[002] As memory devices of all types have evolved, continuous strides have been made in improving their performance in a variety of respects. For example, the storage capacity of memory devices has continued to increase at geometric proportions. This increased capacity, coupled with the geometrically higher operating speeds of electronic systems containing memory devices, has made high memory device bandwidth ever more critical. One application in which memory devices, such as dynamic random access memory (“DRAM”) devices, require a higher bandwidth is their use as system memory in computer systems. As the operating speed of processors has increased, processors are able to read and write data at correspondingly higher speeds. Yet conventional DRAM devices often do not have the bandwidth to read and write data at these higher speeds, thereby slowing the performance of conventional computer systems. This problem is exacerbated by the trend toward multi-core processors and multiple processor computer systems. It is currently estimated that computer systems operating as high-end servers are idle as many as 3 out of every 4 clock cycles because of the limited data bandwidth of system memory devices. In fact, the limited bandwidth of DRAM devices operating as system memory can reduce the performance of computer systems to as low as 10% of the performance of which they would otherwise be capable.

[003] Various attempts have been made to increase the data bandwidth of memory devices. For example, wider internal data buses have been used to transfer data to and from arrays with a higher bandwidth. However, doing so usually requires that write data be serialized and read data deserialized at the memory device

interface. Another approach has been to simply scale up the size of memory devices or conversely shrink their feature sizes, but, for a variety of reasons, scaling has been incapable of keeping up with the geometric increase in the demand for higher data bandwidths. Proposals have also been made to stack several integrated circuit memory device dice in the same package, but doing so threatens to create a large number of other problems that must be overcome.

[004] One potential problem with stacking memory device dice on top of each other is that it may create signal timing skews between the signals transmitted to or from each of the memory devices. Insofar as the distances between each of the memory devices and an interface for the packaged memory devices will vary for each memory device, the time required for signals to be transmitted to and from each of the memory devices will inherently vary. This can be a considerable problem because there may be a large number of memory device dice in the stack, such as, for example, eight memory devices. Additionally, because of process, temperature and supply voltage variations, the timing performances of the memory devices may vary even if they are fabricated on the same wafer. An example of such signal timing skews is illustrated in Figure 1, which shows the period during which read data signals are considered valid at a package interface for each of 4 stacked dynamic random access memory (“DRAM”) device dice DRAM0-DRAM1. This data valid period is sometimes referred to as a data “eye.” As shown therein, the read data for DRAM2 is valid first, followed by DRAM0, DRAM1 and finally DRAM3. The period during which all of the read data, *i.e.*, the composite eye 8 for all of the DRAM, is almost nonexistent. Therefore, it would be very difficult for a memory access device, such as a memory controller or processor, to capture the read data using a single clock signal, particularly as the operating speeds and resulting data transfer rates of memory devices continue to increase.

[005] In the past, the problem of signal skews from different memory devices has been greatly alleviated by transmitting respective read strobe signals from the memory devices along with the respective read data. The strobe signal is then used by the memory access device to capture the read data. Insofar as differences in the timing of read data from each of the memory devices are substantially matched by differences in the timing of the strobe signals, transitions of

the strobe signals are substantially centered in the data eye from each memory device, thereby allowing the memory access device to successfully capture the read data from each of the memory devices. As the operating speed of memory devices have continued to increase, even this approach has not been sufficient. As a result, techniques have been developed to adjust the timing of the strobe signals by either adjusting their transmit times at the memory devices or by delaying them by adjustable amounts in the memory access device. Alternatively, the timing of each of the bits of the read data can be adjusted relative to the timing of the read strobe signal. An example of a memory device that adjusts the timing of each bit of read data in this manner is described in U.S. Patent No. 6,882,304.

[006] The conventional approach of adjusting the timing between a read strobe signal and read data signals could be used for stacked memory device dice. However, doing so would require an extensive amount of timing adjustment circuitry in each memory device, thereby reducing the area of each memory device dice available for providing memory capacity. Adjusting the timing between a read strobe signal and read data signal in each memory device die would also require sending a read strobe signal from each memory device. Furthermore, although the timing problems have been discussed with respect to read data signals, essentially the same type of problems can exist with write data signals, command signals and address signals. If a separate strobe signal was transmitted to or from each memory device for each of these types of signals, the strobe signals would require that the packaged memory devices include a large number of strobe terminals. For example, if eight memory device dice were stacked, 32 terminals would be required to transfer a strobe signal to or from the memory devices of all these types of signals. Yet, it is generally considered undesirable to unduly increase the number of terminals in a memory device because of the lack of available area in a memory device package and the large number of conductors that would be required in the bus or circuit board on which the memory device was mounted.

[007] Therefore, a need exists for a method and apparatus to minimize problems and limitations resulting from timing skews between signals transmitted to or from stacked memory device dice in a manner that maximizes the area of a die

available for memory capacity and does not unduly increase the number of required terminals.

BRIEF DESCRIPTION OF THE DRAWINGS

- [008]** Figure 1 is a timing diagram showing the manner in which the timing of read data signals from each of a plurality of stacked memory device dice are skewed relative to each other.
- [009]** Figure 2 is a block diagram of a processor-based system that includes a memory system containing stacked memory device dice according to an embodiment of the invention.
- [010]** Figure 3 is a timing diagram showing the timing of read data signals shown in Figure 1 along with read data signal having corrected timing according to an embodiment of the invention.
- [011]** Figure 4 is a more detailed block diagram of an embodiment of the processor-based system of Figure 2 showing the memory system in greater detail.
- [012]** Figure 5 is a block diagram showing an embodiment of a timing correction system used in the memory system of Figures 2 and 4.
- [013]** Figure 6 is a block diagram showing an embodiment of a strobe timing adjustment circuit used in the timing correction system of Figure 5.
- [014]** Figure 7 is a block diagram showing an embodiment of a receiver used in the timing correction system of Figure 5.

