TIMOTHY T. SCOTT (SBN 126971/tscott@kslaw.com)
GEOFFREY M. EZGAR (SBN 184243/gezgar@kslaw.com)
LEO SPOONER III (SBN 241541/lspooner@kslaw.com)
KING & SPALDING LLP
333 Twin Dolphin Drive, Suite 400
Redwood Shores, CA 94065
Telephone: (650) 590-0700
Facsimile: (650) 590-1900

SCOTT T. WEINGAERTNER (*pro hac vice*/sweingaertner@kslaw.com)
ROBERT F. PERRY (rperry@kslaw.com)
ALLISON ALTERSOHN (*pro hac vice*/aaltersohn@kslaw.com)
SUSAN KIM (*pro hac vice*/skim@kslaw.com)
DANIEL MILLER (*pro hac vice*/dmiller@kslaw.com)
King & Spalding LLP
1185 Avenue of the Americas
New York, NY 10036-4003
Telephone: (212) 556-2100
Facsimile: (212) 556-2222

Attorneys for Plaintiff
GOOGLE INC.

UNITED STATES DISTRICT COURT FOR THE

NORTHERN DISTRICT OF CALIFORNIA

OAKLAND DIVISION

GOOGLE INC.

Plaintiff,

v.

NETLIST, INC.,

Defendant.

Civil Action No. C08 04144 SBA

[Related to Civil Action No. C09-05718 SBA]

EXHIBITS A, B, F, G, K, L AND M TO DECLARATION OF SCOTT T. WEINGAERTNER IN SUPPORT OF GOOGLE INC.'S OPPOSITION TO NETLIST, INC.'S MOTION FOR SUMMARY JUDGMENT OF INFRINGEMENT AND COUNTER-MOTION FOR SUMMARY JUDGMENT OF NON-INFRINGEMENT

Date: July 27, 2010
Time: 1:00pm
Place: Courtroom 3
Judge: Hon. Saundra Brown Armstrong

Exhibit A

US007289386B2

(12) United States Patent
Bhakta et al.

(10) Patent No.: US 7,289,386 B2
(45) Date of Patent: Oct. 30, 2007

(54) **MEMORY MODULE DECODER**

(75) Inventors: **Jayesh R. Bhakta**, Cerritos, CA (US); **Jeffrey C. Solomon**, Irvine, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 21 days.

(21) Appl. No.: **11/173,175**

(22) Filed: **Jul. 1, 2005**

(65) **Prior Publication Data**

US 2006/0062047 A1 Mar. 23, 2006

Related U.S. Application Data

(63) Continuation-in-part of application No. 11/075,395, filed on Mar. 7, 2005.

(60) Provisional application No. 60/588,244, filed on Jul. 15, 2004, provisional application No. 60/575,595, filed on May 28, 2004, provisional application No. 60/550,668, filed on Mar. 5, 2004.

(51) **Int. Cl.**
G11C 8/00 (2006.01)

(52) **U.S. Cl.** **365/230.06**; 365/51; 365/230.08

(58) **Field of Classification Search** 326/105; 365/230.06, 51, 52, 230.01, 230.03, 230.08; 711/5, 211
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,392,212	A *	7/1983	Miyasaka et al.	365/230.06
5,247,643	A	9/1993	Shottan	
5,426,753	A *	6/1995	Moon	711/5
5,703,826	A	12/1997	Hush et al.	
5,805,520	A	9/1998	Anglada et al.	
5,959,930	A *	9/1999	Sakurai	365/230.03
6,154,418	A	11/2000	Li	
6,453,381	B1	9/2002	Yuan et al.	
6,518,794	B2	2/2003	Coteus et al.	
6,681,301	B1	1/2004	Mehta et al.	
6,785,189	B2	8/2004	Jacobs et al.	
6,807,125	B2	10/2004	Coteus et al.	
6,813,196	B2	11/2004	Park et al.	

(Continued)

OTHER PUBLICATIONS

Barr, Michael. "Programmable Logic: What's it to Ya?," Embedded Systems Programming, Jun. 1999, pp. 75-84. Retrieved by Examiner from http://www.netrino.com/Articles/ProgrammableLogic/index.php.*

(Continued)

Primary Examiner—Son Dinh
Assistant Examiner—Alexander Sofocleous
(74) *Attorney, Agent, or Firm*—Knobbe Martens Olson & Bear LLP

(57) **ABSTRACT**

A memory module connectable to a computer system includes a printed circuit board, a plurality of memory devices coupled to the printed circuit board, and a logic element coupled to the printed circuit board. The plurality of memory devices has a first number of memory devices. The logic element receives a set of input control signals from the computer system. The set of input control signals corresponds to a second number of memory devices smaller than the first number of memory devices. The logic element generates a set of output control signals in response to the set of input control signals. The set of output control signals corresponds to the first number of memory devices.

13 Claims, 18 Drawing Sheets

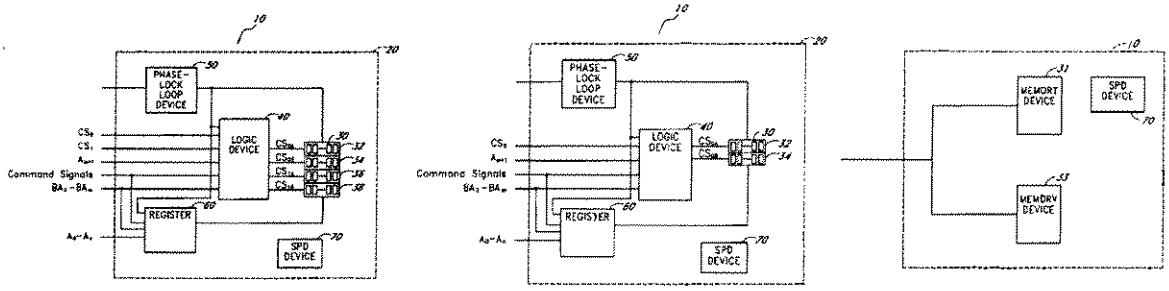

U.S. PATENT DOCUMENTS

6,944,694	B2	9/2005	Pax	
6,981,089	B2 *	12/2005	Dodd et al.	710/308
6,982,893	B2 *	1/2006	Jakobs	365/63
6,996,686	B2	2/2006	Doblar et al.	
7,046,538	B2 *	5/2006	Kinsley et al.	365/52
7,120,727	B2 *	10/2006	Lee et al.	711/5
7,200,021	B2 *	4/2007	Raghuram	365/51
2001/0052057	A1	12/2001	Lai et al.	
2002/0088633	A1	7/2002	Kong et al.	
2003/0063514	A1	4/2003	Faue	
2003/0191995	A1	10/2003	Abrosimov et al.	
2003/0210575	A1	11/2003	Seo et al.	
2004/0037158	A1	2/2004	Coteus et al.	
2005/0036378	A1	2/2005	Kawaguchi et al.	
2005/0281096	A1	12/2005	Bhakta et al.	
2006/0126369	A1 *	6/2006	Raghuram	365/51
2006/0129755	A1 *	6/2006	Raghuram	711/105

OTHER PUBLICATIONS

Cuppu, Vinodh, et al., "A Performance Comparison of Contemporary DRAM Architectures," *IEEE Proceedings of the 26th International Symposium on Computer Architecture*, May 2-4, 1999, Atlanta, Georgia, pp. 1-12.

JEDEC Standard No. 21-C, 4.20-2—168 Pin, PC133 SDRAM Registered Design Specification. Revision 1.4, Release 11a, Feb. 2002.

JEDEC Standard No. 21-C, 4.20-3—144 Pin, PC133 SDRAM Unbuffered SO-DIMM, Reference Design Specification, Revision 1.02, Release 11.

JEDEC Standard No. 21-C, DDR SDRAM PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification, Revision 1.3, Release 11b, Jan. 2002.

JEDEC Standard No. 21-C, 4.20.5—184 Pin, PC1600/2100 DDR SDRAM Unbuffered DIMM Design Specification, Revision 1.1, Release 11b.

JEDEC Standard No. 21-C, 4.20.6—200 Pin, PC2700/PC2100/PC1600 DDR SDRAM Unbuffered SO-DIMM Reference Design Specification, Revision 1.1, Release 11b, Apr. 26, 2002.

JEDEC standard JESD79D, "Double Data Rate (DDR) SDRAM Specification," published Feb. 2004.

Intel Corporation, PC SDRAM Registered DIMM Design Support Document, Revision 1.2, Oct. 1998.

Intel Corporation, 66/100 MHz PC SDRAM 64-Bit Non-ECC/Parity 144 Pin Unbuffered SO-DIMM Specification, Revision 1.0, Feb. 1999.

U.S. Appl. No. 11/173,175, filed Jul. 1, 2005, Bhakta et al.

* cited by examiner

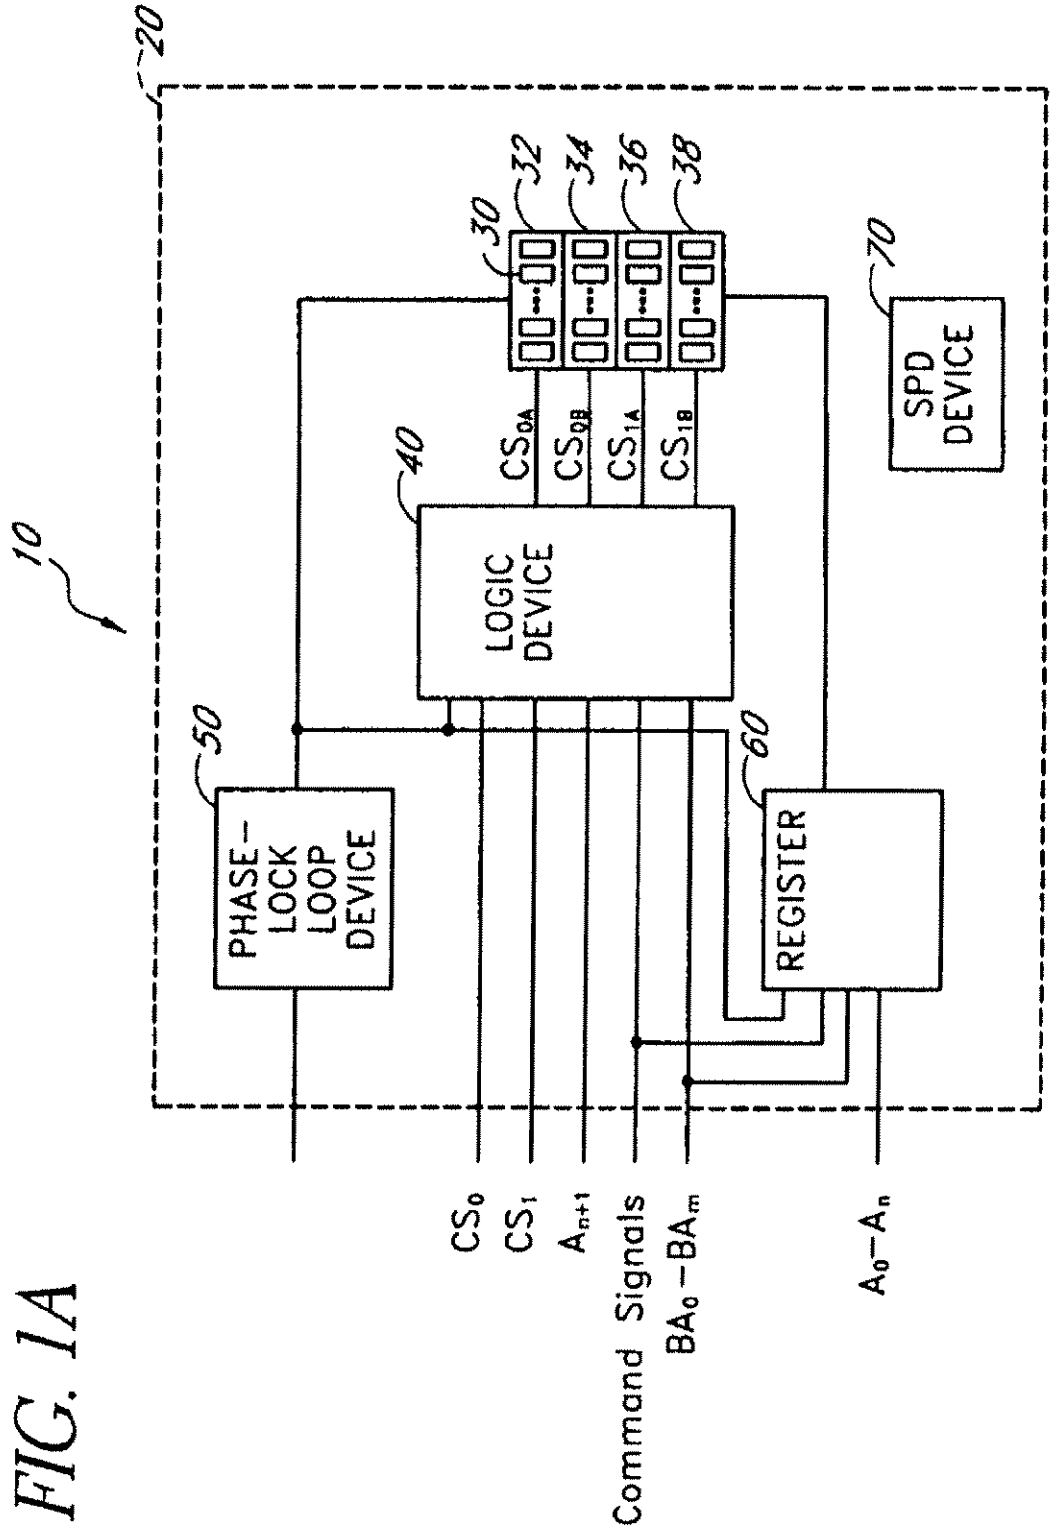
10
20
30
32
34
36
38
40
50
60
70
LOGIC DEVICE
PHASE-LOCK LOOP DEVICE
REGISTER
SPD DEVICE
CS0A
CS0B
CS1A
CS1B
CS0
CS1
An+1
Command Signals
BA0 – BAm
A0 – An

FIG. 1A

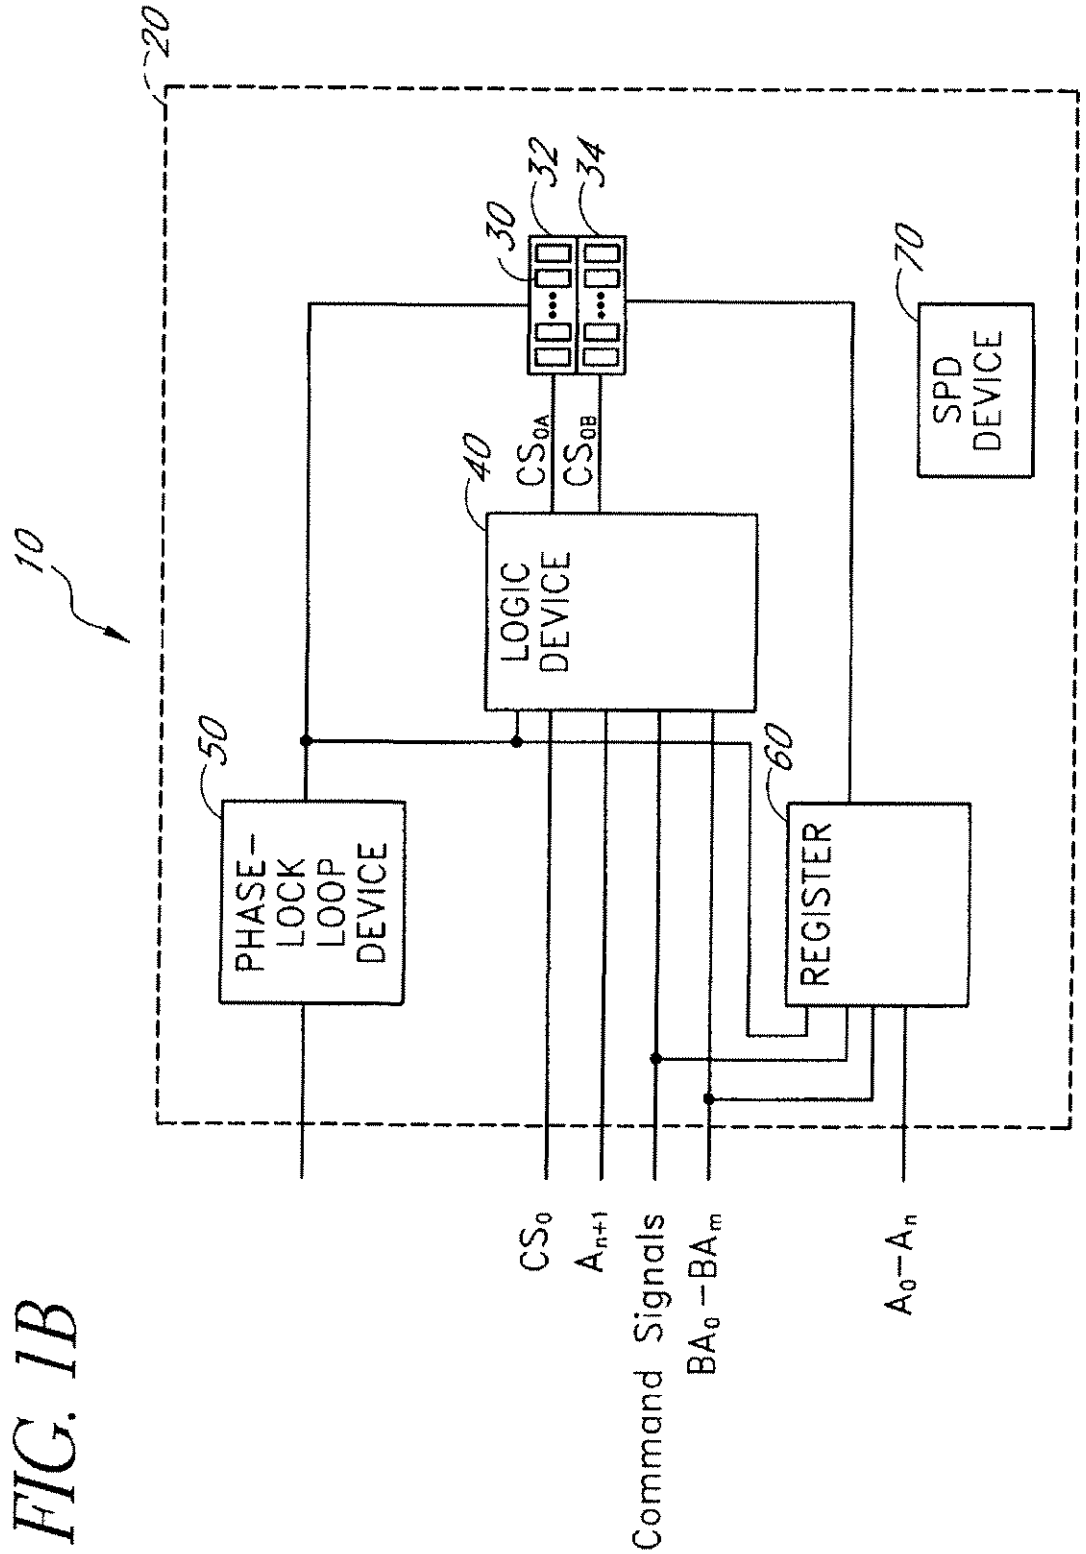
10
20
30
32
34
40
CS_{0A}
CS_{0B}
LOGIC DEVICE
50
PHASE-LOCK LOOP DEVICE
60
REGISTER
70
SPD DEVICE
CS_0
A_{n+1}
Command Signals
$BA_0 - BA_m$
$A_0 - A_n$

FIG. 1B

FIG. 1C

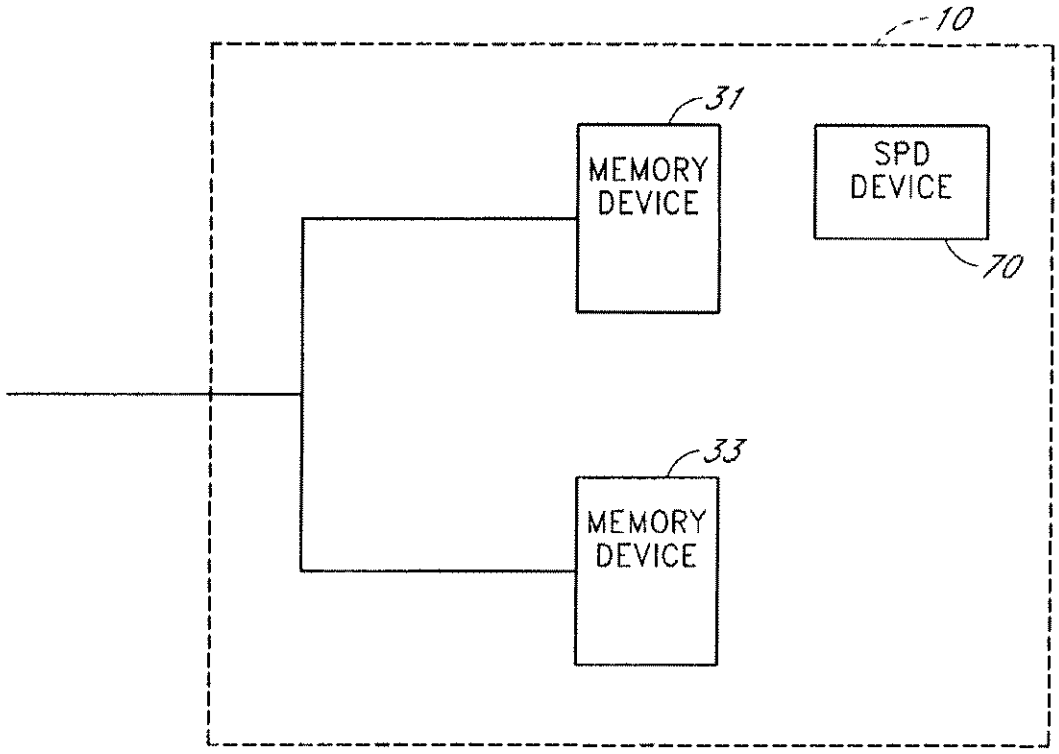

FIG. 2A

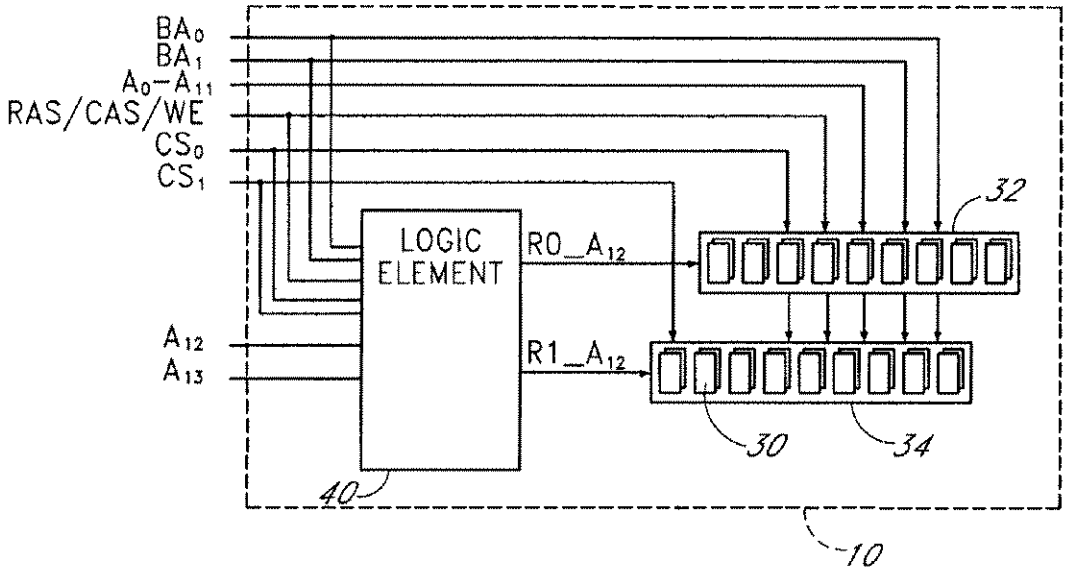
BA0
BA1
A0–A11
RAS/CAS/WE
CS0
CS1
A12
A13
LOGIC ELEMENT
R0_A12
R1_A12
32
30
34
40
10