DETAILED DESCRIPTION

- [015]** A computer system including a high-capacity, high bandwidth memory system 10 according to an embodiment of the invention is shown in Figure 2. The memory system 10 is connected to a processor 12 through a relatively narrow high-speed bus 14 that is divided into downstream lanes and separate upstream lanes (not shown in Figure 2). The memory system 10 includes 4 DRAM dice 20, 22, 24, 26, which may be identical to each other, stacked on top of each other. Although the memory system 10 includes 4 DRAM dice 20, 22, 24, 26, other embodiments of the memory device use a greater or lesser number of DRAM dice.

The stacked DRAM dice 20, 22, 24, 26 are connected to a logic die 30, which serves as the interface with the processor 12. The logic die 30 can be physically positioned relative to DRAM dice 20, 22, 24, 26 in any order, such as by stacking the DRAM dice 20, 22, 24, 26 on top of the logic die 30. However, the logic die 30 could, for example, be positioned in the middle of the stack of DRAM dice 20, 22, 24, 26.

[016] The logic die 30 can implement a variety of functions in the memory system 10, such as to limit the number of functions that must be implemented in the DRAM dice 20, 22, 24, 26. For example, the logic die 30 may perform memory management functions, such as power management and refresh of memory cells in the DRAM dice 20, 22, 24, 26. In some embodiments, the logic die 30 may perform error checking and correcting (“ECC”) functions. In the embodiments described herein, the logic die 30 implements the timing correction of signals that the logic die 30 couples to or receives from the DRAM dice 20, 22, 24, 26, as explained in greater detail below.

[017] The DRAM dice 20, 22, 24, 26 may be connected to each other, and they are connected to the logic die 30 by a relatively wide bus 34. The bus 34 may be implemented with through silicon vias (“TSVs”), which comprise a large number of conductors extending at least partially through the DRAM dice 20, 22, 24, 26 at the same locations on the DRAM dice and connect to respective conductors formed on the dice 20, 22, 24, 26. In one embodiment, each of the DRAM dice 20, 22, 24, 26 are divided into 16 autonomous partitions, each of which may contain 2 or 4 independent memory banks. In such case, the partitions of each dice 20, 22, 24, 26 that are stacked on top of each other may be independently accessed for read and write operations. Each set of 16 stacked partitions may be referred to as a “vault.” Thus, the memory system 10 may contain 16 vaults.

[018] As mentioned above, the logic die 30 contains timing correction circuitry for ensuring that read data signals received from the memory device dice 20, 22, 24, 26 are aligned with each other when they are transmitted from the memory system 10. As shown in Figure 3, the logic die 30 causes the read data signals that would be received from the memory device dice 20, 22, 24, 26 as shown in Figure 1 so that they are substantially aligned with each other. As a result, the read data signals are transmitted from the memory system 10 with a composite eye

28 that is substantially larger than the virtually non-existent composite eye 8 shown in Figure 1.

[019] As explained in greater detail below, one of the functions performed by the logic die 30 is to serialize the read data bits coupled from the DRAM dice 20, 22, 24, 26 into a serial stream of 16 serial data bits coupled through each of 16 parallel bits of one of the upstream lanes 42a-d of the bus 14. Similarly, the logic die 30 may perform the functions of deserializing 16 serial data bits coupled through one of the 16-bit downstream lanes 40a-d of the bus 14 to obtain 256 parallel data bits. The logic die 30 then couples these 256 bits through one of the 32-bit sub-buses 38a-p in a serial stream of 8 bits. However, other embodiments may use different numbers of lanes 40, 42 having different widths or different numbers of sub-buses 38a-p having different widths, and they may couple data bits having different structures. As will be appreciated by one skilled in the art, the stacking of multiple DRAM dice results in a memory device having a very large capacity. Further, the use of a very wide bus connecting the DRAM dice allows data to be coupled to and from the DRAM dice with a very high bandwidth.

[020] A logic die 30 according to an embodiment of the invention is shown in Figure 4 connected to the processor 12 and the DRAM dice 20, 22, 24, 26. As shown in Figure 4, each of the 4 downstream lanes 40a-d is connected to a respective link interface 50a-d. Each link interface 50a-d includes a deserializer 54 that converts each serial stream of 16 data bits on each of the 16-bit lanes 40a-d to 256 parallel bits. Insofar as there are 4 link interfaces 50a-d, the link interfaces can together output 1024 output parallel bits.

[021] Each of the link interfaces 50a-d applies its 256 parallel bits to a respective downstream target 60a-d, which decodes the command and address portions of the received packet and buffers write data in the event a memory request is for a write operation. The downstream targets 60a-d output their respective commands, addresses and possibly write data to a switch 62. The switch 62 contains 16 multiplexers 64 each of which direct the command, addresses and any write data from any of the downstream targets 60a-d to its respective vault of the DRAM dice 20, 22, 24, 26. Thus, each of the downstream targets 60a-d can access any of the 16 vaults in the DRAM dice 20, 22, 24, 26. The multiplexers 64 use the address in the

received memory requests to determine if its respective vault is the target of a memory request. Each of the multiplexers 64 apply the memory request to a respective one of 16 vault controllers 70a-p.

[022] Each vault controller 70a-p includes a respective memory controller 80, each of which includes a write buffer 82, a read buffer 84 and a command pipeline 86. The commands and addresses in memory requests received from the switch 62 are loaded into the command pipeline 86, which subsequently outputs the received commands and corresponding addresses. Any write data in the memory requests are stored in the write buffer 82. The read buffer 84 is used to store read data from the respective vault, as will be explained in greater detail below. Both the write data from the write buffer 82 and the commands and addresses from the command pipeline 86 of each of the vault controllers 70a-p are applied to a memory interface 88. The memory interface 88 couples commands and addresses from the command pipelines 86 to the DRAM dice 20, 22, 24, 26 through a command/address bus 94, and it couples 32-bits of write data from each of the write buffers 82 to the DRAM dice 20, 22, 24, 26 through a 32-bit data bus 92.

[023] Although data are loaded into the write buffer 82 as 256 parallel bits, they are output from the buffer 82 in two sets, each set being 128 parallel bits. These 128 bits are then further serialized by the memory interface 88 to 4 sets of 32-bit data, which are coupled through the data bus 92. In the embodiment shown in Figure 4, write data are coupled to the write buffer 82 in synchronism with a 500 MHz clock so the data are stored in the write buffer at 16 gigabytes (“GB”) per second. The write data are coupled from the write buffer 82 to the DRAM dice 20, 22, 24, 26 using a 2 GHz clock so the data are output from the write buffer 82 at 8 GB/s. Therefore, as long as more than half of the memory requests are not write operations to the same vault, the write buffers 82 will be able to couple the write data to the DRAM dice 20, 22, 24, 26 at least as fast as the data are coupled to the write buffer 82.

[024] In the event a memory request is for a read operation, the command and address for the request are coupled to the DRAM dice 20, 22, 24, 26 in the same manner as a write request, as explained above. In response to a read request, 32 bits of read data are output from the DRAM dice 20, 22, 24, 26 through the 32-bit data

bus 92. The memory interface 88 deserializes the 32 bits of read data from each of the vaults into two sets of 128-bit read data, which are applied to the read buffer 84. After 2 sets of 128-bit read data have been stored in the read buffer 84, the read buffer transmits 256 bits to the switch 62. The switch includes 4 output multiplexers 104 coupled to respective upstream masters 110a-d. Each multiplexer 104 can couple 256 bits of parallel data from any one of the vault controllers 70a-p to its respective upstream master 110a-d. The upstream masters 110a-d format the 256 bits of read data into packet data and couple the packet to respective upstream link interfaces 114a-d. Each of the link interfaces 114a-d include a respective serializer 120 that converts the incoming 256 bits to a serial stream of 16 bits on each bit of a respective one of the 16-bit upstream links 42a-d.

[025] As explained above, the logic die 30 corrects the timing of signals that the logic die 30 couples to or receives from the DRAM dice 20, 22, 24, 26, as explained in greater detail below. This timing correction is implemented by a timing correction system 100, one of which is included in each of the memory interfaces 88. An embodiment of the timing correction system 100 is shown in Figure 5. The system 100 includes a timing control circuit 134 that outputs 4 read data strobe signals RDQS0-RDQS3. The strobe signals RDQS0-RDQS3 are in a pattern of alternating high and low logic levels. Each of the strobe signals RDQS0- RDQS3 is applied to a respective one of 4 strobe timing adjustment circuits 140a-d, each of which receives a respective timing control signal from the timing control circuit 134. Each of the 4 strobe timing adjustment circuits 140a-d transmit respective adjusted read data strobe signals RDQS0'-RDQS3' to a respective one of the DRAM device dice 20, 22, 24, 26. Thus, the strobe timing adjustment circuit 140a transmits its strobe signal RDQS0' to the first DRAM device die 20 through a TSV 152a extending from the logic die 30 to the DRAM device die 20. The strobe signal RDQS0' may also be coupled through additional TSVs 152b-d extending to the other DRAM device dice 22, 24, 26, but they are not internally connected to any circuitry in the respective DRAM device dice 22, 24, 26. Similarly, the strobe timing adjustment circuit 140b transmits its strobe signal RDQS1' to the second DRAM device die 22 through two TSVs 154a,b, the strobe timing adjustment circuit 140c transmits its strobe signal RDQS2' to the third DRAM device die 24 through

three TSVs 156a,b,c, and the strobe timing adjustment circuit 140d transmits its strobe signal RDQS3' to the fourth DRAM device die 26 through four TSVs 158a,b,c,d.

[026] In response to each of the read data strobe signals RDQS0'-RDQS3', the respective DRAM device dice 20, 22, 24, 26 output 32 bits of read data on a common 32-bit bus 150. Each of the bits of the read data is applied through a respective input buffer 162 to a respective one of 32 receivers 160 (only 1 of which is shown in Figure 5). Thus, each of the memory interfaces 88 shown in Figure 4 contains 32 receivers 160. Each receiver 160 deserializes the received bit of read data into 4 bits of read data. Thus, the 32 bits of read data collectively received by the 32 receivers 160 result in 128 bits of read data. In a similar manner, a respective transmitter 164 is connected to each bit of the bus 150 so that each of the memory interfaces 88 shown in Figure 4 contains 32 transmitters 164 (only 1 of which is shown in Figure 5). Each of the transmitters 164 transmits 4 bits of write data through a respective output buffer 166 and serializes the data into 1 bit of write data that is applied to its respective bit of the bus 150.

[027] The timing of the read data being transmitted from each of the DRAM device dice 20, 22, 24, 26 is controlled by the timing of its respective read data strobe signals RDQS0'-RDQS3'. The timing control signals output from the timing control circuit 134 cause the respective strobe timing adjustment circuits 140a-d to properly adjust the timing of the read data strobe signals RDQS0'-RDQS3' so that the read data is received at the logic die 30 from each of the DRAM device dice 20, 22, 24, 26 with the same timing. As a result, the receivers 160 can capture the read data using the same clock signal CLK regardless of which DRAM device dice 20, 22, 24, 26 transmitted the read data. Therefore, the logic die 30 can transmit the read data from the memory system with the same timing without any additional timing adjustments regardless of which DRAM device dice 20, 22, 24, 26 was the originator of the data.

[028] In one embodiment of the timing control circuit 134, a conventional training sequence is used to determine the correct timing of each of the strobe signals RDQS0'-3', such as is disclosed in U.S. Patent No. 6,882,304. The timing control circuit 134 causes the timing of each of the strobe signals RDQS0'-3' to be

incrementally adjusted while the receivers 160 attempt to capture known read data. The timing that best captures the read data is then used during normal operation. The optimum timing can be determined, for example, by using a timing that is midway between all of the timing values that successfully captured the read data.