FIG. 2B

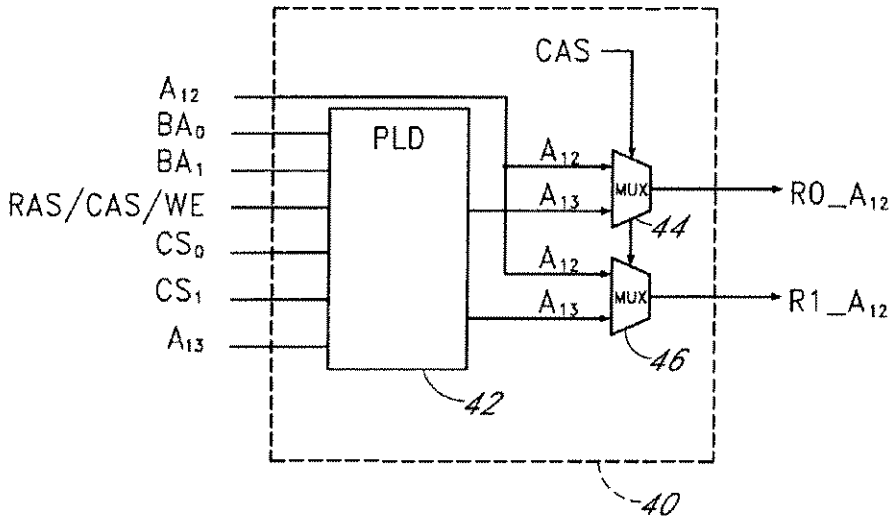
CAS
A12
BA0
BA1
RAS/CAS/WE
CS0
CS1
A13
PLD
A12
A13
MUX
44
R0_A12
A12
A13
MUX
46
R1_A12
42
40

FIG. 3A

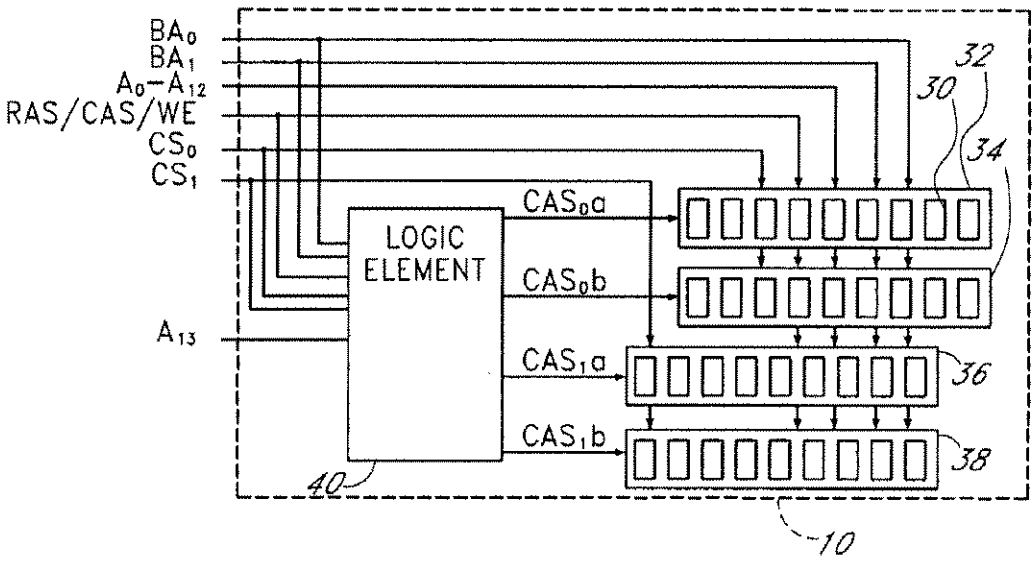
BA0
BA1
A0–A12
RAS/CAS/WE
CS0
CS1
A13
LOGIC ELEMENT
CAS0a
CAS0b
CAS1a
CAS1b
32
30
34
36
38
40
10

FIG. 3B

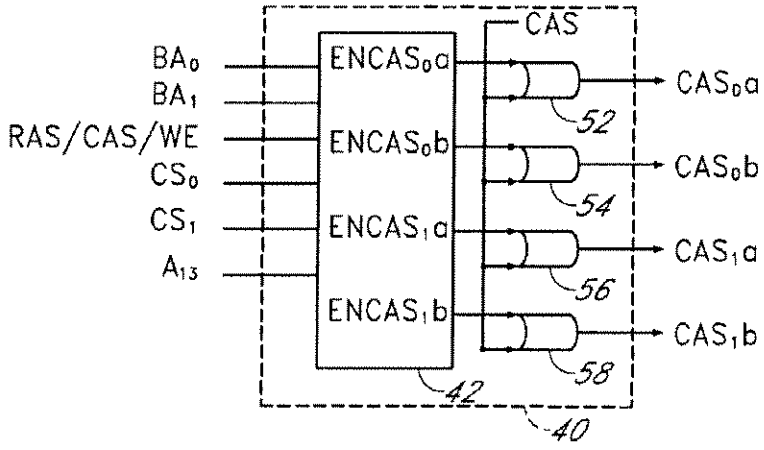
BA0
BA1
RAS/CAS/WE
CS0
CS1
A13
ENCAS0a
ENCAS0b
ENCAS1a
ENCAS1b
CAS
CAS0a
CAS0b
CAS1a
CAS1b
52
54
56
58
42
40

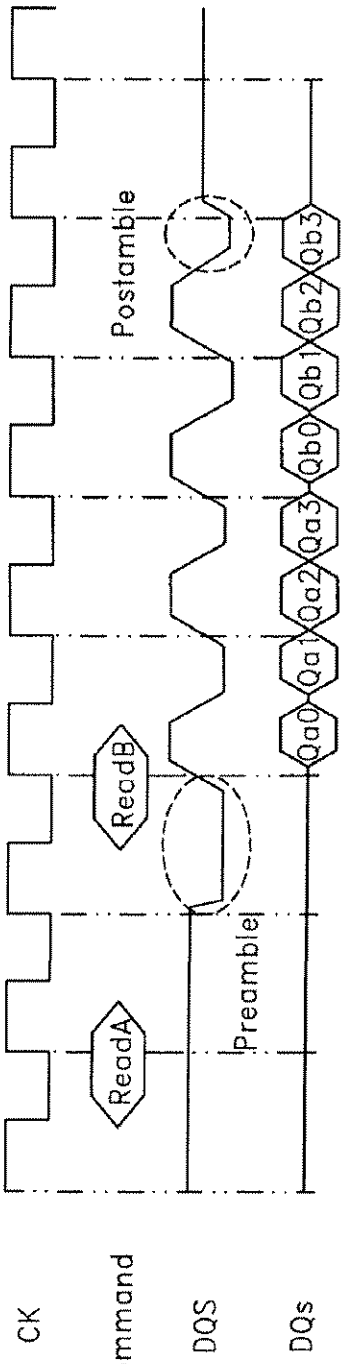
CK
Command
DQS
DQs
ReadA
ReadB
Preamble
Postamble
Qa0
Qa1
Qa2
Qa3
Qb0
Qb1
Qb2
Qb3

FIG. 4A

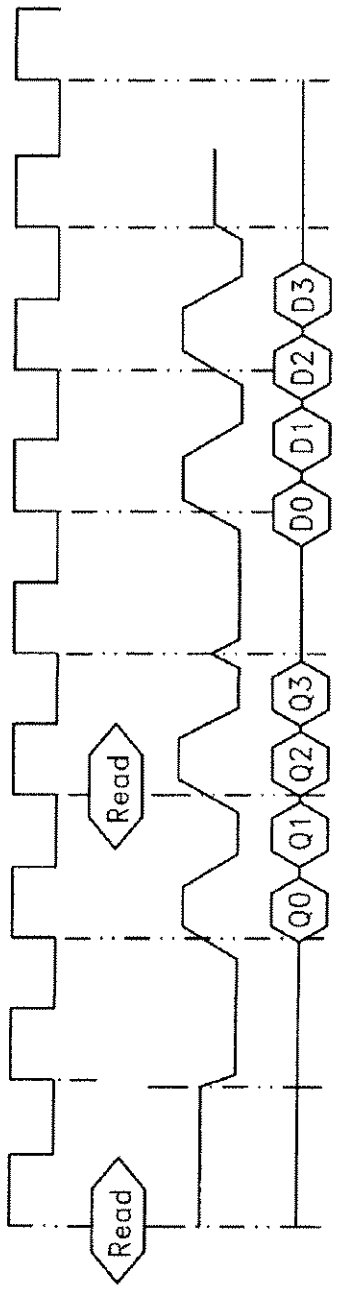
CK
Command
DQS
DQs
Read
Read
Q0
Q1
Q2
Q3
D0
D1
D2
D3

FIG. 4B

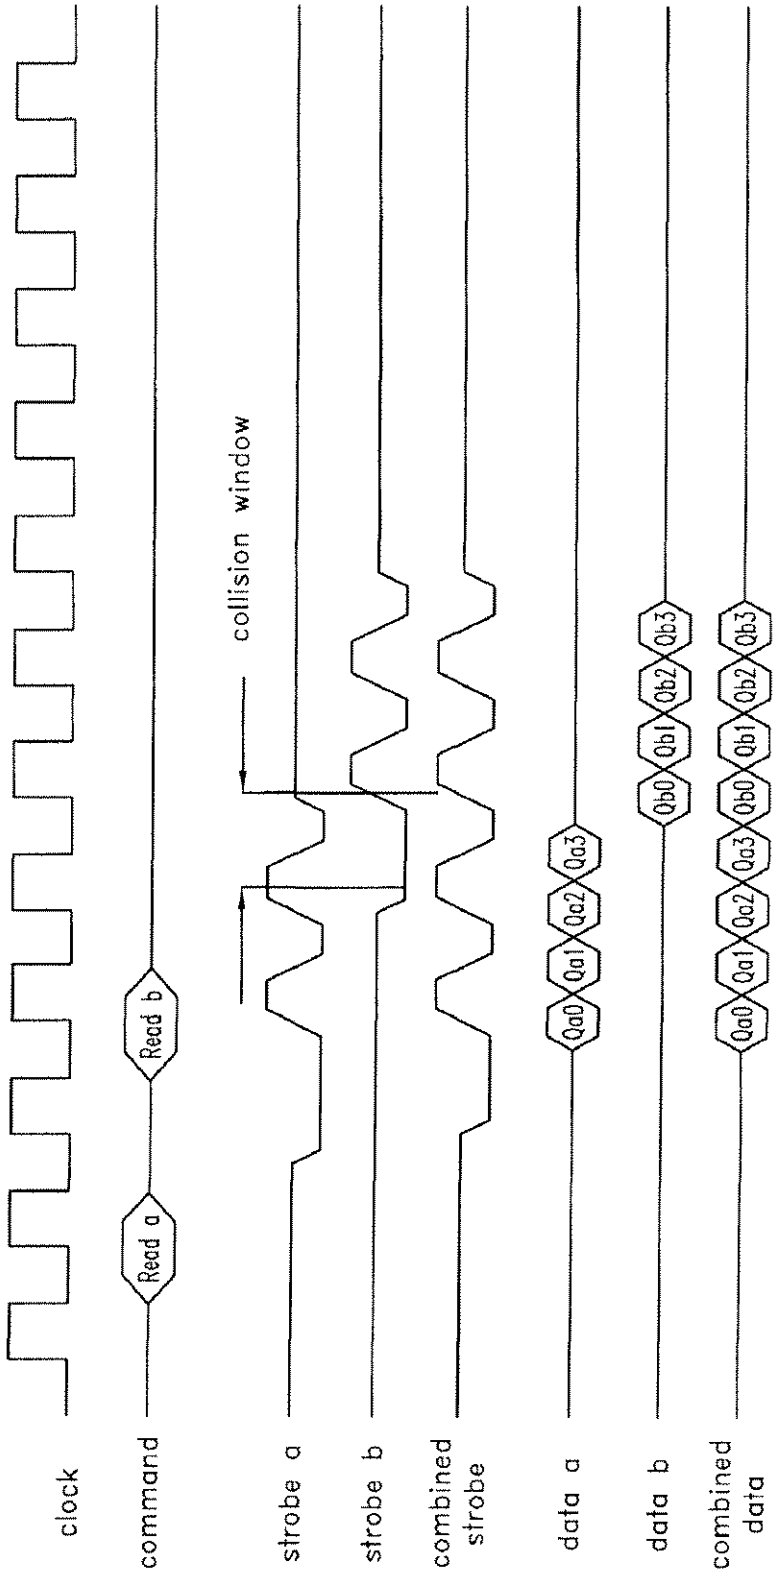
clock
command
Read a
Read b
collision window
strobe a
strobe b
combined strobe
data a
Qa0
Qa1
Qa2
Qa3
data b
Qb0
Qb1
Qb2
Qb3
combined data

FIG. 5

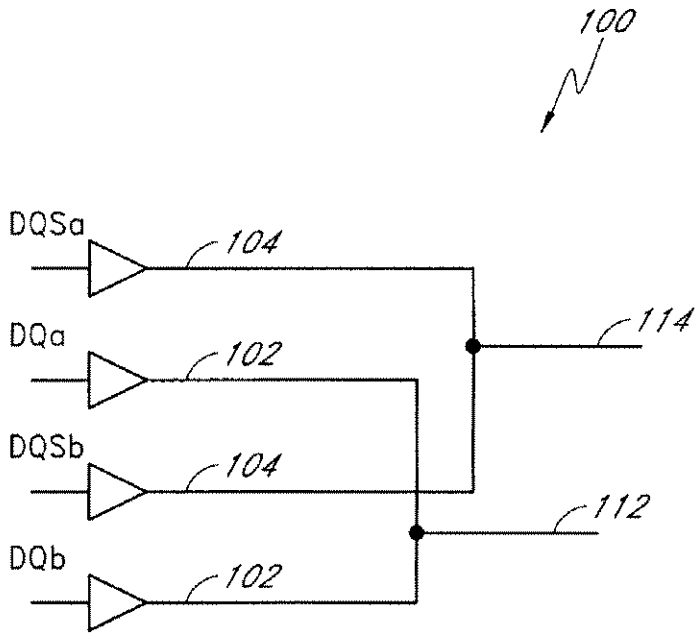
100
DQSa
104
DQa
102
DQSb
104
DQb
102
114
112

FIG. 6A

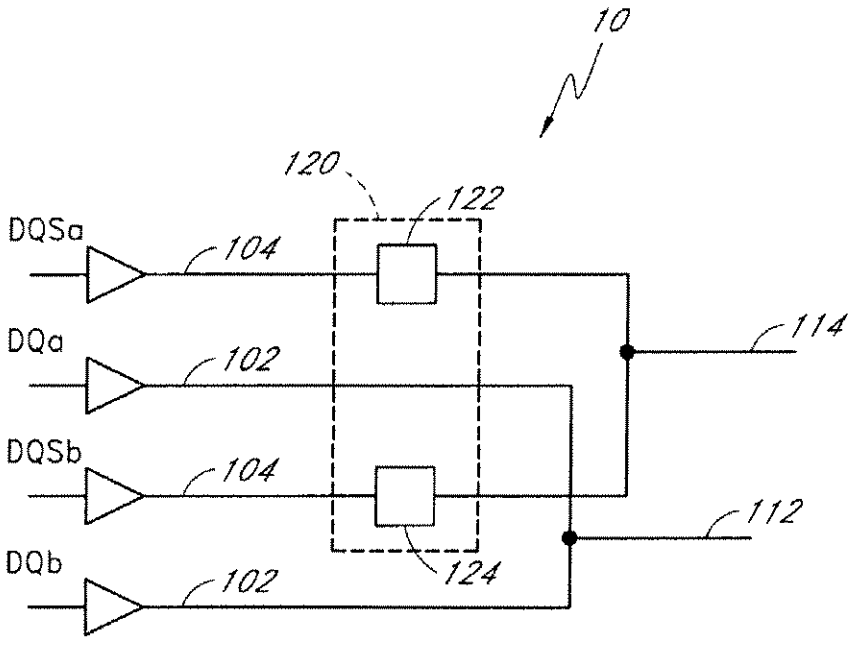
10
120
122
DQSa
104
DQa
102
DQSb
104
DQb
102
124
114
112

FIG. 6B

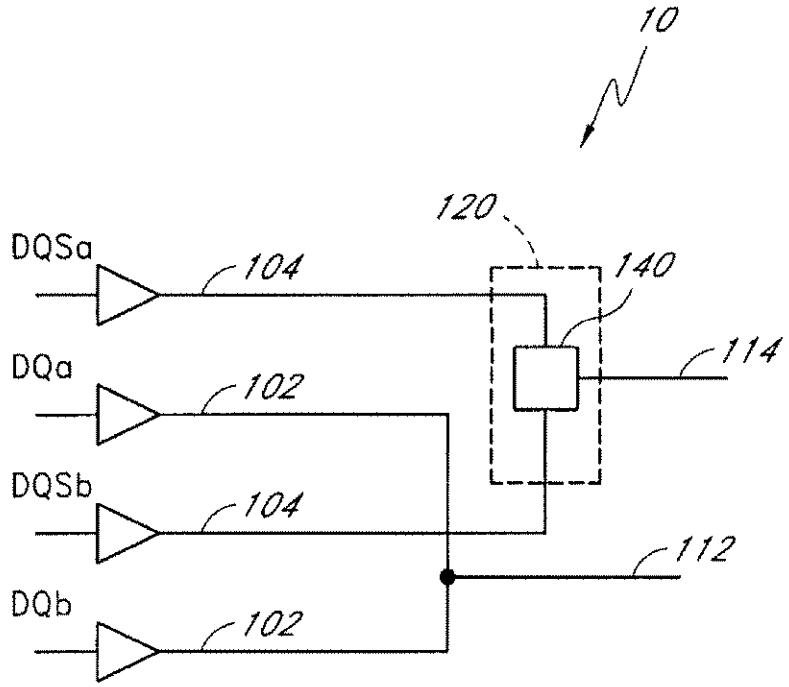
10
120
DQSa
104
140
114
DQa
102
DQSb
104
112
DQb
102

FIG. 6C

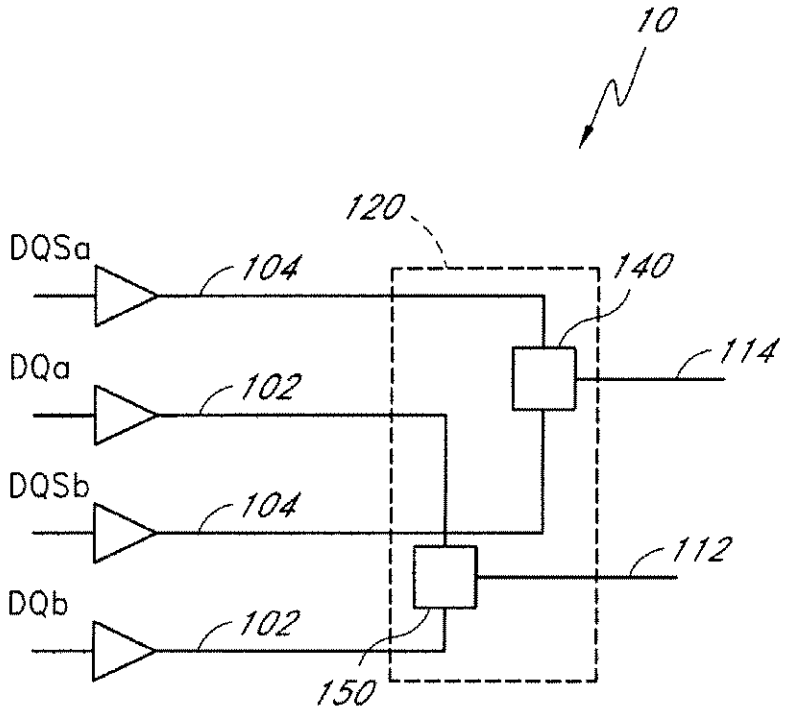
10
120
DQSa
104
140
114
DQa
102
DQSb
104
112
DQb
102
150

FIG. 6D

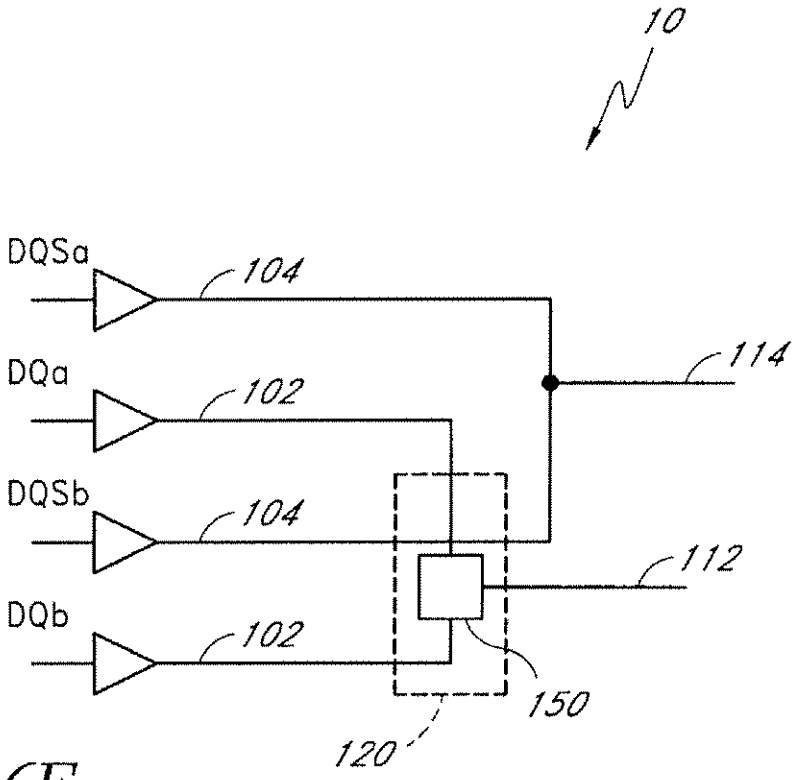

FIG. 6E

FIG. 7

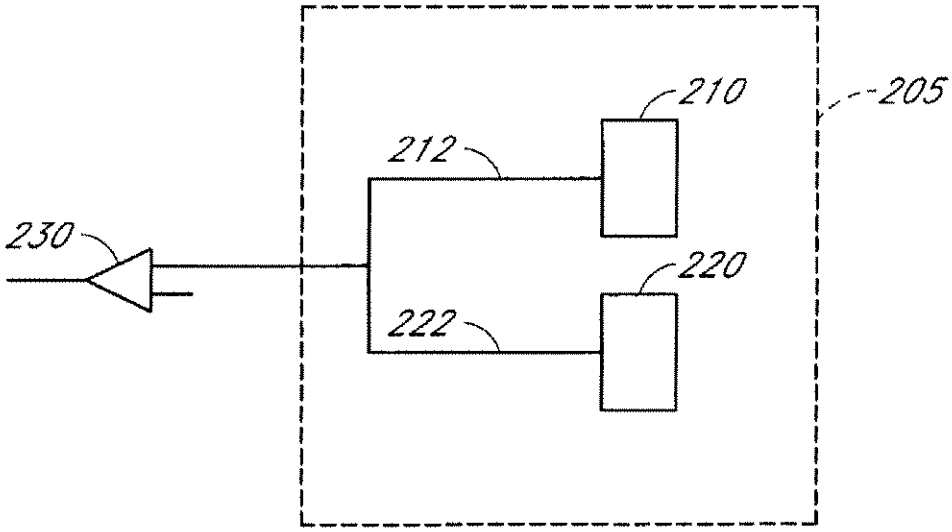

FIG. 8

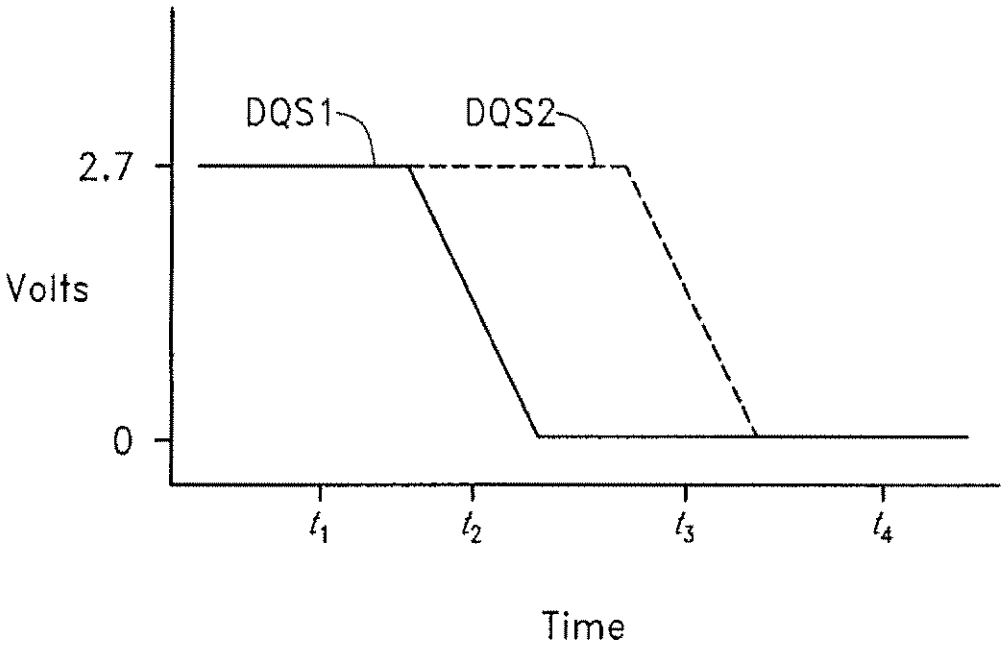

FIG. 9

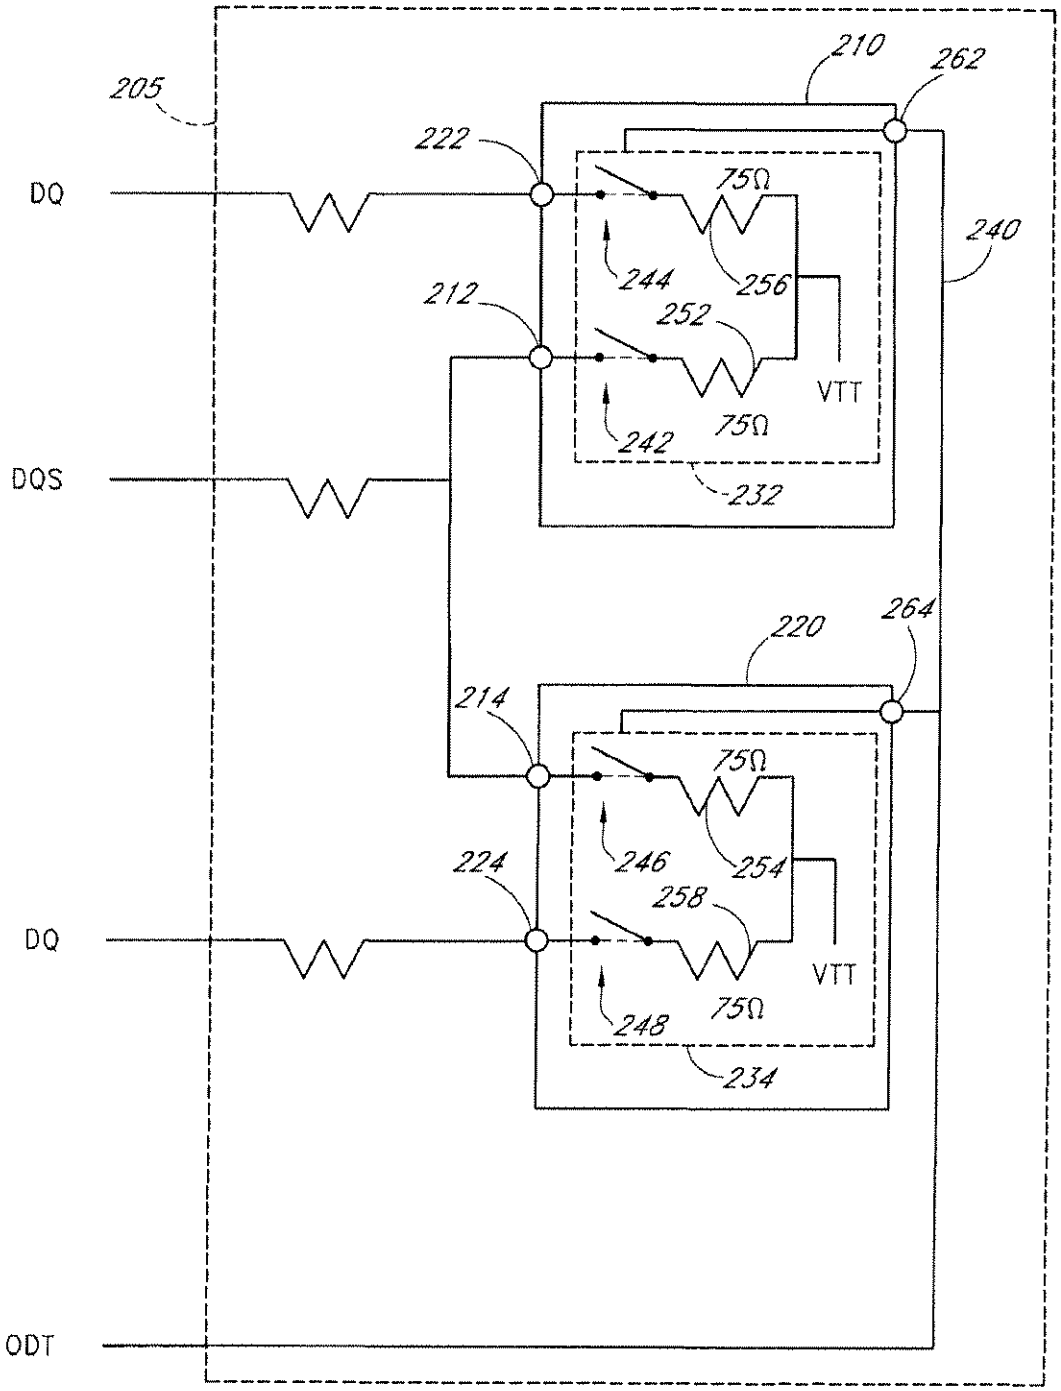
205
210
262
222
DQ
75Ω
240
244
256
212
252
VTT
242
75Ω
DQS
232
220
264
214
75Ω
246
254
224
258
DQ
VTT
248
75Ω
234
ODT

FIG. 10

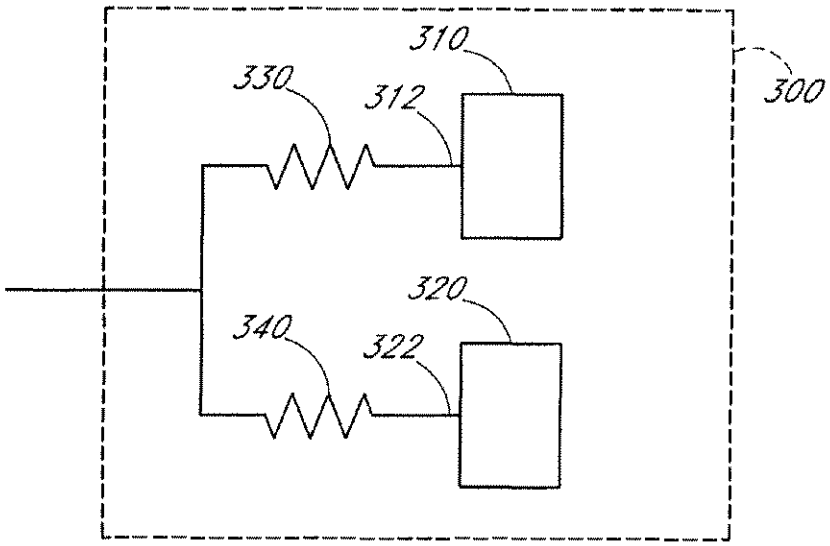

FIG. 11A

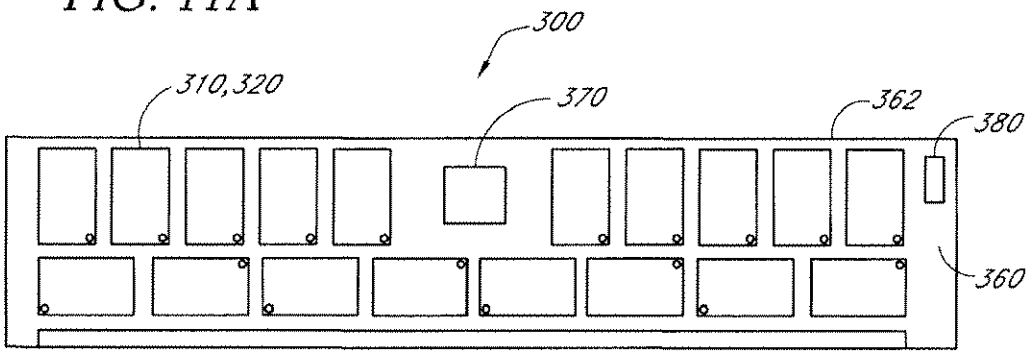

FIG. 11B

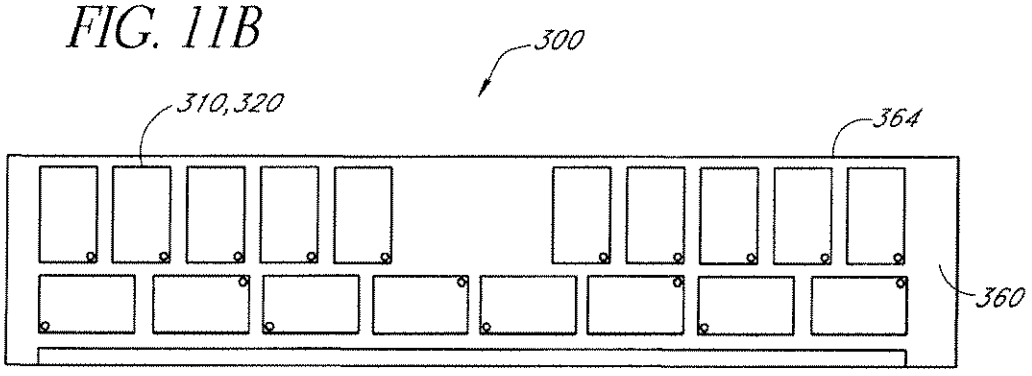

FIG. 12A

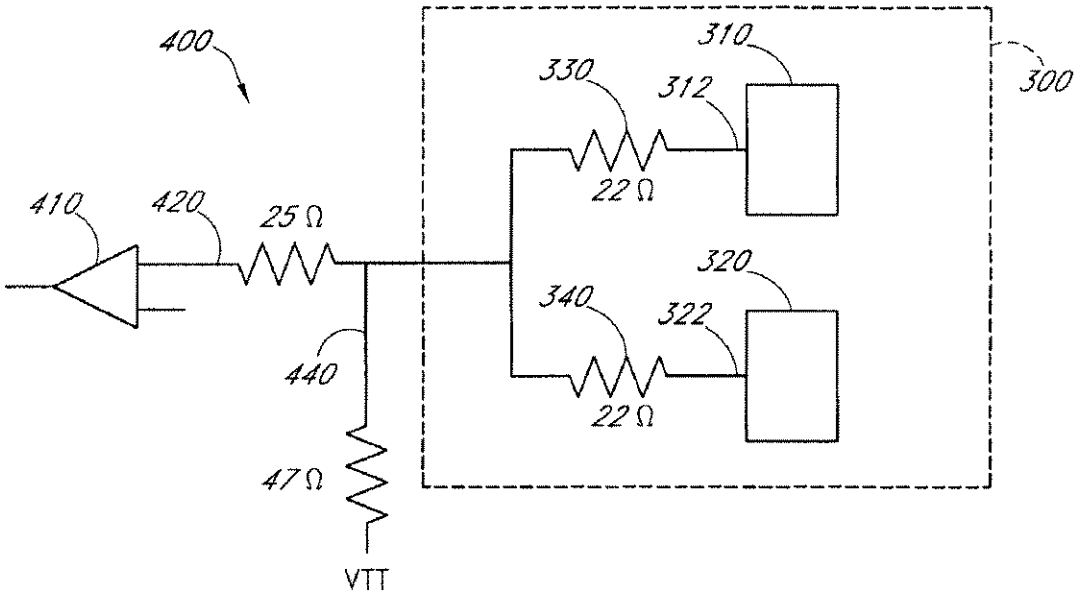

FIG. 12B

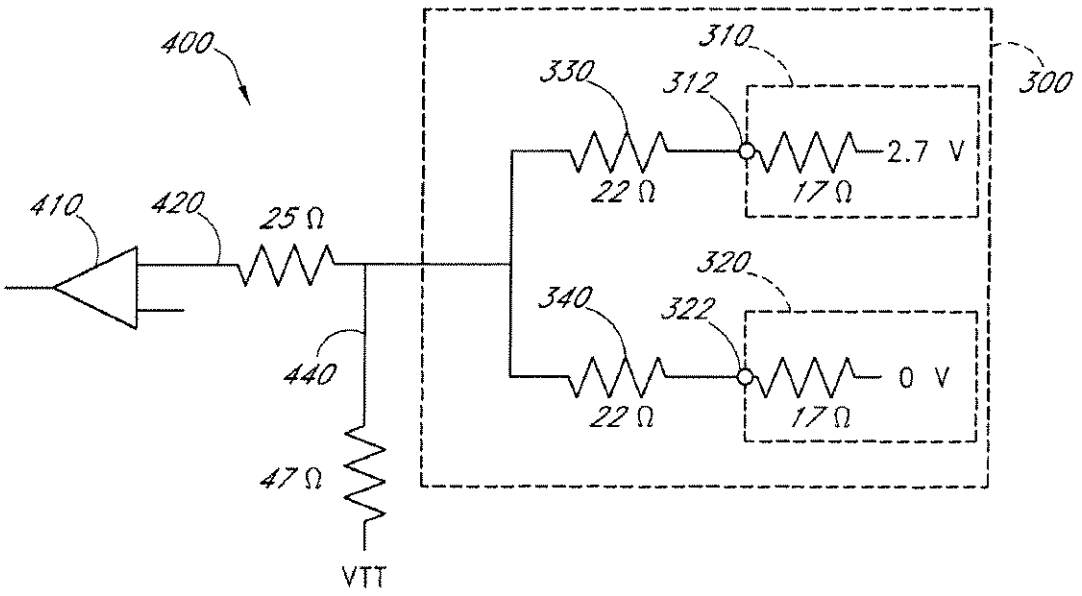

FIG. 13

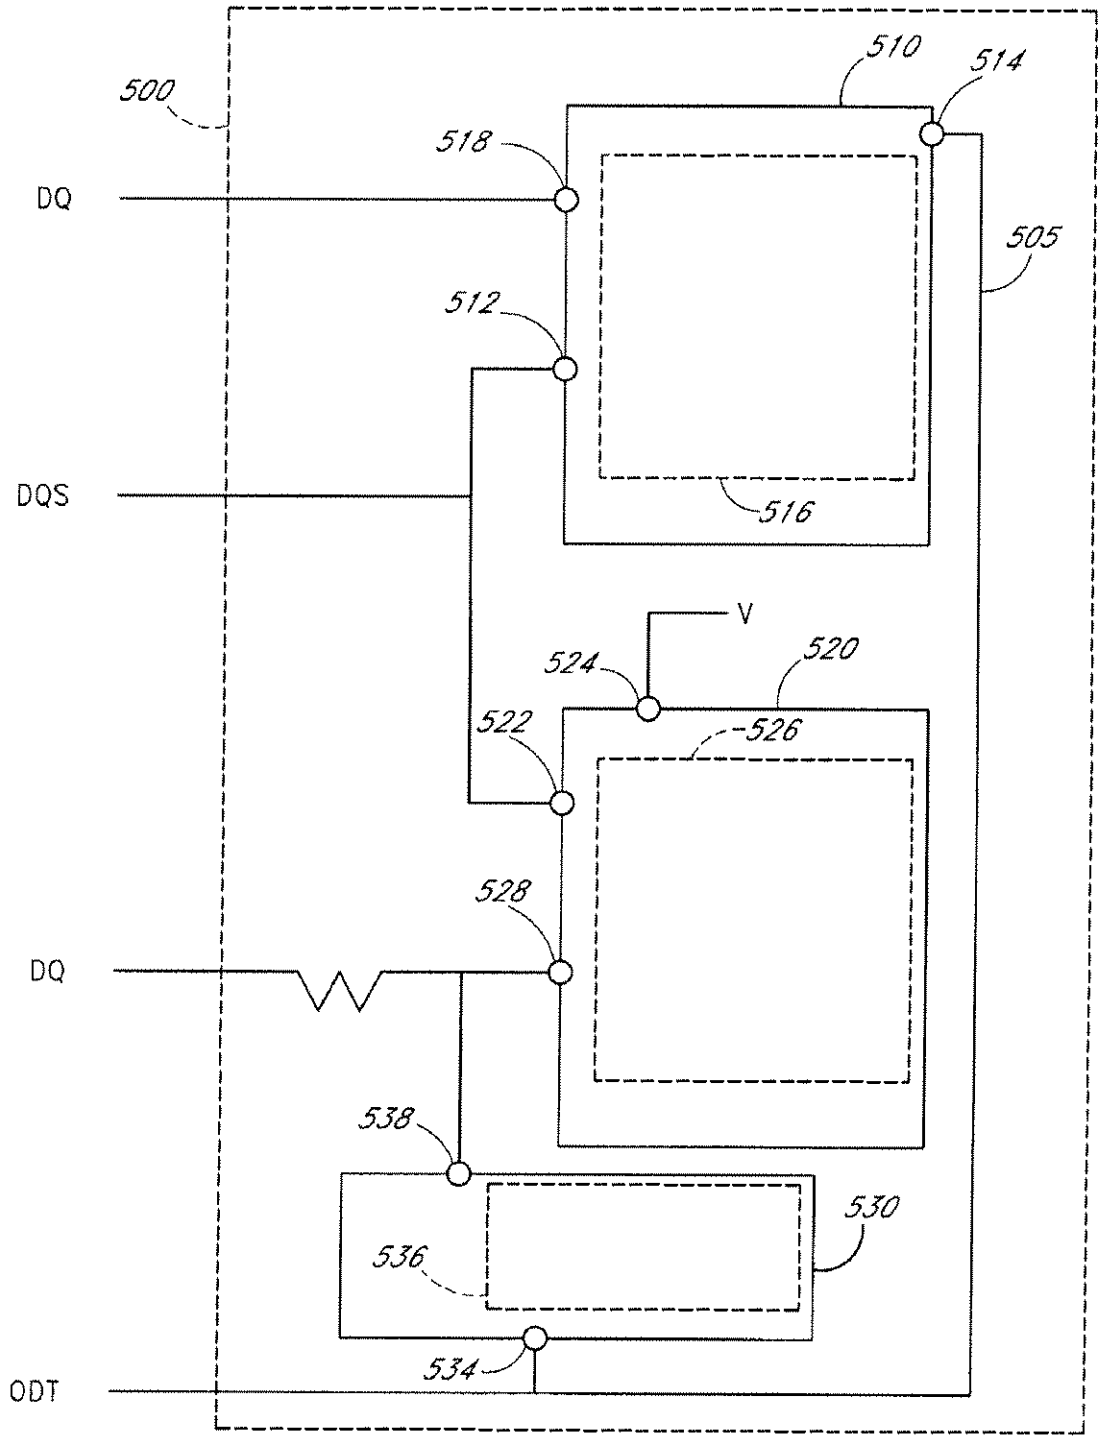
500
510
514
518
DQ
505
512
DQS
516
V
524
520
522
526
528
DQ
538
530
536
534
ODT