[029] The disclosed embodiment of the memory interface 88 uses the timing control circuit 100 only to adjust the timing of the read data. However, in other embodiments, a similar timing control circuit 100 adjusts the timing of write data signals, command signals and/or address signals in a similar manner. In each case, the timing control circuit can vary the timing of these signals transmitted to the DRAM device dice 20, 22, 24, 26 so they are received by the DRAM device dice at the proper time. The proper timing can be determined by conducting a read operations as the timing of these signals are incrementally varied to determine which timing can best be captured in the DRAM device dice 20, 22, 24, 26. Also, although the timing control circuit 100 is used to apply properly timed read strobe signals to 4 DRAM device dice 20, 22, 24, 26, other embodiments are used with lesser or additional numbers of DRAM device dice.

[030] An embodiment of each of the strobe timing adjustment circuits 140a-d is shown in Figure 6. Although only one timing adjustment circuit 140 is shown in Figure 6, it will be understood that 4 of them would be used in the timing control circuit 100 of Figure 5. As explained above, the timing adjustment circuit 140 receives a respective read data strobe signal RDQS, which consists of alternating high and low logic levels, which may be labeled A-H. The RDQS signal is applied to two registers 170, 172, each of which has 4 outputs Q0-Q3. The registers 170, 172 are alternately enabled by respective complementary enable signals received from respective outputs of a flip-flop 176. The flip-flop 176 is configured to toggle so that it switches states responsive to each rising edge of a clock signal CLK_1XSE. In one embodiment, the clock signal CLK_1XSE has a frequency of 500 MHz, so that the flip-flop 176 outputs signal Q high for 4 ns and then output Q* high for 4 ns. Therefore, the flip-flop 176 first enables the register 170 for 4 ns and then enables the register 172 for 4 ns. However, in other embodiments, the clock signal CLK_1XSE has different frequencies. Therefore, the register 170 outputs the signals designated A-D in the read data strobe signal RDQS,

and the register 172 outputs the signals designated E-H in the read data strobe signal RDQS.

[031] The signals output by registers 170, 172 are applied to two multiplexers 180, 182. More specifically, the signals designated A, C, E, G from the registers 170, 172 are applied to respective inputs of the multiplexer 180, and the signals designated B, D, F, H from the registers 170, 172 are applied to respective inputs of the multiplexer 182. The multiplexers 180, 182 are controlled by respective select signals to cause one of its inputs to be applied to its output. The select signals are generated by a selector 188, which may be implemented using a specially configured counter. The selector 188 is clocked by a clock signal CLK_4X, which, in one embodiment, has a frequency of 1 GHz and is a differential clock signal so that the clock signals have different states. However, the clock signals CLK_4X are coupled through a delay circuit 190, which may be, for example, a delay line or a phase interpolator. The delay circuit 190 delays the clock signals CLK_4X by an adjustable delay that is controlled by the respective timing control signal from the timing control circuit 134 (Figure 5). The selector 188 first causes the multiplexer 180 to output the signal A from the register 170, and it then causes the multiplexer 182 to output the signal B from the register 170. Similarly, the selector 188 then causes the multiplexer 180 to output the signal C from the register 170, and it then causes the multiplexer 182 to output the signal D from the register 170. In like manner, the selector 188 causes the multiplexer 180, 182 to sequentially output the signals E-H. The respective outputs of the multiplexers 180, 182 are applied to respective inputs of a serializer 194.

[032] In addition to the clock signals CLK_4X being applied to the selector 188, one of the clock signals CLK_4X is applied to a clock input of the serializer 194. The clock signal alternately selects each of its inputs and couples them to its output. Thus, although signals designated A,C,E,G are applied to one input and signals B,D,F,H are applied to the other input, the serializer 194 outputs a signal in the order A,B,C,D,E,F,G,H. As a result, the serializer 194 outputs the original read strobe signal RDQS applied to the registers 170, 172 except that the timing of the strobe signal RDQS has been adjusted by the timing control signal to generate the timing adjusted read strobe signal RDQS'. This RDQS' signal is transmitted though

an output buffer 196 to its respective DRAM device dice 20, 22, 24, 26, as explained above with reference to Figure 5. Thus, the selector 188, the multiplexers 180, 182 and serializer 188 together implement a serializing circuit.

[033] An embodiment of the receiver 160 shown in Figure 5 is shown in Figure 7. The receiver 160 receives one bit of read data DQ, which, in one embodiment, consists of 8 bits of serial read data, and deserializes it as explained above with reference to Figure 5. These serial bits of read data DQ will be designated bits A-H for purposes of facilitating an explanation of the receiver 160. The read data bit is applied to the respective data input of 8 flip-flops 200a-h after being coupled through an input buffer 204. The flip-flops 200a-h are clocked by a selector 208, which may be implemented by a counter. However, alternating flip-flops 200a,c,e,g are clocked by the rising edge of a signal received from the selector 208, while the flip-flops 200b,d,f,h are clocked by the falling edge of a signal received from the selector 208. Additionally, the same signal is applied to adjacent flip flop 200a,b through 200g,h.

[034] When the selector 208 is clocked, each of its outputs sequentially transitions high responsive to rising edges of the clock signal CLK_2X, which, in one embodiment, has a frequency of 1GHz. Thus the flip-flop 200a first outputs the serial read data bit A, and continues to do so until it is clocked again 4 ns later when the received signal from the selector 208 again transitions high. The flip-flop 200b is clock 1 ns later by the falling edge of the same signal that clocked the flip-flop 200a, so that the flip-flop 200b outputs the serial read data bit B for 4 ns. However, at the same time that the signal applied to the flip-flops 200a,b transitions low, the signal applied to the flip-flops 200c,d transitions high to cause the flip-flop 200c to output the serial read data bit C for 4 ns. In a similar manner, the flip-flops 200d-g sequentially output read data bits designated D-G. At the end of 4 ns, all of the serial data bits A-G will have been output from the flip-flops 200a-h.