FIG. 14

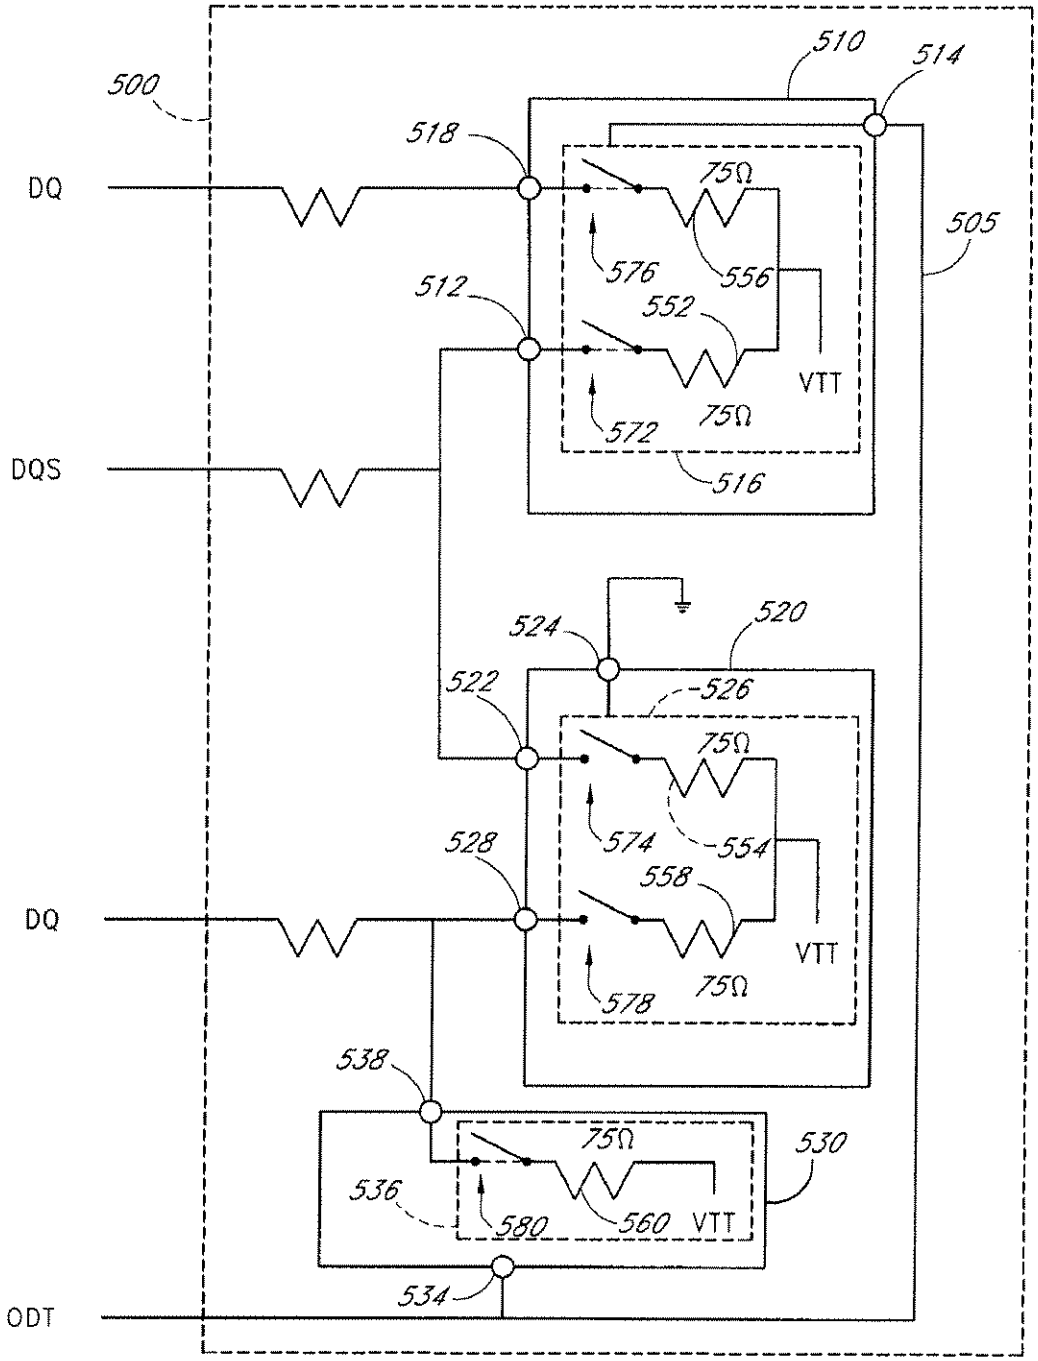
500
510
514
518
DQ
75Ω
505
576
556
512
552
VTT
572
75Ω
DQS
516
524
520
522
526
75Ω
574
554
528
558
DQ
VTT
578
75Ω
538
75Ω
530
536
580
560
VTT
534
ODT

MEMORY MODULE DECODER

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application is a continuation-in-part of U.S. patent application Ser. No. 11/075,395, filed Mar. 7, 2005, which claims the benefit of U.S. Provisional Application No. 60/550,668, filed Mar. 5, 2004 and U.S. Provisional Application No. 60/575,595, filed May 28, 2004. The present application also claims the benefit of U.S. Provisional Application No. 60/588,244, filed Jul. 15, 2004, which is incorporated in its entirety by reference herein.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to memory modules of a computer system, and more specifically to devices and methods for improving the performance, the memory capacity, or both, of memory modules.

2. Description of the Related Art

Certain types of memory modules comprise a plurality of dynamic random-access memory (DRAM) devices mounted on a printed circuit board (PCB). These memory modules are typically mounted in a memory slot or socket of a computer system (e.g., a server system or a personal computer) and are accessed by the processor of the computer system. Memory modules typically have a memory configuration with a unique combination of rows, columns, and banks which result in a total memory capacity for the memory module.

For example, a 512-Megabyte memory module (termed a "512-MB" memory module, which actually has 2^{29} or 536,870,912 bytes of capacity) will typically utilize eight 512-Megabit DRAM devices (each identified as a "512-Mb" DRAM device, each actually having 2^{29} or 536,870,912 bits of capacity). The memory cells (or memory locations) of each 512-Mb DRAM device can be arranged in four banks, with each bank having an array of 2^{24} (or 16,777,216) memory locations arranged as 2^{13} rows and 2^{11} columns, and with each memory location having a width of 8 bits. Such DRAM devices with 64M 8-bit-wide memory locations (actually with four banks of 2^{27} or 134,217,728 one-bit memory cells arranged to provide a total of 2^{26} or 67,108,864 memory locations with 8 bits each) are identified as having a "64 Mb×8" or "64M×8-bit" configuration, or as having a depth of 64M and a bit width of 8. Furthermore, certain commercially-available 512-MB memory modules are termed to have a "64M×8-byte" configuration or a "64M×64-bit" configuration with a depth of 64M and a width of 8 bytes or 64 bits.

Similarly, a 1-Gigabyte memory module (termed a "1-GB" memory module, which actually has 2^{30} or 1,073,741,824 bytes of capacity) can utilize eight 1-Gigabit DRAM devices (each identified as a "1-Gb" DRAM device, each actually having 2^{30} or 1,073,741,824 bits of capacity). The memory locations of each 1-Gb DRAM device can be arranged in four banks, with each bank having an array of memory locations with 2^{14} rows and 2^{11} columns, and with each memory location having a width of 8 bits. Such DRAM devices with 128M 8-bit-wide memory locations (actually with a total of 2^{27} or 134,217,728 memory locations with 8 bits each) are identified as having a "128 Mb×8" or "128M×8-bit" configuration, or as having a depth of 128M and a bit width of 8. Furthermore, certain commercially-available 1-GB memory modules are identified as having a "128M×8-byte" configuration or a "128M×64-bit" configuration with a depth of 128M and a width of 8 bytes or 64 bits.

The commercially-available 512-MB (64M×8-byte) memory modules and the 1-GB (128M×8-byte) memory modules described above are typically used in computer systems (e.g., personal computers) which perform graphics applications since such "×8" configurations are compatible with data mask capabilities often used in such graphics applications. Conversely, memory modules with "×4" configurations are typically used in computer systems such as servers which are not as graphics-intensive. Examples of such commercially available "×4" memory modules include, but are not limited to, 512-MB (128M×4-byte) memory modules comprising eight 512-Mb (128 Mb×4) memory devices.

The DRAM devices of a memory module are generally arranged as ranks or rows of memory, each rank of memory generally having a bit width. For example, a memory module in which each rank of the memory module is 64 bits wide is described as having an "×64" organization. Similarly, a memory module having 72-bit-wide ranks is described as having an "×72" organization.

The memory capacity of a memory module increases with the number of memory devices. The number of memory devices of a memory module can be increased by increasing the number of memory devices per rank or by increasing the number of ranks. For example, a memory module with four ranks has double the memory capacity of a memory module with two ranks and four times the memory capacity of a memory module with one rank. Rather than referring to the memory capacity of the memory module, in certain circumstances, the memory density of the memory module is referred to instead.

During operation, the ranks of a memory module are selected or activated by control signals that are received from the processor. Examples of such control signals include, but are not limited to, rank-select signals, also called chip-select signals. Most computer and server systems support one-rank and two-rank memory modules. By only supporting one-rank and two-rank memory modules, the memory density that can be incorporated in each memory slot is limited.

SUMMARY OF THE INVENTION

In certain embodiments, a memory module is connectable to a computer system. The memory module comprises a printed circuit board, a plurality of memory devices coupled to the printed circuit board, and a logic element coupled to the printed circuit board. The plurality of memory devices has a first number of memory devices. The logic element receives a set of input control signals from the computer system. The set of input control signals corresponds to a second number of memory devices smaller than the first number of memory devices. The logic element generates a set of output control signals in response to the set of input control signals. The set of output control signals corresponds to the first number of memory devices.

In certain embodiments, a method utilizes a memory module in a computer system. The method comprises coupling the memory module to the computer system. The memory module comprises a plurality of memory devices arranged in a first number of ranks. The method further comprises inputting a first set of control signals to the memory module. The first set of control signals corresponds to a second number of ranks smaller than the first number of ranks. The method further comprises generating a second set

of control signals in response to the first set of control signals. The second set of control signals corresponds to the first number of ranks.

In certain embodiments, a memory module is connectable to a computer system. The memory module comprises a plurality of memory devices arranged in a first number of ranks. The memory module comprises means for coupling the memory module to the computer system. The memory module further comprises means for inputting a first set of control signals to the memory module. The first set of control signals corresponds to a second number of ranks smaller than the first number of ranks. The memory module further comprises means for generating a second set of control signals in response to the first set of control signals. The second set of control signals corresponds to the first number of ranks.

In certain embodiments, a memory module is connectable to a computer system. The memory module comprises a first memory device having a first data signal line and a first data strobe signal line. The memory module further comprises a second memory device having a second data signal line and a second data strobe signal line. The memory module further comprises a common data signal line connectable to the computer system. The memory module further comprises an isolation device electrically coupled to the first data signal line, to the second data signal line, and to the common data signal line. The isolation device selectively alternates between electrically coupling the first data signal line to the common data signal line and electrically coupling the second data signal line to the common data signal line.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. **1**A schematically illustrates an exemplary memory module with four ranks of memory devices compatible with certain embodiments described herein.

FIG. **1**B schematically illustrates an exemplary memory module with two ranks of memory devices compatible with certain embodiments described herein.

FIG. **1**C schematically illustrates another exemplary memory module in accordance with certain embodiments described herein.

FIG. **2**A schematically illustrates an exemplary memory module which doubles the rank density in accordance with certain embodiments described herein.

FIG. **2**B schematically illustrates an exemplary logic element compatible with embodiments described herein.

FIG. **3**A schematically illustrates an exemplary memory module which doubles number of ranks in accordance with certain embodiments described herein.

FIG. **3**B schematically illustrates an exemplary logic element compatible with embodiments described herein.

FIG. **4**A shows an exemplary timing diagram of a gapless read burst for a back-to-back adjacent read condition from one memory device.

FIG. **4**B shows an exemplary timing diagram with an extra clock cycle between successive read commands issued to different memory devices for successive read accesses from different memory devices.

FIG. **5** shows an exemplary timing diagram in which the last data strobe of memory device "a" collides with the pre-amble time interval of the data strobe of memory device "b."

FIG. **6**A schematically illustrates a circuit diagram of a conventional memory module showing the interconnections between the DQ data signal lines of two memory devices and their DQS data strobe signal lines.

FIG. **6**B schematically illustrates a circuit diagram of an exemplary memory module comprising an isolation device in accordance with certain embodiments described herein.

FIG. **6**C schematically illustrates an isolation device comprising a logic element which multiplexes the DQS data strobe signal lines from one another.

FIG. **6**D schematically illustrates an isolation device which multiplexes the DQS data strobe signal lines from one another and which multiplexes the DQ data signal lines from one another.

FIG. **6**E schematically illustrates an isolation device which comprises the logic element on the DQ data signal lines but not a corresponding logic element on the DQS data strobe signal lines.

FIG. **7** schematically illustrates an exemplary memory module in which a data strobe (DQS) pin of a first memory device is electrically connected to a DQS pin of a second memory device while both DQS pins are active.

FIG. **8** is an exemplary timing diagram of the voltages applied to the two DQS pins due to non-simultaneous switching.

FIG. **9** schematically illustrates another exemplary memory module in which a DQS pin of a first memory device is connected to a DQS pin of a second memory device.

FIG. **10** schematically illustrates an exemplary memory module in accordance with certain embodiments described herein.

FIGS. **11**A and **11**B schematically illustrate a first side and a second side, respectively, of a memory module with eighteen 64M×4 bit, DDR-I SDRAM FBGA memory devices on each side of a 184-pin glass-epoxy printed circuit board.

FIGS. **12**A and **12**B schematically illustrate an exemplary embodiment of a memory module in which a first resistor and a second resistor are used to reduce the current flow between the first DQS pin and the second DQS pin.

FIG. **13** schematically illustrates another exemplary memory module compatible with certain embodiments described herein.

FIG. **14** schematically illustrates a particular embodiment of the memory module schematically illustrated by FIG. **13**.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

Most high-density memory modules are currently built with 512-Megabit ("512-Mb") memory devices wherein each memory device has a 64M×8-bit configuration. For example, a 1-Gigabyte ("1-GB") memory module with error checking capabilities can be fabricated using eighteen such 512-Mb memory devices. Alternatively, it can be economically advantageous to fabricate a 1-GB memory module using lower-density memory devices and doubling the number of memory devices used to produce the desired word width. For example, by fabricating a 1-GB memory module using thirty-six 256-Mb memory devices with 64M×4-bit configuration, the cost of the resulting 1-GB memory module can be reduced since the unit cost of each 256-Mb memory device is typically lower than one-half the unit cost of each 512-Mb memory device. The cost savings can be significant, even though twice as many 256-Mb memory devices are used in place of the 512-Mb memory devices.

Market pricing factors for DRAM devices are such that higher-density DRAM devices (e.g., 1-Gb DRAM devices) are much more than twice the price of lower-density DRAM devices (e.g., 512-Mb DRAM devices). In other words, the

price per bit ratio of the higher-density DRAM devices is greater than that of the lower-density DRAM devices. This pricing difference often lasts for months or even years after the introduction of the higher-density DRAM devices, until volume production factors reduce the costs of the newer higher-density DRAM devices. Thus, when the cost of a higher-density DRAM device is more than the cost of two lower-density DRAM devices, there is an economic incentive for utilizing pairs of the lower-density DRAM devices to replace individual higher-density DRAM devices.

FIG. **1**A schematically illustrates an exemplary memory module **10** compatible with certain embodiments described herein. The memory module **10** is connectable to a computer system (not shown). The memory module **10** comprises a printed circuit board **20** and a plurality of memory devices **30** coupled to the printed circuit board **20**. The plurality of memory devices **30** has a first number of memory devices. The memory module **10** further comprises a logic element **40** coupled to the printed circuit board **20**. The logic element **40** receives a set of input control signals from the computer system. The set of input control signals correspond to a second number of memory devices smaller than the first number of memory devices. The logic element **40** generates a set of output control signals in response to the set of input control signals. The set of output control signals corresponds to the first number of memory devices.

In certain embodiments, as schematically illustrated in FIG. **1**A, the memory module **10** further comprises a phase-lock loop device **50** coupled to the printed circuit board **20** and a register **60** coupled to the printed circuit board **20**. In certain embodiments, the phase-lock loop device **50** and the register **60** are each mounted on the printed circuit board **20**. In response to signals received from the computer system, the phase-lock loop device **50** transmits clock signals to the plurality of memory devices **30**, the logic element **40**, and the register **60**. The register **60** receives and buffers a plurality of control signals, including address signals (e.g., bank address signals, row address signals, column address signals, gated column address strobe signals, chip-select signals), and transmits corresponding signals to the appropriate memory devices **30**. In certain embodiments, the register **60** comprises a plurality of register devices. While the phase-lock loop device **50**, the register **60**, and the logic element **40** are described herein in certain embodiments as being separate components, in certain other embodiments, two or more of the phase-lock loop device **50**, the register **60**, and the logic element **40** are portions of a single component. Persons skilled in the art are able to select a phase-lock loop device **50** and a register **60** compatible with embodiments described herein.

In certain embodiments, the memory module **10** further comprises electrical components which are electrically coupled to one another and are surface-mounted or embedded on the printed circuit board **20**. These electrical components can include, but are not limited to, electrical conduits, resistors, capacitors, inductors, and transistors. In certain embodiments, at least some of these electrical components are discrete, while in other certain embodiments, at least some of these electrical components are constituents of one or more integrated circuits.

Various types of memory modules **10** are compatible with embodiments described herein. For example, memory modules **10** having memory capacities of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, as well as other capacities, are compatible with embodiments described herein. In addition, memory modules **10** having widths of 4 bytes, 8 bytes, 16 bytes, 32 bytes, or 32 bits, 64 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits). are compatible with embodiments described herein. Furthermore, memory modules **10** compatible with embodiments described herein include, but are not limited to, single in-line memory modules (SIMMs), dual in-line memory modules (DIMMs), small-outline DIMMs (SO-DIMMs), unbuffered DIMMs (UDIMMs), registered DIMMs (RDIMMs), fully-buffered DIMM (FB-DIMM), mini-DIMMs, and micro-DIMMs.

In certain embodiments, the printed circuit board **20** is mountable in a module slot of the computer system. The printed circuit board **20** of certain such embodiments has a plurality of edge connections electrically coupled to corresponding contacts of the module slot and to the various components of the memory module **10**, thereby providing electrical connections between the computer system and the components of the memory module **10**.

Memory devices **30** compatible with embodiments described herein include, but are not limited to, random-access memory (RAM), dynamic random-access memory (DRAM), synchronous DRAM (SDRAM), and double-data-rate DRAM (e.g., DDR-1, DDR-2, DDR-3). In addition, memory devices **30** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with embodiments described herein. Memory devices **30** compatible with embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBGA), micro-BGA (μBGA), mini-BGA (mBGA), and chip-scale packaging (CSP). Memory devices **30** compatible with embodiments described herein are available from a number of sources, including but not limited to. Samsung Semiconductor, Inc. of San Jose, Calif., Infineon Technologies AG of San Jose, Calif., and Micron Technology, Inc. of Boise, Id. Persons skilled in the art can select appropriate memory devices **30** in accordance with certain embodiments described herein.

In certain embodiments, the plurality of memory devices **30** are arranged in a first number of ranks. For example, in certain embodiments, the memory devices **30** are arranged in four ranks, as schematically illustrated by FIG. **1**A. In other embodiments, the memory devices **30** are arranged in two ranks, as schematically illustrated by FIG. **1**B. Other numbers of ranks of the memory devices **30** are also compatible with embodiments described herein.

In certain embodiments, the logic element **40** comprises a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a custom-designed semiconductor device, or a complex programmable-logic device (CPLD). In certain embodiments, the logic element **40** is a custom device. Sources of logic elements **40** compatible with embodiments described herein include, but are not limited to, Lattice Semiconductor Corporation of Hillsboro, Oreg., Altera Corporation of San Jose, Calif., and Xilinx Incorporated of San Jose, Calif. In certain embodiments, the logic element **40** comprises various discrete electrical elements, while in certain other embodiments, the logic element **40** comprises one or more integrated circuits. Persons skilled in the art can select an appropriate logic element **40** in accordance with certain embodiments described herein.

As schematically illustrated by FIGS. **1**A and **1**B, in certain embodiments, the logic element **40** receives a set of input control signals, which includes address signals (e.g., bank address signals, row address signals, column address signals, gated column address strobe signals, chip-select

signals) and command signals (e.g., refresh, precharge) from the computer system. In response to the set of input control signals, the logic element **40** generates a set of output control signals which includes address signals and command signals.

In certain embodiments, the set of output control signals corresponds to a first number of ranks in which the plurality of memory devices **30** of the memory module **10** are arranged, and the set of input control signals corresponds to a second number of ranks per memory module for which the computer system is configured. The second number of ranks in certain embodiments is smaller than the first number of ranks. For example, in the exemplary embodiment as schematically illustrated by FIG. **1**A, the first number of ranks is four while the second number of ranks is two. In the exemplary embodiment of FIG. **1**B, the first number of ranks is two while the second number of ranks is one. Thus, in certain embodiments, even though the memory module **10** actually has the first number of ranks of memory devices **30**, the memory module **10** simulates a virtual memory module by operating as having the second number of ranks of memory devices **30**. In certain embodiments, the memory module **10** simulates a virtual memory module when the number of memory devices **30** of the memory module **10** is larger than the number of memory devices **30** per memory module for which the computer system is configured to utilize.

In certain embodiments, the computer system is configured for a number of ranks per memory module which is smaller than the number of ranks in which the memory devices **30** of the memory module **10** are arranged. In certain such embodiments, the computer system is configured for two ranks of memory per memory module (providing two chip-select signals CS_0, CS_1) and the plurality of memory modules **30** of the memory module **10** are arranged in four ranks, as schematically illustrated by FIG. **1**A. In certain other such embodiments, the computer system is configured for one rank of memory per memory module (providing one chip-select signal CS_0) and the plurality of memory modules **30** of the memory module **10** are arranged in two ranks, as schematically illustrated by FIG. **1**B.

In the exemplary embodiment schematically illustrated by FIG. **1**A, the memory module **10** has four ranks of memory devices **30** and the computer system is configured for two ranks of memory devices per memory module. The memory module **10** receives row/column address signals or signal bits (A_0-A_{n+1}), bank address signals (BA_0-BA_m), chip-select signals (CS_0 and CS_1), and command signals (e.g., refresh, precharge, etc.) from the computer system. The A_0-A_n row/column address signals are received by the register **60**, which buffers these address signals and sends these address signals to the appropriate ranks of memory devices **30**. The logic element **40** receives the two chip-select signals (CS_0, CS_1) and one row/column address signal (A_{n+1}) from the computer system. Both the logic element **40** and the register **60** receive the bank address signals (BA_0-BA_m) and at least one command signal (e.g., refresh, precharge, etc.) from the computer system.

Logic Tables

Table 1 provides a logic table compatible with certain embodiments described herein for the selection among ranks of memory devices **30** using chip-select signals.

TABLE 1

State	CS_0	CS_1	A_{n+1}	Command	CS_{0A}	CS_{0B}	CS_{1A}	CS_{1B}
1	0	1	0	Active	0	1	1	1
2	0	1	1	Active	1	0	1	1
3	0	1	x	Active	0	0	1	1
4	1	0	0	Active	1	1	0	1
5	1	0	1	Active	1	1	1	0
6	1	0	x	Active	1	1	0	0
7	1	1	x	x	1	1	1	1

Note:

1. Cs_0, CS_1, CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B} are active low signals.
2. A_{n+1} is an active high signal.
3. 'x' is a Don't Care condition.
4. Command involves a number of command signals that define operations such as refresh, precharge, and other operations.