[035] The output from each of the first 4 flip flops 200a-d are applied to a first input of a respective multiplexers 210a-d, and the output from each of the second 4 flip flops 200e-h are applied to a second input of the respective multiplexers 210a-d. The multiplexers 210a-d are each controlled by a signal received from a flip-flop 212, which is configured to toggle responsive to a signal

from an AND gate 214. The AND gate 214 receives at one of its inputs the clock signal CLK_1XSE which, it may be recalled, has a frequency of 500 MHz in one embodiment. The other input of the AND gate 214 receives a signal from the output of a flip-flop 216. The flip-flop 216 has a data input that receives a signal from the output of a flip-flop 218. The flip-flop 218 receives a read enable signal RD_EN at its data input, and it is clocked by the clock signal CLK_1XSE.

[036] In operation, when the read enable signal RD_EN transitions high, the next rising edge of the clock signal CLK_1XSE causes the flip-flop 218 to output a high that is applied to the data input of the flip-flop 216. This high output signal from the flip-flop 218 is also applied to the enable input of the selector 208 to allow it to begin clocking the flip-flops 200a-h so that they can collectively output the serial read data bits. On the next rising edge of the clock signal CLK_1XSE, the flip-flop 216 transitions high, thereby causing the output of the AND gate 214 to transition high. The flip-flop 212 is then clocked, and it continues to be clocked by the clock signal CLK_1XSE until the RD_EN signal transitions low at the end of a read operation. Insofar as the flip-flop 212 is configured to toggle, it outputs a signal that is high for 4 ns and then low for 4 ns. As a result, the multiplexers 210a-d output serial data bits A-D for 4 ns, and they then output serial data bits E-H for 4 ns. Thus, the 8 ns that it takes to output bits A-H coincides with the 8 ns that the 8 serial data bits A-H are applied to the receiver 160. Of course, the timing and frequency examples have been provided herein for purposes of explanation, and may be different for other embodiments.

[037] 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 embodiments of the invention are explained in the context of stacked DRAM dice, it will be understood that the stacked die may be other types of memory device die, such as flash memory device die. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

What is claimed is:

1. A memory system, comprising:
a plurality of stacked memory device dice containing a plurality of memory cells; and
a logic die coupled to the memory device dice through a plurality of conductors, the logic die being operable to write data to and read data from the memory device dice, the logic die including a timing correction system that is operable to control the timing at which at least one signal is received by the logic die from each of the memory device dice.
2. The memory system of claim 1 wherein the plurality of stacked memory device dice are connected to each other and to the logic die through a plurality of through silicon vias.
3. The memory system of claim 1 wherein the at least one signal coupled between the logic die and each of the memory device dice comprise a respective set of read data signals transmitted to the logic die by each of the memory device dice.
4. The memory system of claim 2 wherein each of the memory device dice is operable to transmit its respective set of read data signals at a time determined by receipt of a respective read data strobe, and wherein the timing correction system comprises:
a strobe timing adjustment circuit for each of the memory device dice, each of the strobe timing adjustment circuits being structured to output a read strobe signal with a timing that is controlled by a respective timing control signal;
and
a timing control circuit generating the respective timing control signals and applying the timing control signals to the respective memory device dice, the timing control circuit generating the timing control signals causing the

respective sets of read data transmitted by the memory device dice to be received by the logic die with substantially the same timing.

5. The memory system of claim 4 wherein the timing control circuit is operable to generate each of the timing control signals by providing a respective plurality of timing control signals to each of the strobe timing adjustment circuits that cause the respective strobe timing adjustment circuit to vary the timing of the respective read strobe signal over a range thereby causing the respective memory device die to provide its respective set of read data signals to the logic die at times that vary over a range, the timing circuit being operable to use as the timing control signal for application to the respective strobe timing adjustment circuit a timing control signal that causes the set of read data signals to be received by the logic die at a suitable time within the range.

6. The memory system of claim 2 wherein the timing correction system comprises a plurality of data receivers each of which is operable to receive a corresponding read data signal from each of the memory device dice, the receiver being operable to deserialize the received read data signal into a plurality of parallel read data bits.

7. The memory system of claim 6 wherein each of the receivers is operable to capture the received read data signal responsive to a clock signal, all of the plurality of receivers using the same clock signal to capture their respective received read data signal.

8. The memory system of claim 6 wherein each of the receivers comprise:

a plurality of flip-flops coupled to receive the read data signal, each of the flip-flops being operable to store and apply to an output a respective sample of the read data signal;

a selector operable to generate and sequentially apply the clock signals to the respective flip-flops so that the flip-flops store the respective samples of the read data signal; and

a plurality of multiplexers each of which is coupled to the respective outputs of a plurality of the flip-flops, the multiplexers being responsive to a clock signal to alternately output the respective samples from a respective pair of the plurality of flip-flops.

9. The memory system of claim 2 wherein each of the memory device dice comprise a respective dynamic random access memory device dice.

10. A timing correction system, comprising:

a circuit generating a plurality of data strobe signals;

a plurality of strobe timing adjustment circuits receiving respective ones of the data strobe signals, each of the strobe timing adjustment circuits comprising:

a register storing the respective data strobe signal and applying a plurality of bits of the stored data strobe signal to respective output terminals;

a serializing circuit receiving the bits of the data strobe signal from the register, the serializing circuit being operable to output the bits of the data strobe signal in serial form at a time determined by an adjusted clock signal; and

a delay circuit receiving a clock signal and outputting the adjusted clock signal with a delay corresponding to a timing control signal; and

a timing control circuit generating a plurality of timing control signals each of which is applied to the delay circuit in a respective one of the strobe timing adjustment circuits.