In Logic State 1: CS_0 is active low, A_{n+1} is non-active, and Command is active. CS_{0A} is pulled low, thereby selecting Rank 0.

In Logic State 2: CS_0 is active low, A_{n+1} is active, and Command is active. CS_{0B} is pulled low, thereby selecting Rank 1.

In Logic State 3: CS_0 is active low, A_{n+1} is Don't Care, and Command is active high. CS_{0A} and CS_{0B} are pulled low, thereby selecting Ranks 0 and 1.

In Logic State 4: CS_1 is active low, A_{n+1} is non-active, and Command is active. CS_{1A} is pulled low, thereby selecting Rank 2.

In Logic State 5: CS_1 is active low, A_{n+1} is active, and Command is active. CS_{1B} is pulled low, thereby selecting Rank 3.

In Logic State 6: CS_1 is active low, A_{n+1} is Don't Care, and Command is active. CS_{1A} and CS_{1B} are pulled low, thereby selecting Ranks 2 and 3.

In Logic State 7: CS_0 and CS_1 are pulled non-active high, which deselects all ranks, i.e., CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B} are pulled high.

The "Command" column of Table 1 represents the various commands that a memory device (e.g., a DRAM device) can execute, examples of which include, but are not limited to, activation, read, write, precharge, and refresh. In certain embodiments, the command signal is passed through to the selected rank only (e.g., state 4 of Table 1). In such embodiments, the command signal (e.g., read) is sent to only one memory device or the other memory device so that data is supplied from one memory device at a time. In other embodiments, the command signal is passed through to both associated ranks (e.g., state 6 of Table 1). In such embodiments, the command signal (e.g., refresh) is sent to both memory devices to ensure that the memory content of the memory devices remains valid over time. Certain embodiments utilize a logic table such as that of Table 1 to simulate a single memory device from two memory devices by selecting two ranks concurrently.

Table 2 provides a logic table compatible with certain embodiments described herein for the selection among ranks of memory devices **30** using gated CAS signals.

TABLE 2

CS*	RAS*	CAS*	WE*	Density Bit	A_{10}	Command	CAS0*	CAS1*
1	x	x	x	x	x	NOP	x	x
0	1	1	1	x	x	NOP	1	1
0	0	1	1	0	x	ACTIVATE	1	1
0	0	1	1	1	x	ACTIVATE	1	1
0	1	0	1	0	x	READ	0	1
0	1	0	1	1	x	READ	1	0
0	1	0	0	0	x	WRITE	0	1
0	1	0	0	1	x	WRITE	1	0
0	0	1	0	0	0	PRECHARGE	1	1
0	0	1	0	1	0	PRECHARGE	1	1
0	0	1	0	x	1	PRECHARGE	1	1
0	0	0	0	x	x	MODE REG SET	0	0
0	0	0	1	x	x	REFRESH	0	0

In certain embodiments in which the density bit is a row address bit, for read/write commands, the density bit is the value latched during the activate command for the selected bank.

Serial-Presence-Detect Device

Memory modules typically include a serial-presence detect (SPD) device **70** (e.g., an electrically-erasable-programmable read-only memory or EEPROM device) comprising data which characterize various attributes of the memory module, including but not limited to, the number of row addresses the number of column addresses, the data width of the memory devices, the number of ranks, the memory density per rank, the number of memory devices, and the memory density per memory device. The SPD device **70** communicates this data to the basic input/output system (BIOS) of the computer system so that the computer system is informed of the memory capacity and the memory configuration available for use and can configure the memory controller properly for maximum reliability and performance.

For example, for a commercially-available 512-MB (64M×8-byte) memory module utilizing eight 512-Mb memory devices each with a 64M×8-bit configuration, the SPD device contains the following SPD data (in appropriate bit fields of these bytes):

- Byte **3**: Defines the number of row address bits in the DRAM device in the memory module [13 for the 512-Mb memory device].
- Byte **4**: Defines the number of column address bits in the DRAM device in the memory module [11 for the 512-Mb memory device].
- Byte **13**: Defines the bit width of the primary DRAM device used in the memory module [8 bits for the 512-Mb (64M×8-bit) memory device].
- Byte **14**: Defines the bit width of the error checking DRAM device used in the memory module [8 bits for the 512-Mb (64M×8-bit) memory device].
- Byte **17**: Defines the number of banks internal to the DRAM device used in the memory module [4 for the 512-Mb memory device].

In a further example, for a commercially-available 1-GB (128M×8-byte) memory module utilizing eight 1-Gb memory devices each with a 128M×8-bit configuration, as described above, the SPD device contains the following SPD data (in appropriate bit fields of these bytes):

- Byte **3**: Defines the number of row address bits in the DRAM device in the memory module [14 for the 1-Gb memory device].
- Byte **4**: Defines the number of column address bits in the DRAM device in the memory module [11 for the 1-Gb memory device].
- Byte **13**: Defines the bit width of the primary DRAM device used in the memory module [8 bits for the 1-Gb (128M×8-bit) memory device].
- Byte **14**: Defines the bit width of the error checking DRAM device used in the memory module [8 bits for the 1-Gb (128M×8-bit) memory device].
- Byte **17**: Defines the number of banks internal to the DRAM device used in the memory module [4 for the 1-Gb memory device].

In certain embodiments, the SPD device **70** comprises data which characterize the memory module **10** as having fewer ranks of memory devices than the memory module **10** actually has, with each of these ranks having more memory density. For example, for a memory module **10** compatible with certain embodiments described herein having two ranks of memory devices **30**, the SPD device **70** comprises data which characterizes the memory module **10** as having one rank of memory devices with twice the memory density per rank. Similarly, for a memory module **10** compatible with certain embodiments described herein having four ranks of memory devices **30**, the SPD device **70** comprises data which characterizes the memory module **10** as having two ranks of memory devices with twice the memory density per rank. In addition, in certain embodiments, the SPD device **70** comprises data which characterize the memory module **10** as having fewer memory devices than the memory module **10** actually has, with each of these memory devices having more memory density per memory device. For example, for a memory module **10** compatible with certain embodiments deseribed herein, the SPD device **70** comprises data which characterizes the memory module **10** as having one-half the number of memory devices that the memory module **10** actually has, with each of these memory devices having twice the memory density per memory device.

FIG. **1C** schematically illustrates an exemplary memory module **10** in accordance with certain embodiments described herein. The memory module **10** comprises a pair of substantially identical memory devices **31**, **33**. Each memory device **31**, **33** has a first bit width, a first number of banks of memory locations, a first number of rows of memory locations, and a first number of columns of memory locations. The memory module **10** further comprises an SPD device **70** comprising data that characterizes the pair of memory devices **31**, **33**. The data characterize the pair of memory devices **31**, **33** as a virtual memory device having a second bit width equal to twice the first bit width, a second

number of banks of memory locations equal to the first number of banks, a second number of rows of memory locations equal to the first number of rows, and a second number of columns of memory locations equal to the first number of columns.

In certain such embodiments, the SPD device **70** of the memory module **10** is programmed to describe the combined pair of lower-density memory devices **31**, **33** as one virtual or pseudo-higher-density memory device. In an exemplary embodiment, two 512-Mb memory devices, each with a 128M×4-bit configuration, are used to simulate one 1-Gb memory device having a 128M×8-bit configuration. The SPD device **70** of the memory module **10** is programmed to describe the pair of 512-Mb memory devices as one virtual or pseudo-1-Gb memory device.

For example, to fabricate a 1-GB (128M×8-byte) memory module, sixteen 512-Mb (128M×4-bit) memory devices can be used. The sixteen 512-Mb (128M×4-bit) memory devices are combined in eight pairs, with each pair serving as a virtual or pseudo-1-Gb (128M×8-bit) memory device. In certain such embodiments, the SPD device **70** contains the following SPD data (in appropriate bit fields of these bytes):

Byte **3**: 13 row address bits.

Byte **4**: 12 column address bits.

Byte **13**: 8 bits wide for the primary virtual 1-Gb (128M×8-bit) memory device.

Byte **14**: 8 bits wide for the error checking virtual 1-Gb (128M×8-bit) memory device.

Byte **17**: 4 banks.

In this exemplary embodiment, bytes **3**, **4**, and **17** are programmed to have the same values as they would have for a 512-MB (128M×4-byte) memory module utilizing 512-Mb (128M×4-bit) memory devices. However, bytes **13** and **14** of the SPD data are programmed to be equal to 8, corresponding to the bit width of the virtual or pseudo-higher-density 1-Gb (128M×8-bit) memory device, for a total capacity of 1-GB. Thus, the SPD data does not describe the actual-lower-density memory devices, but instead describes the virtual or pseudo-higher-density memory devices. The BIOS accesses the SPD data and recognizes the memory module as having 4 banks of memory locations arranged in 2^{13} rows and 2^{12} columns, with each memory location having a width of 8 bits rather than 4 bits.

In certain embodiments, when such a memory module **10** is inserted in a computer system, the computer system's memory controller then provides to the memory module **10** a set of input control signals which correspond to the number of ranks or the number of memory devices reported by the SPD device **70**. For example, placing a two-rank memory module **10** compatible with certain embodiments described herein in a computer system compatible with one-rank memory modules, the SPD device **70** reports to the computer system that the memory module **10** only has one rank. The logic element **40** then receives a set of input control signals corresponding to a single rank from the computer system's memory controller, and generates and transmits a set of output control signals corresponding to two ranks to the appropriate memory devices **30** of the memory module **10**. Similarly, when a two-rank memory module **10** compatible with certain embodiments described herein is placed in a computer system compatible with either one- or two-rank memory modules, the SPD device **70** reports to the computer system that the memory module **10** only has one rank. The logic element **40** then receives a set of input control signals corresponding to a single rank from the computer system's memory controller, and generates and transmits a set of output control signals corresponding to two ranks to the appropriate memory devices **30** of the memory module **10**. Furthermore, a four-rank memory module **10** compatible with certain embodiments described herein simulates a two-rank memory module whether the memory module **10** is inserted in a computer system compatible with two-rank memory modules or with two- or four-rank memory modules. Thus, by placing a four-rank memory module **10** compatible with certain embodiments described herein in a module slot that is four-rank-ready, the computer system provides four chip-select signals, but the memory module **10** only uses two of the chip-select signals.

Memory Density Multiplication

In certain embodiments, two memory devices having a memory density are used to simulate a single memory device having twice the memory density, and an additional address signal bit is used to access the additional memory. Similarly, in certain embodiments, two ranks of memory devices having a memory density are used to simulate a single rank of memory devices having twice the memory density, and an additional address signal bit is used to access the additional memory. As used herein, such simulations of memory devices or ranks of memory devices are termed as "memory density multiplication," and the term "density transition bit" is used to refer to the additional address signal bit which is used to access the additional memory.

In certain embodiments utilizing memory density multiplication embodiments, the memory module **10** can have various types of memory devices **30** (e.g., DDR1, DDR2, DDR3, and beyond). The logic element **40** of certain such embodiments utilizes implied translation logic equations having variations depending on whether the density transition bit is a row, column, or internal bank address bit. In addition, the translation logic equations of certain embodiments vary depending on the type of memory module **10** (e.g., UDIMM, RDIMM, FBDIMM, etc.). Furthermore, in certain embodiments, the translation logic equations vary depending on whether the implementation multiplies memory devices per rank or multiplies the number of ranks per memory module.

Table 3A provides the numbers of rows and columns for DDR-1 memory devices, as specified by JEDEC standard JESD79D, "Double Data Rate (DDR) SDRAM Specification," published February 2004, and incorporated in its entirety by reference herein.

TABLE 3A

	128-Mb	256-Mb	512-Mb	1-Gb
Number of banks	4	4	4	4
Number of row address bits	12	13	13	14
Number of column address bits for "×4" configuration	11	11	12	12
Number of column address bits for "×8" configuration	10	10	11	11
Number of column address bits for "×16" configuration	9	9	10	10

As described by Table 3A, 512-Mb (128M×4-bit) DRAM devices have 2^{13} rows and 2^{12} columns of memory locations, while 1-Gb (128M×8-bit) DRAM devices have 2^{14} rows and 2^{11} columns of memory locations. Because of the differences in the number of rows and the number of columns for the two types of memory devices, complex address translation procedures and structures would typically be needed to fabricate a 1-GB (128M×8-byte) memory module using sixteen 512-Mb (128M×4-bit) DRAM devices.

Table 3B shows the device configurations as a function of memory density for DDR2 memory devices.

TABLE 3B

	Number of Rows	Number of Columns	Number of Internal Banks	Page Size (×4s or ×8s)
256 Mb	13	10	4	1 KB
512 Mb	14	10	4	1 KB
1 Gb	14	10	8	1 KB
2 Gb	15	10	8	1 KB
4 Gb	to be determined	to be determined	8	1 KB

Table 4 lists the corresponding density transition bit for the density transitions between the DDR2 memory densities of Table 3B.

TABLE 4

Density Transition	Density Transition Bit
256 Mb to 512 Mb	A_{13}
512 Mb to 1 Gb	BA_2
1 Gb to 2 Gb	A_{14}
2 Gb to 4 Gb	to be determined

Because the standard memory configuration of 4-Gb DDR2 SDRAM modules is not yet determined by the appropriate standards-setting organization, Tables 3B and 4 have "to be determined" in the appropriate table entries.

In certain embodiments, the logic translation equations are programmed in the logic element **40** by hardware, while in certain other embodiments, the logic translation equations are programmed in the logic element **40** by software. Examples 1 and 2 provide exemplary sections of Verilog code compatible with certain embodiments described herein. As described more fully below, the code of Examples 1 and 2 includes logic to reduce potential problems due to "back-to-back adjacent read commands which cross memory device boundaries or "BBARX." Persons skilled in the art are able to provide additional logic translation equations compatible with embodiments described herein.

An exemplary section of Verilog code compatible with memory density multiplication from 512 Mb to 1 Gb using DDR2 memory devices with the BA_2 density transition bit is listed below in Example 1. The exemplary code of Example 1 corresponds to a logic element **40** which receives one chip-select signal from the computer system and which generates two chip-select signals.

EXAMPLE 1

```
always @(posedge clk_in)
  begin
    rs0N_R <= rs0_in_N;      // cs0
    rasN_R <= ras_in_N;
    casN_R <= cas_in_N;
    weN_R <= we_in_N;
  end
// Gated Chip Selects
assign    pcs0a_1 = (~rs0_in_N & ~ras_in_N & ~cas_in_N)                    // ref,md reg set
        | (~rs0_in_N & ras_in_N & cas_in_N)                                // ref exit, pwr dn
        | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & a10_in)           // pchg all
        | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & ~a10_in & ~ba2_in)// pchg single bnk
        | (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N & ~ba2_in)           // activate
        | (~rs0_in_N & ras_in_N & ~cas_in_N & ~ba2_in)                     // xfr
        ;
assign    pcs0b_1 = (~rs0_io_N & ~ras_in_N & ~cas_in_N)                    // ref,tm1 reg set
        | (~rs0_in_N & ras_in_N & cas_in_N)                                // ref exit, pwr dn
        | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & a10_in)           // pchg all
        | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & ~a10_in & ba2_in)// pchg single bnk
        | (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N & ba2_in)            // activate
        | (~rs0_in_N & ras_in_N & ~cas_in_N & ba2_in)                      // xfr
        ;
//-----------------------------------
 always @(posedge_clk_in)
  begin
    a4_r <= a4_in ;
    a5_r <= a5_in ;
    a6_r <= a6_in ;
    a10_r <= a10_in ;
    ba0_r <= ba0_in ;
    ba1_r <= ba1_in ;
    ba2_r <= ba2_in ;
    q_mrs_cmd_eyc1 <= q_mrs_cmd ;
  end
////////////////////////////////////////////////////////////////
// determine the cas latency
////////////////////////////////////////////////////////////////
assign q_mrs_cmd_r = (!rasN_R & !casN_R & !weN_R)
    & !rs0N_R
```

-continued

```
    & (!ba0_r & !ba1_r)
    ;       // md reg set cmd
 always @(posedge clk_in)
  if (~reset_N)       // hur
    cl3 <= 1'b1 ;
  else if (q_mrs_cmd_cyc1)       // load mode reg cmd
  begin
    cl3 <= (~a6_r & a5_r & a4_r ) ;
  end
 always @(posedge clk_in)
  if (~reset_N)       // reset
    cl2 <= 1'b0 ;
  else if (q_mrs_cmd_cyc1)       // load mode reg cmd
  begin
    cl2 <= (~a6_r & a5_r & ~a4_r ) ;
  end
always @(posedge clk_in)
  if (~reset_N)       // reset
    cl4 <= 1'b0 ;
  else if (q_mrs_cmd_cyc1)       // load mode reg cmd
  begin
    cl4 <= (a6_r & ~a5_r & ~a4_r ) ;
  end
always @(posedge clk_in)
  if (~reset_N)       cl5 <= 1'b0 ;
  else if (q_mrs_cmd_cyc1)       // load mode reg cmd
  begin
    cl5 <= (a6_r & ~a5_r & a4_r ) ;
  end
assign    pre_cyc2_enfet = (wr_cmd_cyc1 & acs_cyc1 & cl3)   // wr brst cl3 preamble
          ;
assign    pre_cyc3_enfet = (rd_cmd_cyc2 & cl3)              // rd brst cl3 preamble
          | (wr_cmd_cyc2 & cl3)                             // wr brst cl3 1st pair
          | (wr_cmd_cyc2 & cl4)                             // wr brst cl4 preamble
          ;
assign    pre_cyc4_enfet = (wr_cmd_cyc3 & cl3)              // wr brst cl3 2nd pair
          | (wr_cmd_cyc3 & cl4)                             // wr brst cl4 1st pair
          | (rd_cmd_cyc3 & cl3)                             // rd brst cl3 1st pair
          | (rd_cmd_cyc3 & cl4)                             // rd brst cl4 preamble
          ;
assign    pre_cyc5_enfet = (rd_cmd_cyc4 & cl3)              // rd brst cl3 2nd pair
          | (wr_cmd_cyc4 & cl4)                             // wr brst cl4 2nd pair
          | (rd_cmd_cyc4 & cl4)                             // rd brst cl4 1st pair
          ;
// dq
assign    pre_dq_cyc = pre_cyc2_enfet
          | pre_cyc3_enfet
          | pre_cyc4_enfet
          | pre_cyc5_enfet
          ;
assign    pre_dq_ncyc = enfet_cyc2
          | enfet_cyc3
          | enfet_cyc4
          | enfet_cyc5
          ;
// dqs
assign    pre_dqsa_cyc = (pre_cyc2_enfet & ~ba2_r)
          | (pre_cyc3_enfet & ~ba2_cyc2)
          | (pre_cyc4_enfet & ~ba2_cyc3)
          | (pre_cyc5_enfet & ~ba2_cyc4)
          ;
assign    pre_dqsb_cyc = (pre_cyc2_enfet & ba2_r)
          | (pre_cyc3_enfet & ba2_cyc2)
          | (pre_cyc4_enfet & ba2_cyc3)
          | (pre_cyc5_enfet & ba2_cyc4)
          ;
assign    pre_dqsa_ncyc = (enfet_cyc2 & ~ba2_cyc2)
          | (enfet_cyc3 & ~ba2_cyc3)
          | (enfet_cyc4 & ~ba2_cyc4)
          | (enfet_cyc5 & ~ba2_cyc5)
          ;
assign    pre_dqsb_ncyc = (enfet_cyc2 & ba2_cyc2)
          | (enfet_cyc3 & ba2_cyc3)
          | (enfet_cyc4 & ba2_cyc4)
          | (enfet_cyc5 & ba2_cyc5)
          ;
always @(posedge clk_in)
  begin
    acs_cyc2 <= acs_cyc1 ;       // cs active
```

-continued

```
    ba2_cyc2 <= ba2_r ;
    ba2_cyc3 <= ba2_cyc2 ;
    ba2_cyc4 <= ba2_cyc3 ;
    ba2_cyc5 <= ba2_cyc4 ;
    rd_cmd_cyc2 <= rd_cmd_cyc1 & acs_cyc1;
    rd_cmd_cyc3 <= rd_cmd_cyc2 ;
    rd_cmd_cyc4 <= rd_cmd_cyc3 ;
    rd_cmd_cyc5 <= rd_cmd_cyc4 ;
    rd_cmd_cyc6 <= rd_cmd_cyc5 ;
    rd_cmd_cyc7 <= rd_cmd_cyc6 ;
    wr_cmd_cyc2 <= wr_cmd_cyc1 & acs_cyc1;
    wr_cmd_cyc3 <= wr_cmd_cyc2 ;
    wr_cmd_cyc4 <= wr_cmd_cyc3 ;
    wr_cmd_cyc5 <= wr_cmd_cyc4 ;
  end
 always @(negedge clk_in)
  begin
    dq_ncyc <= dq_cyc;
    dqs_ncyc_a <= dqs_cyc_a;
    dqs_ncyc_b <= dqs_cyc_b;
  end
// DQ FET enables
assign      enq_fet1 = dq_cyc | dq_ncyc              ;
assign      enq_fet2 = dq_cyc | dq_ncyc ;
assign      enq_fet3 = dq_cyc | dq_ncyc              ;
assign        enq_fet4 = dq_cyc | dq_ncyc ;
assign        enq_fet5 = dq_cyc | dq_ncyc ;
// DQS FET enables
assign        ens_fet1a = dqs_cyc_a | dqs_ncyc_a   ;
assign        ens_fet2a = dqs_cyc_a | dqs_ncyc_a   ;
assign        ens_fet3a = dqs_cyc_a | dqs_ncyc_a   ;
assign        ens_fet1b = dqs_cyc_b | dqs_ncyc_b   ;
assign        ens_fet2b = dqs_cyc_b | dqs_ncyc_b   ;
assign        ens_fet3b = dqs_cyc_b | dqs_ncyc_b   ;
```

Another exemplary section of Verilog code compatible with memory density multiplication from 256 Mb to 512 Mb using DDR2 memory devices and gated CAS signals with the row A_{13} density transition bit is listed below in Example 2. The exemplary code of Example 2 corresponds to a logic element **40** which receives one gated CAS signal from the computer system and which generates two gated CAS signals.

EXAMPLE 2

```
// latched a13 flags cs0, banks 0-3
 always @(posedge clk_in)
  if (actv_cmd_R & ~rs0N_R & ~bnk1_R & ~bnk0_R) // activate
  begin
        l_a13_00 <= a13_r ;
  end
 always @(posedge clk_in)
  if (actv_cmd_R & ~rs0N_R & ~bnk1_R & bnk0_R) // activate
  begin
        l_a13_01 <= a13_r ;
  end
 always @(posedge clk_in)
  if (actv_cmd_R & ~rs0N_R & bnk1_R & ~bnk0_R) // activate
  begin
        l_a13_10 <= a13_r ;
  end
 always @(posedge clk_in)
  if (actv_cmd_R & ~rs0N_R & bnk1_R & bnk0_R) // activate
  begin
        l_a13_11 <= a13_r ;
  end
// gated cas
assign cas_i = ~(casN_R);
assign cas0_o = ( ~rasN_R & cas_i)
```

-continued

```
          | ( rasN__R & ~l__a13__00 & ~bnk1__R & ~bnk0__R & cas__i)
          | ( rasN__R & ~l__a13__01 & ~bnk1__R & bnk0__R & cas__i)
          | ( rasN__R & ~l__a13__10 & bnk1__R & ~bnk0__R & cas__i)
          | ( rasN__R & ~l__a13__11 & bnk1__R & bnk0__R & cas__i)
          ;
assign cas1__o = ( ~rasN__R & cas__i)
          | ( rasN__R & l__a13__00      & ~bnk1__R & ~bnk0__R & cas__i)
          | ( rasN__R & l__a13__01      & ~bnk1__R & bnk0__R & cas__i)
          | ( rasN__R & l__a13__10      & bnk1__R & ~bnk0__R & cas__i)
          | (rasN__R & l__a13__11       & bnk1__R & bnk0__R & cas__i)
          ;
assign    pcas__0__N = ~cas0__o;
assign    pcas__1__N = ~cas1__o;
assign    rd0__o__R1 = rasN__R & cas0__o & weN__R & ~rs0N__R; // rnk0 rd cmd cyc
assign    rd1__o__R1 = rasN__R & cas1__o & weN__R & ~rs0N__R; // rnk1 rd cmd cyc
assign    wr0__o__R1 = rasN__R & cas0__o & ~weN__R & ~rs0N__R; // rnk0 wr cmd cyc
assign    wr1__o__R1 = rasN__R & cas1__o & ~weN__R & ~rs0N__R ; // rnk1 wr cmd cyc
 always @(posedge clk__in)
  begin
          rd0__o__R2 <= rd0__o__R1 ;
          rd0__o__R3 <= rd0__o__R2;
          rd0__o__R4 <= rd0__o__R3;
          rd0__o__R5 <= rd0__o__R4;
          rd1__o__R2 <= rd1__o__R1 ;
          rd1__o__R3 <= rd1__o__R2;
          rd1__o__R4 <= rd1__o__R3;
          rd1__o__R5 <= rd1__o__R4;
          wr0__o__R2 <= wr0__o__R1 ;
          wr0__o__R3 <= wr0__o__R2;
          wr0__o__R4 <= wr0__o__R3;
          wr1__o__R2 <= wr1__o__R1 ;
          wr1__o__R3 <= wr1__o__R2;
          wr1__o__R4 <= wr1__o__R3;
  end
 always @(posedge clk__in)
  begin
  if (
       (rd0__o__R2 & ~rd1__o__R4)                          // pre-am rd if no ped on rnk 1
    | rd0__o__R3                                           // 1st cyc of rd brst
    | rd0__o__R4                                           // 2nd cyc of rd brst
    | (rd0__o__R5 & ~rd1__o__R2 & ~rd1__o__R3)             // post-rd cyc if no ped on rnk 1
    | (wr0__o__R1)                                         // pre-am wr
    | wr0__o__R2 | wr0__o__R3                              // wr brst 1st & 2nd cyc
    | (wr0__o__R4)                                         // post-wr cyc (chgef9)
    | wr1__o__R1 | wr1__o__R2 | wr1__o__R3 | wr1__o__R4 // rank 1 (chgef9)
    )
          en__fet__a <= 1'b1;                              // enable fet
  else
          en__fet__a <= 1'b0;                              // disable fet
  end
 always @(posedge clk__in)
          begin
          if (
           (rd1__o__R2 & ~rd0__o__R4)
          | rd1__o__R3
          | rd1__o__R4
          | (rd1__o__R5 & ~rd0__o__R2 & ~rd0__o__R3)
        | (wr1__o__R1)                                     // (chgef8)
        | wr1__o__R2 | wr1__o__R3
      | (wr1__o__R4)                                       // post-wr cyc (chgef9)
      | wr0__o__R1 | wr0__o__R2 | wr0__o__R3 | wr0__o__R4  // rank 0 (chgef9)
          )
          en__fet__b <= 1'b1;                              //
          else
          en__fet__b <= 1'b0;
          end
```

FIG. **2**A schematically illustrates an exemplary memory module **10** which doubles the rank density in accordance with certain embodiments described herein. The memory module **10** has a first memory capacity. The memory module **10** comprises a plurality of substantially identical memory devices **30** configured as a first rank **32** and a second rank **34**. In certain embodiments, the memory devices **30** of the first rank **32** are configured in pairs, and the memory devices **30** of the second rank **34** are also configured in pairs. In certain embodiments, the memory devices **30** of the first rank **32** are configured with their respective DQS pins tied together and the memory devices **30** of the second rank **34** are configured with their respective DQS pins tied together, as described more fully below. The memory module **10** further comprises a logic element **40** which receives a first set of address and control signals from a memory controller (not shown) of the computer system. The first set of address and control signals is compatible with a second memory capacity substantially

equal to one-half of the first memory capacity. The logic element **40** translates the first set of address and control signals into a second set of address and control signals which is compatible with the first memory capacity of the memory module **10** and which is transmitted to the first rank **32** and the second rank **34**.

The first rank **32** of FIG. **2**A has 18 memory devices **30** and the second rank **34** of FIG. **2**A has 18 memory devices **30**. Other numbers of memory devices **30** in each of the ranks **32**, **34** are also compatible with embodiments described herein.

In the embodiment schematically illustrated by FIG. **2**A, the memory module **10** has a width of 8 bytes (or 64 bits) and each of the memory devices **30** of FIG. **2**A has a bit width of 4 bits. The 4-bit-wide ("×4") memory devices **30** of FIG. **2**A have one-half the width, but twice the depth of 8-bit-wide ("×8") memory devices. Thus, each pair of "×4" memory devices **30** has the same density as a single "×8" memory device, and pairs of "×4" memory devices **30** can be used instead of individual "×8" memory devices to provide the memory density of the memory module **10**. For example, a pair of 512-Mb 128M×4-bit memory devices has the same memory density as a 1-Gb 128M×8-bit memory device.

For two "×4" memory devices **30** to work in tandem to mimic a "×8" memory device, the relative DQS pins of the two memory devices **30** in certain embodiments are advantageously tied together, as described more fully below. In addition, to access the memory density of a high-density memory module **10** comprising pairs of "×4" memory devices **30**, an additional address line is used. While a high-density memory module comprising individual "×8" memory devices with the next-higher density would also utilize an additional address line, the additional address lines are different in the two memory module configurations.

For example, a 1-Gb 128M×8-bit DDR-1 DRAM memory device uses row addresses A_{13}-A_0 and column addresses A_{11} and A_9-A_0. A pair of 512-Mb 128M×4-bit DDR-1 DRAM memory devices uses row addresses A_{12}-A_0 and column addresses A_{12}, A_{11}, and A_9-A_0. In certain embodiments, a memory controller of a computer system utilizing a 1-GB 128M×8 memory module **10** comprising pairs of the 512-Mb 128M×4 memory devices **30** supplies the address and control signals including the extra row address (A_{13}) to the memory module **10**. The logic element **40** receives the address and control signals from the memory controller and converts the extra row address (A_{13}) into an extra column address (A_{12}).

FIG. **2**B schematically illustrates an exemplary logic element **40** compatible with embodiments described herein. The logic element **40** is used for a memory module **10** comprising pairs of "×4" memory devices **30** which mimic individual "×8" memory devices. In certain embodiments, each pair has the respective DQS pins of the memory devices **30** tied together. In certain embodiments, as schematically illustrated by FIG. **2**B, the logic element **40** comprises a programmable-logic device (PLD) **42**, a first multiplexer **44** electrically coupled to the first rank **32** of memory devices **30**, and a second multiplexer **46** electrically coupled to the second rank **34** of memory devices **30**. In certain embodiments, the PLD **42** and the first and second multiplexers **44**, **46** are discrete elements, while in, other certain embodiments, they are integrated within a single integrated circuit. Persons skilled in the art can select an appropriate PLD **42**, first multiplexer **44**, and second multiplexer **46** in accordance with embodiments described herein.

In the exemplary logic element **40** of FIG. **2**B, during a row access procedure (CAS is high), the first multiplexer **44** passes the A_{12} address through to the first rank **32**, the second multiplexer **46** passes the A_{12} address through to the second rank **34**, and the PLD **42** saves or latches the A_{13} address from the memory controller. In certain embodiments, a copy of the A_{13} address is saved by the PLD **42** for each of the internal banks (e.g., 4 internal banks) per memory device **30**. During a subsequent column access procedure (CAS is low), the first multiplexer **44** passes the previously-saved A_{13} address through to the first rank **32** as the A_{12} address and the second multiplexer **46** passes the previously-saved A_{13} address through to the second rank **34** as the A_{12} address. The first rank **32** and the second rank **34** thus interpret the previously-saved A_{13} row address as the current A_{12} column address. In this way, in certain embodiments, the logic element **40** translates the extra row address into an extra column address in accordance with certain embodiments described herein.

Thus, by allowing two lower-density memory devices to be used rather than one higher-density memory device, certain embodiments described herein provide the advantage of using lower-cost, lower-density memory devices to build "next-generation" higher-density memory modules. Certain embodiments advantageously allow the use of lower-cost readily-available 512-Mb DDR-2 SDRAM devices to replace more expensive 1-Gb DDR-2 SDRAM devices. Certain embodiments advantageously reduce the total cost of the resultant memory module.

FIG. **3**A schematically illustrates an exemplary memory module **10** which doubles number of ranks in accordance with certain embodiments described herein. The memory module **10** has a first plurality of memory locations with a first memory density. The memory module **10** comprises a plurality of substantially identical memory devices **30** configured as a first rank **32**, a second rank **34**, a third rank **36**, and a fourth rank **38**. The memory module **10** further comprises a logic element **40** which receives a first set of address and control signals from a memory controller (not shown). The first set of address and control signals is compatible with a second plurality of memory locations having a second memory density. The second memory density is substantially equal to one-half of the first memory density. The logic element **40** translates the first set of address and control signals into a second set of address and control signals which is compatible with the first plurality of memory locations of the memory module **10** and which is transmitted to the first rank **32**, the second rank **34**, the third rank **36**, and the fourth rank **38**.

Each rank **32**, **34**, **36**, **38** of FIG. **3**A has 9 memory devices **30**. Other numbers of memory devices **30** in each of the ranks **32**, **34**, **36**, **38** are also compatible with embodiments described herein.

In the embodiment schematically illustrated by FIG. **3**A, the memory module **10** has a width of 8 bytes (or 64 bits) and each of the memory devices **30** of FIG. **3**A has a bit width of 8 bits. Because the memory module **10** has twice the number of 8-bit-wide ("×8") memory devices **30** as does a standard 8-byte-wide memory module, the memory module **10** has twice the density as does a standard 8-byte-wide memory module. For example, a 1-GB 128M×8-byte memory module with 36 512-Mb 128M×8-bit memory devices (arranged in four ranks) has twice the memory density as a 512-Mb 128M×8-byte memory module with 18 512-Mb 128M×8-bit memory devices (arranged in two ranks).

To access the additional memory density of the high-density memory module **10**, the two chip-select signals (CS_0, CS_1) are used with other address and control signals to gate a set of four gated CAS signals. For example, to access the additional ranks of four-rank 1-GB 128Mx8-byte DDR-1 DRAM memory module, the CS_0 and CS_1 signals along with the other address and control signals are used to gate the CAS signal appropriately, as schematically illustrated by FIG. **3**A. FIG. **3**B schematically illustrates an exemplary logic element **40** compatible with embodiments described herein. In certain embodiments, the logic element **40** comprises a programmable-logic device (PLD) **42** and four "OR" logic elements **52**, **54**, **56**, **58** electrically coupled to corresponding ranks **32**, **34**, **36**, **38** of memory devices **30**.

In certain embodiments, the PLD **42** comprises an ASIC, an FPGA, a custom-designed semiconductor device, or a CPLD. In certain embodiments, the PLD **42** and the four "OR" logic elements **52**, **54**, **56**, **58** are discrete elements, while in other certain embodiments, they are integrated within a single integrated circuit. Persons skilled in the art can select an appropriate PLD **42** and appropriate "OR" logic elements **52**, **54**, **56**, **58** in accordance with embodiments described herein.

In the embodiment schematically illustrated by FIG. **3**B, the PLD **42** transmits each of the four "enabled CAS" ($ENCAS_0a$, $ENCAS_0b$, $ENCAS_1a$, $ENCAS_1b$) signals to a corresponding one of the "OR" logic elements **52**, **54**, **56**, **58**. The CAS signal is also transmitted to each of the four "OR" logic elements **52**, **54**, **56**, **58**. The CAS signal and the "enabled CAS" signals are "low" true signals. By selectively activating each of the four "enabled CAS" signals which are inputted into the four "OR" logic elements **52**, **54**, **56**, **58**, the PLD **42** is able to select which of the four ranks **32**, **34**, **36**, **38** is active.

In certain embodiments, the PLD **42** uses sequential and combinatorial logic procedures to produce the gated CAS signals which are each transmitted to a corresponding one of the four ranks **32**, **34**, **36**, **38**. In certain other embodiments, the PLD **42** instead uses sequential and combinatorial logic procedures to produce four gated chip-select signals (e.g., CS_0a, CS_0b, CS_1a, and CS_1b) which are each transmitted to a corresponding one of the four ranks **32**, **34**, **36**, **38**.

Back-to-Back Adjacent Read Commands

Due to their source synchronous nature, DDR SDRAM (e.g., DDR1, DDR2, DDR3) memory devices operate with a data transfer protocol which surrounds each burst of data strobes with a pre-amble time interval and a post-amble time interval. The pre-amble time interval provides a timing window for the receiving memory device to enable its data capture circuitry when a known valid level is present on the strobe signal to avoid false triggers of the memory device's capture circuit. The post-amble time interval provides extra time after the last strobe for this data capture to facilitate good signal integrity. In certain embodiments, when the computer system accesses two consecutive bursts of data from the same memory device, termed herein as a "back-to-back adjacent read," the post-amble time interval of the first read command and the pre-amble time interval of the second read command are skipped by design protocol to increase read efficiency. FIG. **4**A shows an exemplary timing diagram of this "gapless" read burst for a back-to-back adjacent read condition from one memory device.

In certain embodiments, when the second read command accesses data from a different memory device than does the first read command, there is at least one time interval (e.g., clock cycle) inserted between the data strobes of the two memory devices. This inserted time interval allows both read data bursts to occur without the post-amble time interval of the first read data burst colliding or otherwise interfering with the pre-amble time interval of the second read data burst. In certain embodiments, the memory controller of the computer system inserts an extra clock cycle between successive read commands issued to different memory devices, as shown in the exemplary timing diagram of FIG. **4**B for successive read accesses from different memory devices.

In typical computer systems, the memory controller is informed of the memory boundaries between the ranks of memory of the memory module prior to issuing read commands to the memory module. Such memory controllers can insert wait time intervals or clock cycles to avoid collisions or interference between back-to-back adjacent read commands which cross memory device boundaries, which are referred to herein as "BBARX."

In certain embodiments described herein in which the number of ranks of the memory module is doubled or quadrupled, the logic element **40** generates a set of output control signals so that the selection decoding is transparent to the computer system. However, in certain such embodiments, there are memory device boundaries of which the computer system is unaware, so there are occasions in which BBARX occurs without the cognizance of the memory controller of the computer system. As shown in FIG. **5**, the last data strobe of memory device "a" collides with the pre-amble time interval of the data strobe of memory device "b," resulting in a "collision window."

FIG. **6**A schematically illustrates a circuit diagram of a conventional memory module **100** showing the interconnections between the DQ data signal lines **102** of the memory devices "a" and "b" (not shown) and their DQS data strobe signal lines **104**. In certain embodiments, the electrical signal lines are etched on the printed circuit board. As shown in FIG. **6**A, each of the memory devices has their DQ data signal lines **102** electrically coupled to a common DQ line **112** and the DQS data strobe signal lines **104** electrically coupled to a common DQS line **114**.

In certain embodiments, BBARX collisions are avoided by a mechanism which electrically isolates the DQS data strobe signal lines **104** of the memory devices from one another during the transition from the first read data burst of one rank of memory devices to the second read data burst of another rank of memory devices. FIG. **6**B schematically illustrates a circuit diagram of an exemplary memory module **10** comprising an isolation device **120** in accordance with certain embodiments described herein. As shown in FIG. **6**B, each of the memory devices **30** otherwise involved in a BBARX collision have their DQS data strobe signal lines **104** electrically coupled to the common DQS line **114** through the isolation element **120**. The isolation device **120** of certain embodiments multiplexes the DQS data strobe signal lines **104** of the two ranks of memory devices **30** from one another to avoid a BBARX collision.

In certain embodiments, as schematically illustrated by FIG. **6**B, the isolation device **120** comprises a first switch **122** electrically coupled to a first data strobe signal line (e.g., DQSa) of a first memory device (not shown) and a second switch **124** electrically coupled to a second data strobe signal line (e.g., DQSb) of a second memory device (not shown). Exemplary switches compatible with embodiments described herein include, but are not limited to field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex. In certain embodiments, the time

for switching the first switch **122** and the second switch **124** is between the two read data bursts (e.g., after the last DQS data strobe of the read data burst of the first memory device and before the first DQS data strobe of the read data burst of the second memory device). During the read data burst for a first memory device, the first switch **122** is enabled. After the last DQS data strobe of the first memory device and before the first DQS data strobe of the second memory device, the first switch **122** is disabled and the second switch **124** is enabled.

In certain embodiments, as schematically illustrated by FIG. **6**C, the isolation device **120** comprises a logic element **140** which multiplexes the DQS data strobe signal lines **104** from one another. Exemplary logic elements compatible with embodiments described herein include, but are not limited to multiplexers, such as the SN74AUC2G53 2:1 analog multiplexer/demultiplexer available from Texas Instruments, Inc. of Dallas, Tex. The logic element **140** receives a first DQS data strobe signal from the first memory device and a second DQS data strobe signal from a second memory device and selectively switches one of the first and second DQS data strobe signals to the common DQS data strobe signal line **114**. Persons skilled in the art can select other types of isolation devices **120** compatible with embodiments described herein.

In certain embodiments, as schematically illustrated by FIG. **6**D, the isolation device **120** also multiplexes the DQ data signal lines **102** of the two memory devices from one another. For example, in certain embodiments, the isolation device **120** comprises a pair of switches on the DQ data signal lines **102**, similar to the switches **122**, **124** on the DQS data strobe signal lines **104** schematically illustrated by FIG. **6**B. In certain other embodiments, the isolation device **120** comprises a logic element **150**, as schematically illustrated by FIG. **6**D. In certain embodiments, the same types of switches and/or logic elements are used for the DQ data signal lines **102** as for the DQS data strobe signal lines **104**. The logic element **150** receives a first DQ data signal from the first memory device and a second DQ data signal from the second memory device and selectively switches one of the first and second DQ data signals to the common DQ data signal line **112**. Persons skilled in the art can select other types of isolation devices **120** compatible with embodiments described herein.

In certain embodiments, the isolation device **120** advantageously adds propagation delays to the DQ data signals which match the DQS strobe signals being multiplexed by the isolation device **120**. In certain embodiments, the isolation device **120** advantageously presents a reduced impedance load to the computer system by selectively switching between the two ranks of memory devices to which it is coupled. This feature of the isolation device **120** is used in certain embodiments in which there is no memory density multiplication of the memory module (e.g., for a computer system with four chip-select signals), but where the impedance load of the memory module may otherwise limit the number of ranks or the number of memory devices per memory module. As schematically illustrated by FIG. **6**E, the isolation device **120** of certain such embodiments comprises the logic element **150** on the DQ data signal lines but not a corresponding logic element on the DQS data strobe signal lines.

In certain embodiments, the control and timing of the isolation device **120** is performed by an isolation-control logic element (e.g., application-specific integrated circuit, custom programmable logic device, field-programmable gate array, etc.) which is resident on the memory module **10**. In certain embodiments, the isolation-control logic element is the same logic element **40** as schematically illustrated in FIGS. **1**A and **1**B, is part of the isolation device **120** (e.g., logic element **140** or logic element **150** as schematically illustrated by FIG. **6**D), or is a separate component. The isolation-control logic element of certain embodiments controls the isolation device **120** by monitoring commands received by the memory module **10** from the computer system and producing "windows" of operation whereby the appropriate components of the isolation device **120** are switched to enable and disable the DQS data strobe signal lines **104** to mitigate BBARX collisions. In certain other embodiments, the isolation-control logic element monitors the commands received by the memory module from the computer system and selectively enables and disables the DQ data signal lines **102** to reduce the load impedance of the memory module **10** on the computer system. In still other embodiments, this logic element performs both of these functions together.

Tied Data Strobe Signal Pins

For proper operation, the computer system advantageously recognizes a 1-GB memory module comprising 256-Mb memory devices with 64M×4-bit configuration as a 1-GB memory module having 512-Mb memory devices with 64M×8-bit configuration (e.g., as a 1-GB memory module with 128M×8-byte configuration). This advantageous result is desirably achieved in certain embodiments by electrically connecting together two output signal pins (e.g., DQS or data strobe pins) of the two 256-Mb memory devices such that both output signal pins are concurrently active when the two memory devices are concurrently enabled. The DQS or data strobe is a bi-directional signal that is used during both read cycles and write cycles to validate or latch data. As used herein, the terms "tying together" or "tied together" refer to a configuration in which corresponding pins (e.g., DQS pins) of two memory devices are electrically connected together and are concurrently active when the two memory devices are concurrently enabled (e.g., by a common chip-select or CS signal). Such a configuration is different from standard memory module configurations in which the output signal pins (e.g., DQS pins) of two memory devices are electrically coupled to the same source, but these pins are not concurrently active since the memory devices are not concurrently enabled. However, a general guideline of memory module design warns against tying together two output signal pins in this way.

FIGS. **7** and **8** schematically illustrate a problem which may arise from tying together two output signal pins. FIG. **7** schematically illustrates an exemplary memory module **205** in which a first DQS pin **212** of a first memory device **210** is electrically connected to a second DQS pin **222** of a second memory device **220**. The two DQS pins **212**, **222** are both electrically connected to a memory controller **230**.

FIG. **8** is an exemplary timing diagram of the voltages applied to the two DQS pins **212**, **222** due to non-simultaneous switching. As illustrated by FIG. **8**, at time t_1, both the first DQS pin **212** and the second DQS pin **222** are high, so no current flows between them. Similarly, at time t_4, both the first DQS pin **212** and the second DQS pin **222** are low, so no current flows between them. However, for times between approximately t_2 and approximately t_3, the first DQS pin **212** is low while the second DQS pin **222** is high. Under such conditions, a current will flow between the two DQS pins **212**, **222**. This condition in which one DQS pin is low while the other DQS pin is high can occur for fractions of a second (e.g., 0.8 nanoseconds) during the dynamic random-access

memory (DRAM) read cycle. During such conditions, the current flowing between the two DQS pins **212**, **222** can be substantial, resulting in heating of the memory devices **210**, **220**, and contributing to the degradation of reliability and eventual failure of these memory devices.