11. The timing correction system of claim 10 wherein the delay circuit comprises a phase interpolator.

12. The timing correction system of claim 10 wherein the delay circuit comprises a delay line.

13. The timing correction system of claim 10 wherein the register comprises:

a first register storing a first plurality of successive bits of the data strobe signal; and

a second register storing a second plurality of successive bits of the data strobe signal, the second plurality of successive bits of the data strobe signal immediately following the first plurality of successive bits of the data strobe signal.

14. The timing correction system of claim 13 wherein the serializing circuit comprises:

a first multiplexer coupled to receive alternating bits of the data strobe signal output from the first register and alternating bits of the data strobe signal output from the second register, the first multiplexer being operable responsive to a first control signal to couple each of the received bits of the data strobe signal to a first output terminal;

a second multiplexer coupled to receive alternating bits of the data strobe signal output from the first register that are not applied to the first multiplexer and alternating bits of the data strobe signal output from the second register that are not applied to the first multiplexer, the second multiplexer being operable responsive to a second control signal to couple each of the received bits of the data strobe signal to a second output terminal;

a selector operable responsive to a clock signal to generate the first and second control signals, the first and second control signals causing the first and second multiplexers to sequentially couple each of the received bits of the data strobe signal to the first and second output terminals, respectively; and

a serializer circuit coupled to the first and second output terminals of the first and second multiplexers, respectively, the serializer circuit being operable to alternately output bits of the data strobe signal received from the first and second multiplexers at a time determined by the adjusted clock signal.

15. The timing correction system of claim 10 wherein the timing control circuit is operable to generate each of the timing control signals by providing a respective plurality of timing control signals to each of the strobe timing adjustment circuits that cause the respective strobe timing adjustment circuit to vary the timing of the respective read strobe signal over a range thereby causing the respective memory device die to provide its respective set of read data signals to the logic die at times that vary over a range, the timing circuit being operable to use as the timing control signal for application to the respective strobe timing adjustment circuit a timing control signal that causes the set of read data signals to be received by the logic die at a suitable time within the range.

16. A system, comprising:
a memory access device; and
a memory system coupled to the memory access device, the memory system comprising:

a plurality of stacked memory device dice containing a plurality of memory cells and being operable to transmit read data responsive to a received read strobe signal, the timing at which each of the memory device dice transmitting the read data being determined by the time of receiving the respective read strobe signal; and

a logic die coupled to the memory access device and to the memory device dice through a plurality of conductors, the logic die being operable to write data to and read data from the memory device dice, the logic die including a timing correction system that is operable to control the timing that the read strobe signals are applied to the respective memory device dice, the logic die being operable to transmit the read data to the memory access device substantially upon receipt from the memory device dice.

17. The system of claim 16 wherein the timing correction system comprises:

a strobe timing adjustment circuit for each of the memory device dice, each of the strobe timing adjustment circuits being structured to output the respective read strobe signal with a timing that is controlled by a respective timing control signal; and

a timing control circuit generating the respective timing control signals and applying the timing control signals to the respective memory device dice, the timing control circuit generating the timing control signals causing the respective sets of read data transmitted by the memory device dice to be received by the logic die with substantially the same timing.

18. The system of claim 17 wherein the timing control circuit is operable to generate each of the timing control signals by providing a respective plurality of timing control signals to each of the strobe timing adjustment circuits that cause the respective strobe timing adjustment circuit to vary the timing of the respective read strobe signal over a range thereby causing the respective memory device die to provide its respective set of read data signals to the logic die at times that vary over a range, the timing circuit being operable to use as the timing control signal for application to the respective strobe timing adjustment circuit a timing control signal that causes the set of read data signals to be received by the logic die at a suitable time within the range.

19. The system of claim 16 wherein the logic die comprises a plurality of data receivers each of which is operable to receive a corresponding read data signal from each of the memory device dice, the receiver being operable to deserialize the received read data signal into a plurality of parallel read data bits and to apply the parallel read data bits to the memory access device.

20. The system of claim 16 wherein each of the memory device dice comprise a respective dynamic random access memory device dice.

21. A method of coupling read data from a plurality of stacked memory device dice each of which is connected to a logic die, the method comprising:

transmitting respective read strobe signals to each of the memory device dice from the logic die;

transmitting read data to the logic die from each of the memory device dice responsive to the memory device dice receiving the respective read strobe signal; and

adjusting the timing of transmitting respective read strobe signals from the logic die to each of the memory device dice to cause the read data to be received by the logic die at substantially the same time from each of the memory device dice.

22. The method of claim 21 wherein the act of adjusting the timing of transmitting respective read strobe signals to each of the memory device dice comprises conducting a training sequence to determine which timing allows the logic die to suitably capture the read data.

23. The method of claim 22 wherein the training sequence comprises:

repetitively transmitting the read strobe signal to each of the memory device dice, the read strobe signals being transmitted at different times over a range of times;

receiving read data at the logic die responsive to each of the read strobe signals transmitted to each of the memory device dice;

determining which timing of each of the read strobe signals caused the read data to be received by the logic die at substantially the same time from each of the memory device dice; and

subsequently using the determined timing to transmit each of the read strobe signals to the respective memory device dice.

24. The method of claim 21 wherein the act of adjusting the timing of transmitting each of the read strobe signals to a respective one of the memory device dice comprises:

providing a read strobe signal having a plurality of serial bits;

deserializing the read strobe signal to convert the plurality of serial bits into a plurality of corresponding parallel bits; and

transmitting each of the bits to the respective memory device dice as a serial bit stream starting at a time adjusted to cause the read data to be received by the logic die at substantially the same time that respective read data from the other memory device dice are received at the logic die.

25. The method of claim 21, further comprising the respective read data from each of the memory device dice responsive to a clock signal, the same clock signal being used to capture the respective read data from all of the memory device dice.




FIGURE 1

217




FIGURE 2

3/7




FIGURE 3




FIGURE 4

5/7




FIGURE 5

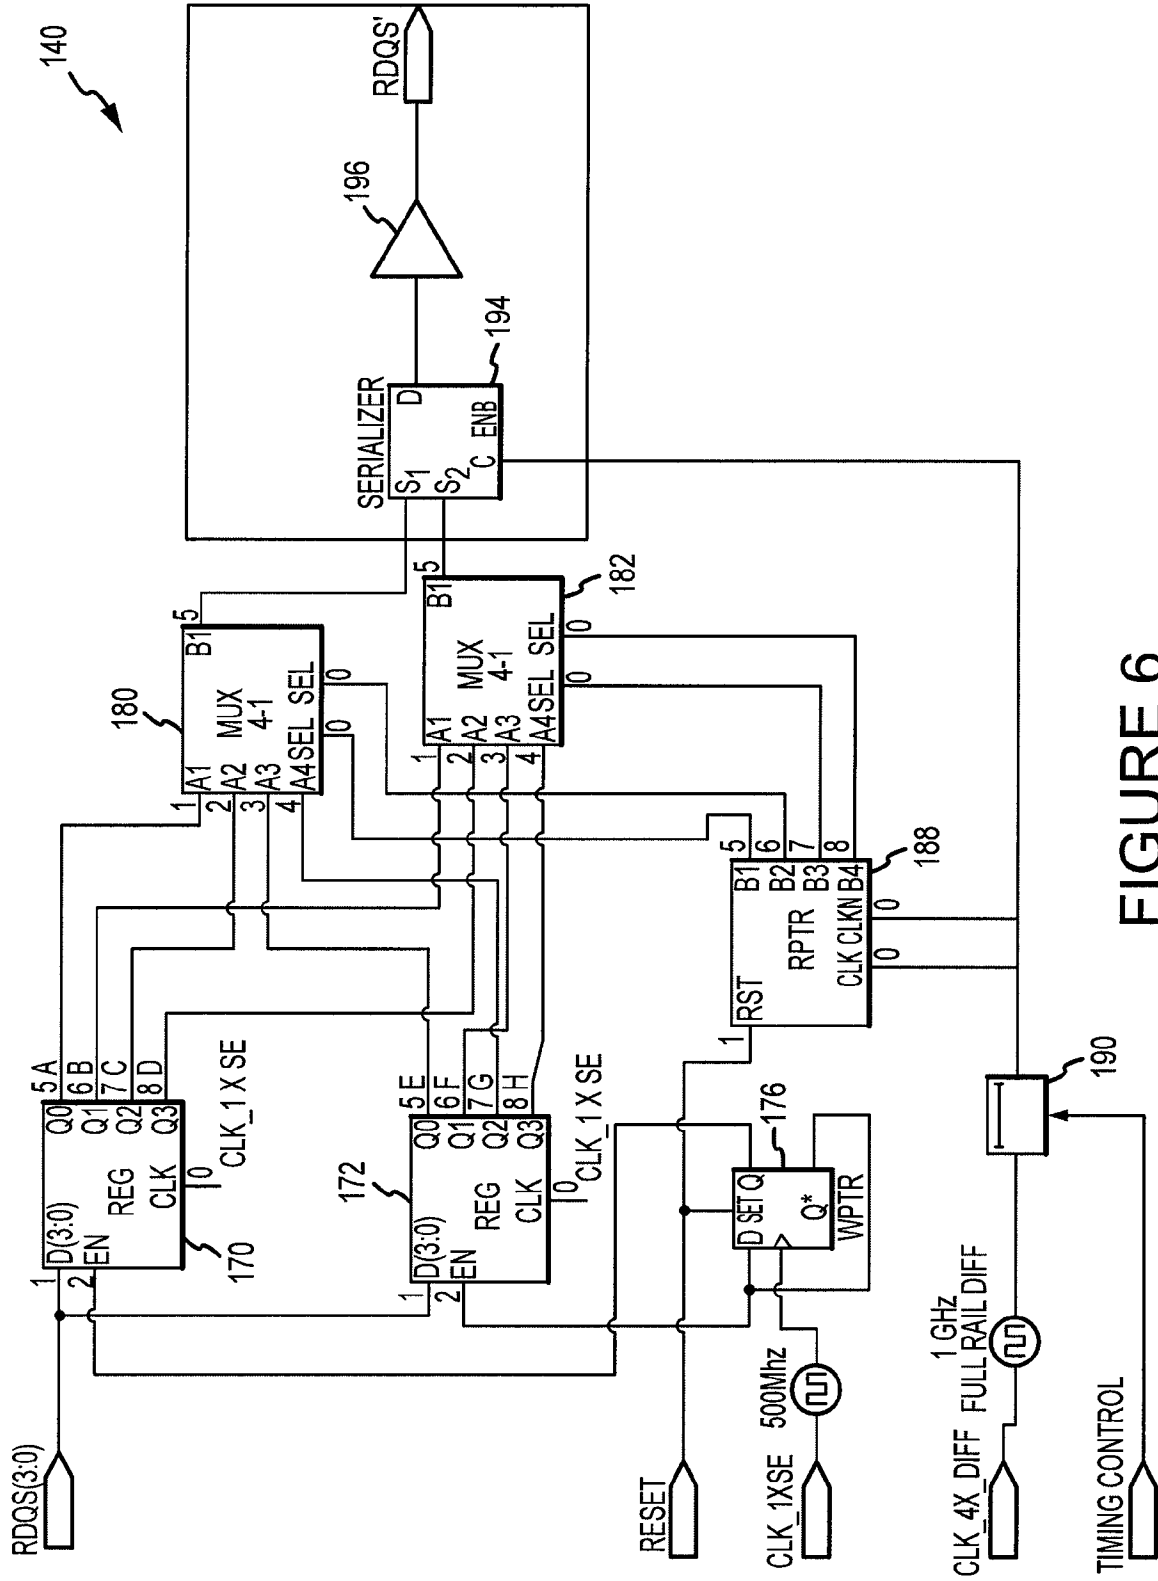


FIGURE 6

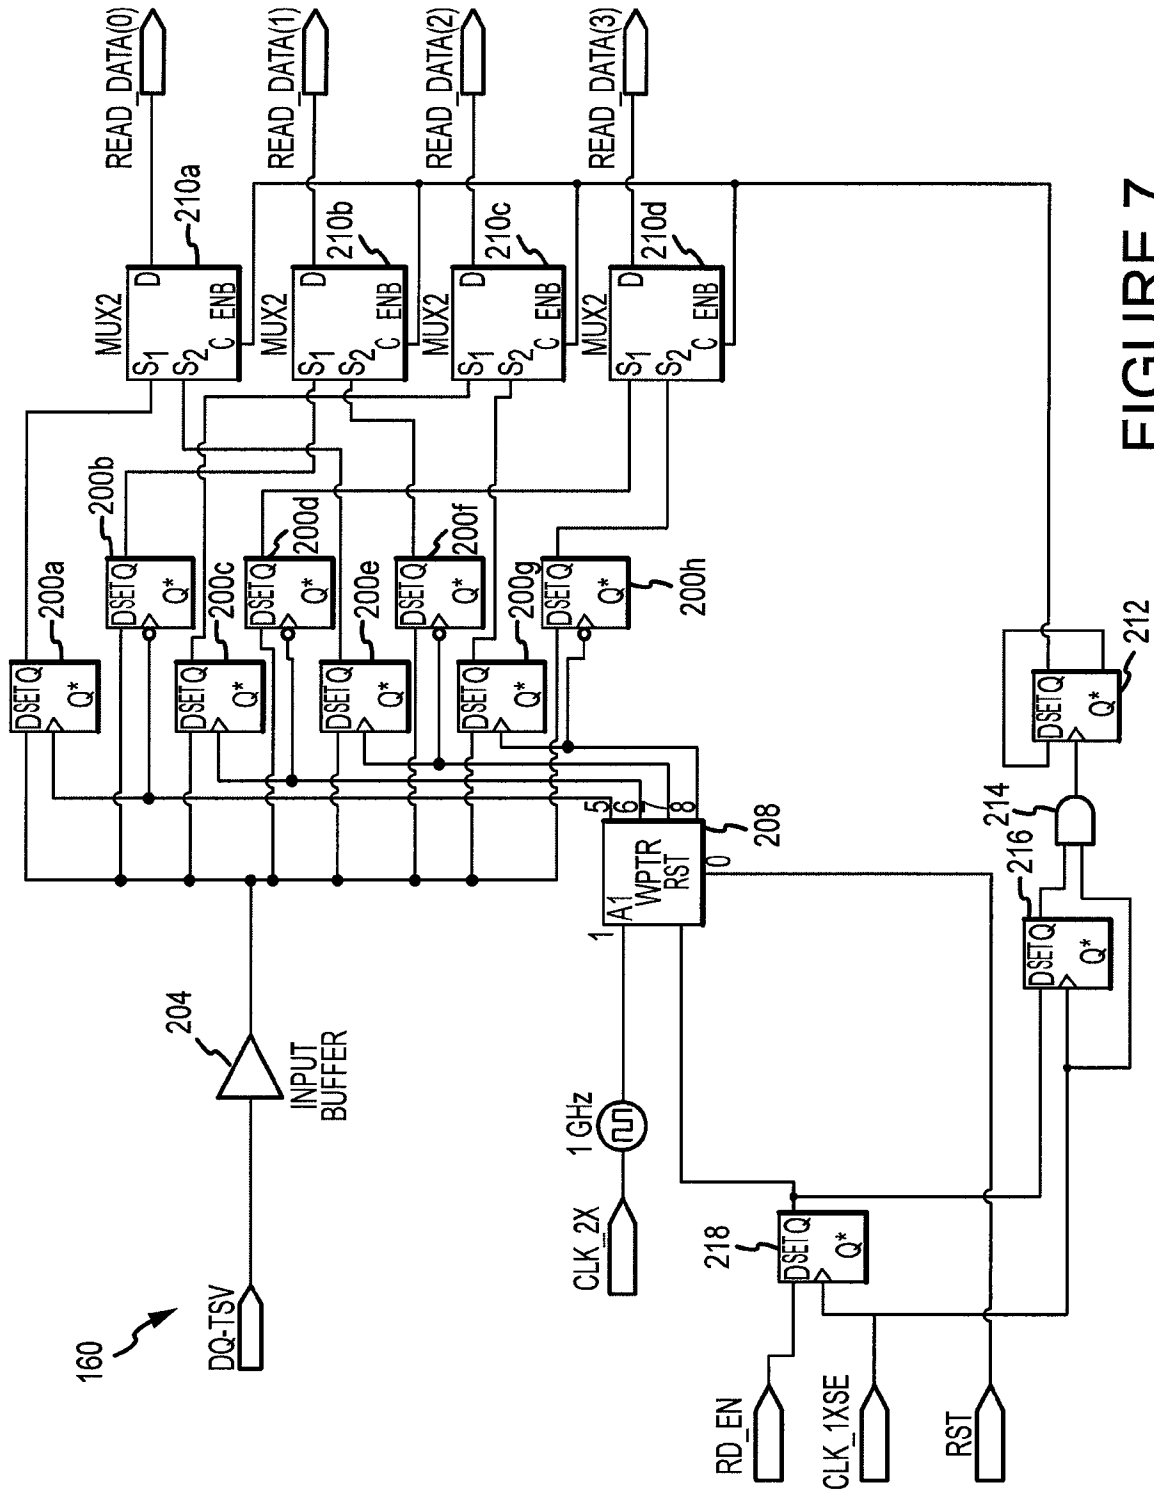


FIGURE 7