A second problem may also arise from tying together two output signal pins. FIG. **9** schematically illustrates another exemplary memory module **205** in which a first DQS pin **212** of a first memory device **210** is electrically connected to a second DQS pin **214** of a second memory device **220**. The two DQS pins **212**, **214** of FIG. **9** are both electrically connected to a memory controller (not shown). The DQ (data input/output) pin **222** of the first memory device **210** and the corresponding DQ pin **224** of the second memory device **220** are each electrically connected to the memory controller by the DQ bus (not shown). Typically, each memory device **210**, **220** will have a plurality of DQ pins (e.g., eight DQ pins per memory device), but for simplicity, FIG. **9** only shows one DQ pin for each memory device **210**, **220**.

Each of the memory devices **210**, **220** of FIG. **9** utilizes a respective on-die termination or "ODT" circuit **232**, **234** which has termination resistors (e.g., 75 ohms) internal to the memory devices **210**, **220** to provide signal termination. Each memory device **210**, **220** has a corresponding ODT signal pin **262**, **264** which is electrically connected to the memory controller via an ODT bus **240**. The ODT signal pin **262** of the first memory device **210** receives a signal from the ODT bus **240** and provides the signal to the ODT circuit **232** of the first memory device **210**. The ODT circuit **232** responds to the signal by selectively enabling or disabling the internal termination resistors **252**, **256** of the first memory device **210**. This behavior is shown schematically in FIG. **9** by the switches **242**, **244** which are either closed (dash-dot line) or opened (solid line). The ODT signal pin **264** of the second memory device **220** receives a signal from the ODT bus **240** and provides the signal to the ODT circuit **234** of the second memory device **220**. The ODT circuit **234** responds to the signal by selectively enabling or disabling the internal termination resistors **254**, **258** of the second memory device **220**. This behavior is shown schematically in FIG. **9** by the switches **246**, **248** which are either closed (dash-dot line) or opened (solid line). The switches **242**, **244**, **246**, **248** of FIG. **9** are schematic representations of the operation of the ODT circuits **232**, **234**, and do not signify that the ODT circuits **232**, **234** necessarily include mechanical switches.

Examples of memory devices **210**, **220** which include such ODT circuits **232**, **234** include, but are not limited to, DDR2 memory devices. Such memory devices are configured to selectively enable or disable the termination of the memory device in this way in response to signals applied to the ODT signal pin of the memory device. For example, when the ODT signal pin **262** of the first memory device **210** is pulled high, the termination resistors **252**, **256** of the first memory device **210** are enabled. When the ODT signal pin **262** of the first memory device **210** is pulled low (e.g., grounded), the termination resistors **252**, **256** of the first memory device **210** are disabled. By selectively disabling the termination resistors of an active memory device, while leaving the termination resistors of inactive memory devices enabled, such configurations advantageously preserve signal strength on the active memory device while continuing to eliminate signal reflections at the bus-die interface of the inactive memory devices.

In certain configurations, as schematically illustrated by FIG. **9**, the DQS pins **212**, **214** of each memory device **210**, **220** are selectively connected to a voltage VTT through a corresponding termination resistor **252**, **254** internal to the corresponding memory device **210**, **220**. Similarly, in certain configurations, as schematically illustrated by FIG. **9**, the DQ pins **222**, **224** are selectively connected to a voltage VTT through a corresponding termination resistor **256**, **258** internal to the corresponding memory device **210**, **220**. In certain configurations, rather than being connected to a voltage VTT, the DQ pins **212**, **214** and/or the DQS pins **222**, **224** are selectively connected to ground through the corresponding termination resistors **252**, **254**, **256**, **258**. The resistances of the internal termination resistors **252**, **254**, **256**, **258** are selected to clamp the voltages so as to reduce the signal reflections from the corresponding pins. In the configuration schematically illustrated by FIG. **9**, each internal termination resistor **252**, **254**, **256**, **258** has a resistance of approximately 75 ohms.

When connecting the first memory device **210** and the second memory device **220** together to form a double word width, both the first memory device **210** and the second memory device **220** are enabled at the same time (e.g., by a common CS signal). Connecting the first memory device **210** and the second memory device **220** by tying the DQS pins **212**, **214** together, as shown in FIG. **9**, results in a reduced effective termination resistance for the DQS pins **212**, **214**. For example, for the exemplary configuration of FIG. **9**, the effective termination resistance for the DQS pins **212**, **214** is approximately 37.5 ohms, which is one-half the desired ODT resistance (for 75-ohm internal termination resistors) to reduce signal reflections since the internal termination resistors **252**, **254** of the two memory devices **210**, **220** are connected in parallel. This reduction in the termination resistance can result in signal reflections causing the memory device to malfunction.

FIG. **10** schematically illustrates an exemplary memory module **300** in accordance with certain embodiments described herein. The memory module **300** comprises a first memory device **310** having a first data strobe (DQS) pin **312** and a second memory device **320** having a second data strobe (DQS) pin **322**. The memory module **300** further comprises a first resistor **330** electrically coupled to the first DQS pin **312**. The memory module **300** further comprises a second resistor **340** electrically coupled to the second DQS pin **322** and to the first resistor **330**. The first DQS pin **312** is electrically coupled to the second DQS pin **322** through the first resistor **330** and through the second resistor **340**.

In certain embodiments, the memory module **300** is a 1-GB unbuffered Double Data Rate (DDR) Synchronous Dynamic RAM (SDRAM) high-density dual in-line memory module (DIMM). FIGS. **11A** and **11B** schematically illustrate a first side **362** and a second side **364**, respectively, of such a memory module **300** with eighteen 64M×4-bit, DDR-1 SDRAM FBGA memory devices on each side of a 184-pin glass-epoxy printed circuit board (PCB) **360**. In certain embodiments, the memory module **300** further comprises a phase-lock-loop (PLL) clock driver **370**, an EEPROM for serial-presence detect (SPD) data **380**, and decoupling capacitors (not shown) mounted on the PCB in parallel to suppress switching noise on VDD and VDDQ power supply for DDR-1 SDRAM. By using synchronous design, such memory modules **300** allow precise control of data transfer between the memory module **300** and the system controller. Data transfer can take place on both edges of the DQS signal at various operating frequencies and programming latencies. Therefore, certain such memory modules **300** are suitable for a variety of high-performance system applications.

In certain embodiments, the memory module **300** comprises a plurality of memory devices configured in pairs, each pair having a first memory device **310** and a second memory device **320**. For example, in certain embodiments, a 128Mx72-bit DDR SDRAM high-density memory module **300** comprises thirty-six 64Mx4-bit DDR-1 SDRAM integrated circuits in FBGA packages configured in eighteen pairs. The first memory device **310** of each pair has the first DQS pin **312** electrically coupled to the second DQS pin **322** of the second memory device **320** of the pair. In addition, the first DQS pin **312** and the second DQS pin **322** are concurrently active when the first memory device **310** and the second memory device **320** are concurrently enabled.

In certain embodiments, the first resistor **330** and the second resistor **340** each has a resistance advantageously selected to reduce the current flow between the first DQS pin **312** and the second DQS pin **322** while allowing signals to propagate between the memory controller and the DQS pins **312**, **322**. In certain embodiments, each of the first resistor **330** and the second resistor **340** has a resistance in a range between approximately 5 ohms and approximately 50 ohms. For example, in certain embodiments, each of the first resistor **330** and the second resistor **340** has a resistance of approximately 22 ohms. Other resistance values for the first resistor **330** and the second resistor **340** are also compatible with embodiments described herein. In certain embodiments, the first resistor **330** comprises a single resistor, while in other embodiments, the first resistor **330** comprises a plurality of resistors electrically coupled together in series and/or in parallel. Similarly, in certain embodiments, the second resistor **340** comprises a single resistor, while in other embodiments, the second resistor **340** comprises a plurality of resistors electrically coupled together in series and/or in parallel.

FIGS. **12**A and **12**B schematically illustrate an exemplary embodiment of a memory module **300** in which the first resistor **330** and the second resistor **340** are used to reduce the current flow between the first DQS pin **312** and the second DQS pin **322**. As schematically illustrated by FIG. **12**A, the memory module **300** is part of a computer system **400** having a memory controller **410**. The first resistor **330** has a resistance of approximately 22 ohms and the second resistor **340** has a resistance of approximately 22 ohms. The first resistor **330** and the second resistor **340** are electrically coupled in parallel to the memory controller **410** through a signal line **420** having a resistance of approximately 25 ohms. The first resistor **330** and the second resistor **340** are also electrically coupled in parallel to a source of a fixed termination voltage (identified by VTT in FIGS. **12**A and **12**B) by a signal line **440** having a resistance of approximately 47 ohms. Such an embodiment can advantageously be used to allow two memory devices having lower bit widths (e.g., 4-bit) to behave as a single virtual memory device having a higher bit width (e.g., 8-bit).

FIG. **12**B schematically illustrates exemplary current-limiting resistors **330**, **340** in conjunction with the impedances of the memory devices **310**, **320**. During an exemplary portion of a data read operation, the memory controller **410** is in a high-impedance condition, the first memory device **310** drives the first DQS pin **312** high (e.g., 2.7 volts), and the second memory device **320** drives the second DQS pin **322** low (e.g., 0 volts). The amount of time for which this condition occurs is approximated by the time between t_2 and t_3 of FIG. **8**, which in certain embodiments is approximately twice the tDQSQ (data strobe edge to output data edge skew time, e.g., approximately 0.8 nanoseconds). At least a portion of this time in certain embodiments is caused by simultaneous switching output (SSO) effects.

In certain embodiments, as schematically illustrated by FIG. **12**B, the DQS driver of the first memory device **310** has a driver impedance R_1 of approximately 17 ohms, and the DQS driver of the second memory device **320** has a driver impedance R_4 of approximately 17 ohms. Because the upper network of the first memory device **310** and the first resistor **330** (with a resistance R_2 of approximately 22 ohms) is approximately equal to the lower network of the second memory device **320** and the second resistor **340** (with a resistance R_3 of approximately 22 ohms), the voltage at the midpoint is approximately 0.5*(2.7−0)=1.35 volts, which equals VTT, such that the current flow across the 47-ohm resistor of FIG. **12**B is approximately zero.

The voltage at the second DQS pin **322** in FIG. **12**B is given by $V_{DQS2}=2.7*R_4/(R_1+R_2+R_3+R_4)=0.59$ volts and the current flowing through the second DQS pin **322** is given by $I_{DQS2}=0.59/R_4=34$ milliamps. The power dissipation in the DQS driver of the second memory device **320** is thus $P_{DQS2}=34$ mA*0.59 V=20 milliwatts. In contrast, without the first resistor **330** and the second resistor **340**, only the 17-ohm impedances of the two memory devices **310**, **320** would limit the current flow between the two DQS pins **312**, **322**, and the power dissipation in the DQS driver of the second memory device **320** would be approximately 107 milliwatts. Therefore, the first resistor **330** and the second resistor **340** of FIGS. **12**A and **12**B advantageously limit the current flowing between the two memory devices during the time that the DQS pin of one memory device is driven high and the DQS pin of the other memory device is driven low.

In certain embodiments in which there is overshoot or undershoot of the voltages, the amount of current flow can be higher than those expected for nominal voltage values. Therefore, in certain embodiments, the resistances of the first resistor **330** and the second resistor **340** are advantageously selected to account for such overshoot/undershoot of voltages.

For certain such embodiments in which the voltage at the second DQS pin **322** is $V_{DQS2}=0.59$ volts and the duration of the overdrive condition is approximately 0.8 nanoseconds at maximum, the total surge is approximately 0.59 V*1.2 ns=0.3 V-ns. For comparison, the JEDEC standard for overshoot/undershoot is 2.4 V-ns, so certain embodiments described herein advantageously keep the total surge within predetermined standards (e.g., JEDEC standards).

FIG. **13** schematically illustrates another exemplary memory module **500** compatible with certain embodiments described herein. The memory module **500** comprises a termination bus **505**. The memory module **500** further comprises a first memory device **510** having a first data strobe pin **512**, a first termination signal pin **514** electrically coupled to the termination bus **505**, a first termination circuit **516**, and at least one data pin **518**. The first termination circuit **516** selectively electrically terminating the first data strobe pin **512** and the first data pin **518** in response to a first signal received by the first termination signal pin **514** from the termination bus **505**. The memory module **500** further comprises a second memory device **520** having a second data strobe pin **522** electrically coupled to the first data strobe pin **512**, a second termination signal pin **524**, a second termination circuit **526**, and at least one data pin **528**. The second termination signal pin **524** is electrically coupled to a voltage, wherein the second termination circuit **526** is responsive to the voltage by not terminating the second data strobe pin **522** or the second data pin **528**. The memory module **500** further comprises at least one termination

assembly **530** having a third termination signal pin **534**, a third termination circuit **536**, and at least one termination pin **538** electrically coupled to the data pin **528** of the second memory device **520**. The third termination signal pin **534** is electrically coupled to the termination bus **505**. The third termination circuit **536** selectively electrically terminates the data pin **528** of the second memory device **520** through the termination pin **538** in response to a second signal received by the third termination signal pin **534** from the termination bus **505**.

FIG. **14** schematically illustrates a particular embodiment of the memory module **500** schematically illustrated by FIG. **13**. The memory module **500** comprises an on-die termination (ODT) bus **505**. The memory module **500** comprises a first memory device **510** having a first data strobe (DQS) pin **512**, a first ODT signal pin **514** electrically coupled to the ODT bus **505**, a first ODT circuit **516**, and at least one data (DQ) pin **518**. The first ODT circuit **516** selectively electrically terminates the first DQS pin **512** and the DQ pin **518** of the first memory device **510** in response to an ODT signal received by the first ODT signal pin **514** from the ODT bus **505**. This behavior of the first ODT circuit **516** is schematically illustrated in FIG. **14** by the switches **572**, **576** which are selectively closed (dash-dot line) or opened (solid line).

The memory module **500** further comprises a second memory device **520** having a second DQS pin **522** electrically coupled to the first DQS pin **512**, a second ODT signal pin **524**, a second ODT circuit **526**, and at least one DQ pin **528**. The first DQS pin **512** and the second DQS pin **522** are concurrently active when the first memory device **510** and the second memory device **520** are concurrently enabled. The second ODT signal pin **524** is electrically coupled to a voltage (e.g., ground), wherein the second ODT circuit **526** is responsive to the voltage by not terminating the second DQS pin **522** or the second DQ pin **524**. This behavior of the second ODT circuit **526** is schematically illustrated in FIG. **14** by the switches **574**, **578** which are opened.

The memory module **500** further comprises at least one termination assembly **530** having a third ODT signal pin **534** electrically coupled to the ODT bus **505**, a third ODT circuit **536**, and at least one termination pin **538** electrically coupled to the DQ pin **528** of the second memory device **520**. The third ODT circuit **536** selectively electrically terminates the DQ pin **528** of the second memory device **520** through the termination pin **538** in response to an ODT signal received by the third ODT signal pin **534** from the ODT bus **505**. This behavior of the third ODT circuit **536** is schematically illustrated in FIG. **14** by the switch **580** which is either closed (dash-dot line) or opened (solid line).

In certain embodiments, the termination assembly **530** comprises discrete electrical components which are surface-mounted or embedded on the printed-circuit board of the memory module **500**. In certain other embodiments, the termination assembly **530** comprises an integrated circuit mounted on the printed-circuit board of the memory module **500**. Persons skilled in the art can provide a termination assembly **530** in accordance with embodiments described herein.

Certain embodiments of the memory module **500** schematically illustrated by FIG. **14** advantageously avoid the problem schematically illustrated by FIG. **7** of electrically connecting the internal termination resistances of the DQS pins of the two memory devices in parallel. As described above in relation to FIG. **9**, FIGS. **13** and **14** only show one DQ pin for each memory device for simplicity. Other embodiments have a plurality of DQ pins for each memory device. In certain embodiments, each of the first ODT circuit **516**, the second ODT circuit **526**, and the third ODT circuit **536** are responsive to a high voltage or signal level by enabling the corresponding termination resistors and are responsive to a low voltage or signal level (e.g., ground) by disabling the corresponding termination resistors. In other embodiments, each of the first ODT circuit **516**, the second ODT circuit **526**, and the third ODT circuit **536** are responsive to a high voltage or signal level by disabling the corresponding termination resistors and are responsive to a low voltage or signal level (e.g., ground) by enabling the corresponding termination resistors. Furthermore, the switches **572**, **574**, **576**, **578**, **580** of FIG. **14** are schematic representations of the enabling and disabling operation of the ODT circuits **516**, **526**, **536** and do not signify that the ODT circuits **516**, **526**, **536** necessarily include mechanical switches.

The first ODT signal pin **514** of the first memory device **510** receives an ODT signal from the ODT bus **505**. In response to this ODT signal, the first ODT circuit **516** selectively enables or disables the termination resistance for both the first DQS pin **512** and the DQ pin **518** of the first memory device **510**. The second ODT signal pin **524** of the second memory device **520** is tied (e.g., directly hard-wired) to the voltage (e.g., ground), thereby disabling the internal termination resistors **554**, **558** on the second DQS pin **522** and the second DQ pin **528**, respectively, of the second memory device **520** (schematically shown by open switches **574**, **578** in FIG. **14**). The second DQS pin **522** is electrically coupled to the first DQS pin **512**, so the termination resistance for both the first DQS pin **512** and the second DQS pin **522** is provided by the termination resistor **552** internal to the first memory device **510**.

The termination resistor **556** of the DQ pin **518** of the first memory device **510** is enabled or disabled by the ODT signal received by the first ODT signal pin **514** of the first memory device **510** from the ODT bus **505**. The termination resistance of the DQ pin **528** of the second memory device **520** is enabled or disabled by the ODT signal received by the third ODT signal pin **534** of the termination assembly **530** which is external to the second memory device **520**. Thus, in certain embodiments, the first ODT signal pin **514** and the third ODT signal pin **534** receive the same ODT signal from the ODT bus **505**, and the termination resistances for both the first memory device **510** and the second memory device **520** are selectively enabled or disabled in response thereto when these memory devices are concurrently enabled. In this way, certain embodiments of the memory module **500** schematically illustrated by FIG. **14** provides external or off-chip termination of the second memory device **520**.

Certain embodiments of the memory module **500** schematically illustrated by FIG. **14** advantageously allow the use of two lower-cost readily-available 512-Mb DDR-2 SDRAM devices to provide the capabilities of a more expensive 1-GB DDR-2 SDRAM device. Certain such embodiments advantageously reduce the total cost of the resultant memory module **500**.

Certain embodiments described herein advantageously increase the memory capacity or memory density per memory slot or socket on the system board of the computer system. Certain embodiments advantageously allow for higher memory capacity in systems with limited memory slots. Certain embodiments advantageously allow for flexibility in system board design by allowing the memory module **10** to be used with computer systems designed for different numbers of ranks (e.g., either with computer systems designed for two-rank memory modules or with com-

puter systems designed for four-rank memory modules). Certain embodiments advantageously provide lower costs of board designs.

In certain embodiments, the memory density of a memory module is advantageously doubled by providing twice as many memory devices as would otherwise be provided. For example, pairs of lower-density memory devices can be substituted for individual higher-density memory devices to reduce costs or to increase performance. As another example, twice the number of memory devices can be used to produce a higher-density memory configuration of the memory module. Each of these examples can be limited by the number of chip select signals which are available from the memory controller or by the size of the memory devices. Certain embodiments described herein advantageously provide a logic mechanism to overcome such limitations.

Various embodiments of the present invention have been described above. Although this invention has been described with reference to these specific embodiments, the descriptions are intended to be illustrative of the invention and are not intended to be limiting. Various modifications and applications may occur to those skilled in the art without departing from the true spirit and scope of the invention.

What is claimed is:

1. A memory module connectable to a computer system, the memory module comprising:

a printed circuit board;

a plurality of memory devices coupled to the printed circuit board, the plurality of memory devices having a first number of memory devices; and

a logic element coupled to the printed circuit board, the logic element receiving a set of input control signals from the computer system, the set of input control signals corresponding to a second number of memory devices smaller than the first number of memory devices, the logic element generating a set of output control signals in response to the set of input control signals, the set of output control signals corresponding to the first number of memory devices, wherein the plurality of memory devices are arranged in a first number of ranks, and the set of input control signals corresponds to a second number of ranks of memory modules, the second number of ranks less than the first number of ranks, wherein the logic element further responds to a first command signal from the computer system by generating a second command signal transmitted to the plurality of memory devices, the first command signal corresponding to the second number of ranks and the second command signal corresponding to the first number of ranks.

2. The memory module of claim **1**, wherein the first command signal is a refresh signal or a precharge signal.

3. The memory module of claim **1**, wherein the memory devices comprise dynamic random-access memory (DRAM) devices.

4. The memory module of claim **1**, wherein the set of input control signals comprises a first number of chip-select signals and wherein the set of output control signals comprises a second number of chip-select signals, wherein the first number of chip-select signals is less than the second number of chip-select signals, the memory module simulating a virtual memory module having the second number of memory devices.

5. The memory module of claim **1**, wherein the logic element comprises an application-specific integrated circuit.

6. The memory module of claim **1**, wherein the logic element comprises a field-programmable gate array.

7. The memory module of claim **1**, wherein the logic element comprises a custom-designed semiconductor device.

8. The memory module of claim **1**, wherein the logic element comprises a complex programmable-logic device.

9. The memory module of claim **1**, wherein the first number of ranks is four, and the second number of ranks is two.

10. The memory module of claim **1**, wherein the first number of ranks is two, and the second number of ranks is one.

11. The memory module of claim **1**, wherein the set of input control signals comprises two chip-select signals and an address signal and the set of output control signals comprises four chip-select signals.

12. The memory module of claim **1**, wherein the printed circuit board is mountable in a module slot of the computer system, the printed circuit board having a plurality of edge connections electrically coupled to corresponding contacts of the module slot.

13. The memory module of claim **1**, wherein the plurality of memory devices are arranged to provide a first memory density per rank, and the set of output control signals corresponds to a second memory density per rank, the second memory density greater than the first memory density per rank.

* * * * *

Exhibit B

AMENDED EXHIBIT A TO JOINT CLAIM CONSTRUCTION AND PREHEARING STATEMENT

List of Constructions To Which Both Parties Have Agreed (U.S. Pat. No. 7,289,386)

Claim	Term	Agreed-Upon Construction
1	memory devices	devices in which data is stored and retrieved
1	coupled to the printed circuit board	electrically connected to the printed circuit board
1	rank	a group of memory devices enabled to receive and transmit data by a common chip-select signal
1	command signal	a signal that initiates a predetermined type of computer operation, such as read, write, refresh, or precharge
1	number of ranks of memory modules	number of ranks of memory devices
11	chip-select signal	a control signal that enables the input and output of data to and/or from a memory device

50657894.doc

Exhibit F

JEDEC STANDARD

FBDIMM
Advanced Memory Buffer (AMB)

JESD82-20A

(Revision of JESD82-20, March 2007)

SPECIAL DISCLAIMER: JEDEC has received information that certain patents or patent applications may be relevant to this standard, and, as of the publication date of this standard, no statements regarding an assurance or refusal to license such patents or patent applications have been provided.

http://www.jedec.org/download/search/FBDIMM/Patents.xls

JEDEC does not make any determination as to the validity or relevancy of such patents or patent applications. Prospective users of the standard should act accordingly.

MARCH 2009

JEDEC SOLID STATE TECHNOLOGY ASSOCIATION

13 AMB Quad Rank Support

13.1 Background

FB-DIMM supports 2 ranks per DIMM in the Architecture and Protocol spec, the AMB implementations, and the Host Controller designs.

This specification describes the implementation to achieve 4 ranks per DIMM. This is accomplished by making the DIMM appear to have the next higher density DRAMs installed, such that 2 ranks appear to the host controller as 1 rank. Quad Rank Mode A is also included in the AMB, allowing the host controller to utilize that mode of operation. The AMB must support both modes. The difference is in how the Rank Select 1 is determined.

13.2 AMB Signal Changes and DRAM Connections

13.2.1 Quad Rank Signal Requirements

Additional control signals (chip select and ODT) are required for quad rank operation. The Quad rank requirements are as follows:

Table 51 — Function Mapping Legend

Signal	Existing (SR & DR)	Quad Rank Requirement
Chip Select (CS#)	Two chip selects provided, with A and B copies.	Chip Selects are independent per rank, requiring 4 independently controlled Chip Selects. There are not separate A and B signals provided. One signal from the AMB drives both sides, and is terminated on each end.
ODT	One ODT provided, with A & B copies. Both ranks controlled by the one ODT signal.	The original ODT with A and B copies are maintained. Optionally, two more ODT signals are provided that are used for two new ranks. ODT usage is raw card dependent.
CKE	Two CKE signals provided, with A and B copies. AMB (and host) has individual control of CKE on each rank.	For raw cards that require 4 ODT signals, two CKE signals provided, but only the "A" copies, which drive both sides. For raw cards that require only 2 ODT signals, both the A and B copies of the CKE signals are available. CKE0 drives ranks 0 & 2, while CKE1 drives ranks 1 & 3. The AMB (and host) can control the CKE for rank pairs.
DRAM Address	A[15:0] provided with A and B copies.	A15 is not required, as 4Gbit devices are not supported by this quad rank spec. They are reused as ECCA2 and ECCA6, which are separately controlled copies of A2 and A6 sent only to the ECC DRAMs. Normally they are exact copies of A2 and A6, but during MRS commands, can be controlled separately by the AMB via the DCALADDR register.

13.2 AMB Signal Changes and DRAM Connections (cont'd)

13.2.2 Address and Control Signal Reuse

The existing AMB address and control signals are re-used for quad rank by:

- Reassigning an address bit not required by 2Gbit and below DRAMs
- Providing only one copy of chip selects, rather than an "A" and "B" copy.
- Optionally providing only one copy of the CKE signals, rather than an "A" and "B" copy.

13.2.2.1 Reassigning Address Bits

An AMB provides A0 through A15, and BA0 through BA2 signals to the DRAMs, but lower density DRAMs do not require all of these address signals. Specifically, A15 is not used until 4Gbit DRAMs which are not likely to be available in DDR2, and are not supported by this quad rank spec.

13.2.2.2 One Copy of Chip Select and (optionally) CKE Signals

The AMB has two copies of all DDR interface address, command, and control signals, one for the DRAMs on the left side of the DIMM and one for the DRAMs on the right side of the DIMM. These are labeled "A" and "B" at the end of the signal names. Each signal is terminated to Vtt via a resistor at the physical end of the DIMM.

For quad rank x8 DIMMs, the chip select signals are routed to only 9 DRAMs each, allowing one chip select signal to drive both sides for each rank. A dual termination scheme is used, providing a termination resistor on each end of the DIMM. Quad rank x4 would route each chip select and to 18 DRAMs each.

Optionally, the CKE0A and CKE1A signals also drive both the left and right sides for quad rank, with a dual termination scheme.

13.2.2.3 ODT Requirements of Quad Rank

There are two options for the ODT signals for quad rank operation. Option 1 provides two additional ODT signals that are used for ranks 2 and 3. Option 2 uses only the two original ODT signals.

In both cases the ODT for the rank 0 and 1 DRAMs are identical for quad rank as for single and dual rank DIMMs with 2 physical copies provided (ODT0A and ODT0B) to minimize the loading and/or simplify the routing.

In option 1, ranks 2 and 3 are controlled by the two new ODT signals, ODT1 and ODT2 respectively. This scheme requires a total of 4 ODT pins.

The ODT functionality is programmable, and it will depend on particular DIMM implementation.

Note that the functionality of original ODTA/B signals for rank 0 and rank 1 is completely unchanged, including the register programming.

13.2.2.4 ECCA2 and ECCA6

The signal routing topology of the ECC DRAMs may require different strength of ODT on the ECC DRAMs from the data DRAMs. The ODT strength is programmed through MRS commands, which will program all DRAMs of a rank with the same value via the DRAM address bus. The ODT strength is programmed via EMRS1 bits A6 and A2.

To allow the ECC DRAMs to have their ODT strength programmed separately, the AMB creates a separate A2 and A6 for the ECC DRAMs. These signals normally follow the A2 and A6 functionality, but on all MRS cycles, AMB can send out separate values.

ECCA2 and ECCA6 are connected to all 4 rank's ECC bits. ECCA2 and ECCA6 are non-inverted outputs when address inversion is enabled.

13.2 AMB Signal Changes and DRAM Connections (cont'd)

13.2.3 Quad Rank Signal Mapping

The following table shows the mapping of the chip select, ODT, and CKE signals to each rank. The Signal name is the logical signal. The mapping to physical AMB balls is listed in a follow on table.

Table 52 — Quad Rank Signal mapping per rank, ODT Option 1

Signal	Rank0	Rank1	Rank2	Rank3
CS0#	CS			
CS1#		CS		
CS2#			CS	
CS3#				CS
ODTA		ODT		
ODTB	ODT			
ODT1			ODT	
ODT2				ODT
CKE0	CKE		CKE	
CKE1		CKE		CKE

Table 53 — Quad Rank Signal mapping per rank, ODT Option 2

Signal	Rank0	Rank1	Rank2	Rank3
CS0#	CS			
CS1#		CS		
CS2#			CS	
CS3#				CS
ODT0A/B[a]	ODT	ODT		
VSS			ODT	ODT
CKE0A/B[1]	CKE		CKE	
CKE1A/B[1]		CKE		CKE

NOTES:
a.For ODT Option 2, CKE and ODT A and B copies are routed to the left and right sides of the DIMM in the same manner as dual rank DIMMs

13.2 AMB Signal Changes and DRAM Connections (cont'd)

13.2.4 Mapping to AMB Balls

The following table shows the mapping between the existing AMB balls and the quad rank signals.

Table 54 — AMB Pin usage for each DIMM type

AMB Ball	Non-QR (SR or DR)	Quad Rank ODT Opt 1	Quad Rank ODT Opt 2	Quad Rank Comment
CS0#A	CS0#A	CS0#[a]	CS0#[1]	Rank 0
CS0#B	CS0#B	CS2#[1]	CS2#[1]	Rank 1
CS1#A	CS1#A	CS1#[1]	CS1#[1]	Rank 2
CS1#B	CS1#B	CS3#[1]	CS3#[1]	Rank 3
ODTA	ODTA	ODTA[b]	ODTA[2]	First ODT
ODTB	ODTB	ODTB[2]	ODTB[2]	Second ODT
A15A	A15A	ECCA2	ECCA2	A2 for the ECC DRAMs
A15B	A15B	ECCA6	ECCA6	A6 for the ECC DRAMs
CKE0A	CKE0A	CKE0	CKE0A	Ranks 0&2
CKE0B	CKE0B	ODT1	CKE0B	Optional third ODT, or CKE0B for Ranks 0&2
CKE1A	CKE1A	CKE1	CKE1A	Ranks 1&3
CKE1B	CKE1B	ODT2	CKE1B	Optional third ODT, or CKE1B for Ranks 1&3

NOTES:
a. There are no separate A & B copies for the CS# signals in Quad Rank
b. Logically identical

13.2.5 Signal State at Reset

CKE and ODT signals must be low during the DRAM power ramp. Since CKE signals are muxed with ODT signals, there is no power up difference. On a quad rank DIMM, the BIOS should set the register bit to configure the CKE0B & CKE1B signals to the ODT functionality (if applicable) prior to taking CKE high.

13.3 Rank Decode

FBD Quad Rank allows two addressing modes for the additional ranks. The AMB must implement both modes. The host may choose which mode to implement.

Mode C is specific to FBD and uses a bank or row address bit as the additional rank select bit. This mode allows for the full 8 DIMMs per channel. This mode appears to the host controller as a dual rank DIMM.

Mode A uses DS2 as the additional rank select bit. It limits the number of DIMMs per channel to 4 when quad rank is being used.

13.3 Rank Decode (cont'd)

13.3.1 Quad Rank Mode C

Quad rank mode C operates by making two physical ranks show up as one logical rank. A bank address bit as well as the existing RS bit determines which rank the AMB will access. The host controller will see the DIMM as dual rank, with each rank being one density higher than the physical DRAMs.

The host controller sends one Rank Select bit with the DRAM commands that are rank specific. This is called RS in the FBD spec, and will be called RS0 here. The AMB must create an RS1 bit internally to determine which of the 4 ranks to access. RS1 will be used to divide each logical rank into 2 physical ranks on the DIMM. Logical Rank 0 will be divided into physical ranks 0 and 2, while Logical Rank 1 will be divided into physical ranks 1 and 3.

Table 55 — RS1:0 to rank decode

RS1 (created inside the AMB from BA0)	RS0 (RS field in FBD commands)	Rank Accessed
0	0	0
0	1	1
1	0	2
1	1	3

In the AMB the BA0 bit from the host controller is used to create RS1. This places the even numbered banks in one physical rank, and the odd numbered banks in the other physical rank.

The AMB must recreate the BA0 signal to the DRAMs (called dBA0 here), which is done differently depending on the DRAM density.

Table 56 — dBA0 (DRAM BA0) Selection by DRAM density

DRAM Density	dBA0 (DRAM BA0) bit:
512Mbit	Host BA2
1Gbit	Host Row A14
2Gbit	Host Row A15

The following table shows the physical layout of each DRAM type, and how the host controller will view the DIMM.

Table 57 — DIMM addressing

Physical x8 DRAMs on DIMM					Host Controller View				
Rank	Density	Banks	Row Bits	Col Bits	Rank	Density	Banks	Row Bits	Col Bits
4	512Mbit	4	14	10	2	1Gbit	***8***	14	10
4	1Gbit	8	14	10	2	2Gbit	8	***15***	10
4	2Gbit*	8	15	10	2	4Gbit	8	***16***	10

* Some QR DIMMs may not support 2Gbit DRAMs

13.3 Rank Decode (cont'd)

13.3.1 Quad Rank Mode C (cont'd)

512Mbit QR DIMMs appear to the host controller logically as Dual Rank 1Gbit DIMMs. The difference between 512Mbit parts and 1Gbit parts is the addition of BA2 (Bank Address 2) for both row and column commands. The AMB will use BA2 in the FBD command as dBA0 to the DRAMs. Note that BA2 is sent on both row and column commands to the DRAMs so this is a simple mapping.

1Gbit QR DIMMs appear to the host controller logically as Dual Rank 2Gbit DIMMs. The difference between a 1Gbit part and 2Gbit parts is the addition of A14 for the Activate Command. The AMB will use Row address A14 in the FBD command as dBA0 to the DRAMs.

2Gbit QR DIMMs appear to the host controller logically as Dual Rank 4Gbit DIMMs. The difference between a 2Gbit part and 4Gbit parts is the addition of A15 for the Activate Command. The AMB will use Row address A15 in the FBD command as dBA0 to the DRAMs.

For 512Mbit parts, the dBA0 bit is created from BA2, which is sent for all commands which are rank specific, so this is a simple mapping. For 1Gbit and 2Gbit DRAMs, a Row Address is used to create the DRAM dBA0. This bit is ONLY sent during activate commands. This requires the AMB to remember the dBA0 bit from the last activate and use it for subsequent Read, Write, and Precharge Single commands. The DRAM's BA0 essentially becomes a row address in Quad Rank Mode C.

The AMB must store the dBA0 bit separately for each combination of BA2:0 and RS0 from the host controller. This is required because the host could have sent a different dBA0 for each possible open bank. This requires 16 dBA0s to be stored in the AMB, corresponding to the 16 banks that could be open at one time by the host controller. The dBA0 is stored when it is sent with the Activate Command, and used for any subsequent read, write, and precharge command to the same bank of the same logical rank.

On each activate, the AMB stores the value being sent on dBA0 in one of 16 locations selected by the host RS bit and BA[2:0] bits. This represents each different bank that the host controller could consider opened.

One each read, write, or precharge single command the AMB selects the proper stored value by using the RS bit and BA[2:0] bits send by the host in the read, write, or precharge command.

Table 58 — Activate Command mapping

Host (FBD channel) bit	512Mbit DRAM	1Gbit DRAM	2Gbit DRAM	All other modes
BA0	(RS1)	(RS1)	(RS1)	BA0
BA1	BA1	BA1	BA1	BA1
BA2	BA0	BA2	BA2	BA2
A[13:0]	A[13:0]	A[13:0]	A[13:0]	A[13:0]
A14	(A14)	BA0	A14	A14
A15	(not used)	(not used)	BA0	A15
RS	selects ranks	selects ranks	selects ranks	selects ranks

13.3 Rank Decode (cont'd)

13.3.1 Quad Rank Mode C (cont'd)

Table 59 — Read, Write & Precharge Single Command mapping

Host (FBD channel) bit	512Mbit DRAM	1Gbit DRAM	2Gbit DRAM	All other modes
BA0	(RS1)	(RS1)	(RS1)	BA0
BA1	BA1	BA1	BA1	BA1
BA2	BA0	BA2	BA2	BA2
A[13:0]	A[13:0]	A[13:0]	A[13:0]	A[13:0]
RS	selects ranks	selects ranks	selects ranks	selects ranks
Stored BA0 bit from the activate	(not used)	BA0	BA0	(not used)

Table 60 — Rank Selection and dBA0 generation for each command type

Command	Rank Decoding	DRAM BA0
Activate	Uses Host BA0	512Mbit: Host BA2 used 1Gbit: Host A14 used 2Gbit: Host A15 used
Read	Uses Host BA0	512Mbit: Host BA2 used For 1Gbit and 2Gbit parts this is a stored value from the last activate command to the same bank and rank. Which stored bit is determined by the host BA[2:0] & RS bits. Which stored bit is used is determined by the host BA[2:0] & RS bits
Write	Uses Host BA0	512Mbit: Host BA2 used For 1Gbit and 2Gbit parts this is a stored value from the last activate command to the same bank and rank. Which stored bit is used is determined by the host BA[2:0] & RS bits.
Precharge Single	Uses Host BA0	512Mbit: Host BA2 used For 1Gbit and 2Gbit parts this is a stored value from the last activate command to the same bank and rank. Which stored bit is used is determined by the host BA[2:0] & RS bits.
Auto Refresh	Command is sent to two ranks. RS bit determines whether it is sent to Ranks 0 & 2 or Ranks 1 & 3.	N/A
Precharge All	Command is sent to two ranks. RS bit determines whether it is sent to Ranks 0 & 2 or Ranks 1 & 3.	N/A

Table 60 — Rank Selection and dBA0 generation for each command type

Command	Rank Decoding	DRAM BA0
Enter Self Refresh	Command is sent to two ranks. RS bit determines whether it is sent to Ranks 0 & 2 or Ranks 1 & 3.	N/A
Exit Self Refresh / power down	Command is sent to two ranks. RS bit determines whether it is sent to Ranks 0 & 2 or Ranks 1 & 3. Note that this command only involves the CKE signals which are physically shared by Ranks 0 & 2, and by Ranks 1 & 3.	N/A
Enter Power Down	Command is sent to two ranks. RS bit determines whether it is sent to Ranks 0 and 2 or Ranks 1 and 3. Note that this command only involves the CKE signals which are physically shared by Ranks 0 & 2, and by Ranks 1 & 3	N/A
CKE per DIMM	Sent to all ranks on the DIMM	N/A
CKE per rank	Command is sent to two ranks. The appropriate RS bit determines whether it is sent to Ranks 0 and 2 or Ranks 1 and 3. Note that this command only involves the CKE signals which are physically shared by Ranks 0 & 2, and by Ranks 1 & 3.	N/A

13.3.1.1 Host responsibilities for read timing

The host controller is responsible for assuring that there is a turnaround cycle between reads to the same DIMM but different ranks. This includes switching between physical ranks that appear as one logical rank in the Mode C addressing mode.

13.3.2 FBD2 Mode A Rank Decode

Mode A uses DS2 (DIMM Select 2) as the additional rank decode. This limits the total number of DIMMs to 4 per channel, as decoded by DS[1:0]. For Mode A, the host controller knows about the 4 individual ranks, and addresses them directly.

The decode for DS[2:0] within the AMB normally comes from the strapping signals SA[2:0], which provides the physical DIMM number, setting the address for the SPD and AMB. This DIMM number is compared to the DS[2:0] in the command frames. In Quad Rank Mode A, a quad rank DIMM decode ignores the DS2 bit, using it as Rank Select 1. By the same token, in Quad Rank Mode A, a quad rank DIMM decode will also ignore the WS2 bit which was used by the AMB on the DIMM to determine if it should write the Wdata into its write FIFO. Single and dual rank DIMMs will continue to decode DS2 as a DIMM select bit.

Table 61 — Quad Rank Mode A DS[2:0] and RS mapping

DIMM Type	DS2 (RS1)	DS1	DS0	RS
Quad Rank DIMM	Rank Select 1	DIMM Decode	DIMM Decode	Rank Select 0
Dual Rank DIMM	DIMM Decode	DIMM Decode	DIMM Decode	Rank Select
Single Rank DIMM	DIMM Decode	DIMM Decode	DIMM Decode	Not Used

13.3 Rank Decode (cont'd)

13.3.2 FBD2 Mode A Rank Decode (cont'd)

In Mode A, the AMB need not store the rank information for each bank, since the rank select is sent with all DRAM commands for all densities.

Operations other than DRAM commands operate in the same manner as a single or dual rank DIMM.

A quad rank DIMM responds to configuration read and write cycles at its selected DIMM number, decoding DS[2:0]. It does not respond to configuration read and write cycles at the alternate address with DS2=1.

A quad rank DIMM sends status frames providing data only on the lane selected by SA[2:0]. It does NOT repeat the status information on the additional lane selected when DS2=1.

The following example is for a quad rank DIMM with the strapping of SA[2:0] = 001 (DIMM 1).

Table 62 — AMB response to commands

Command	**AMB response, when SA[2:0]=001**
Activate Write Read Precharge All Precharge Single Auto Refresh	Responds to DS[2:0]=001 (Ranks 0 and 1) Responds to DS[2:0]=101 (Ranks 2 and 3) Commands are sent to a single rank.
Enter Self Refresh Exit Self Refresh/Exit Power Down Enter Power Down	Responds to DS2:0]=x01 Since CKE is shared between ranks 0&2 and ranks 1&3, these commands target 2 ranks at a time, based on the RS bit. RS=0 targets command to ranks 0 & 2. RS=1 targets command to ranks 1 & 3.
Channel commands Config read and write CKE per DIMM CKE per Rank	Responds to DS[2:0]=001 only
Status Packet	Drives data onto logical lane 1 only

13.4 ODT timing on reads

The read timing for ODT is basically the same as the write timing. The ODT is enabled for the DQS preamble, and remains active for 3.5 clocks for BL=4 and 5.5 clocks for BL=8. The equation for the DRAM ODT signals is:

Turn on: RL - 3 clocks from the read command

Turn off: RL - 2 + (BL/2) clocks from the read command

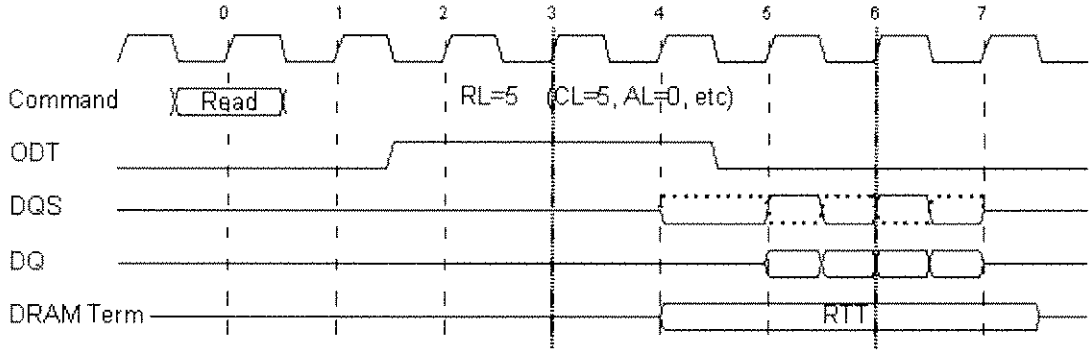

Figure 26 — DRAM ODT TIming during Reads

13.4.1 AMB data bus termination

The AMB data bus termination remains at 65 ohms nominal.

13.5 Registers

DDR2ODTC: DDR2 DRAM On-Die Termination Control

NOTE This is an existing register, in which previously only 8 bits were used. In reality, since the AMB only has one ODT pin set for ranks 0 and rank 1, only 4 of the 8 bits have any function. These original 8 bits are not changed.

For the quad rank option using 4 ODT signals, 8 more bits are added for control of the ODT1 and ODT2 signals for reads. There is one bit for each rank reads.

Device: NodeID **Function: 4** **Offset: FCh**			
Bit	**Attr**	**Default**	**Description**
15	RWST	0	ODT2 control during reads from CS3
14	RWST	0	ODT1 control during reads from CS3
13	RWST	0	ODT2 control during reads from CS2
12	RWST	0	ODT1 control during reads from CS2
11	RWST	0	ODT2 control during reads from CS1
10	RWST	0	ODT1 control during reads from CS1
9	RWST	0	ODT2 control during reads from CS0
8	RWST	0	ODT1 control during reads from CS0 1: ODT pin will drive high 0: ODT pin will drive low
Note that bits 7:0 below do not change from the current spec			
7:6	RWST	0h	R1ODTWR: ODT control during writes to CS1 & CS3 x1: ODT0A/B, ODT1, & ODT2 pins will drive high x0: ODT0A/B, ODT1, & ODT2 pins will drive low Note: Must always be set to 01 for SR and DR DIMMs
5:4	RWST	0h	R1ODTRD: ODT0 control during reads to CS1 & CS3 x1: ODT0A/B pins will drive high x0: ODT0A/B pins will drive low Note: Must always be set to 00
3:2	RWST	0h	R0ODTWR: ODT control during writes to CS0 & CS2 x1: ODT0A/B, ODT1, & ODT2 pins will drive high x0: ODT0A/B, ODT1, & ODT2 pins will drive low Note: Must always be set to 01 for SR and DR DIMMs
1:0	RWST	0h	R0ODTRD: ODT0 control during reads to CS0 & CS2 x1: ODT0A/B pins will drive high x0: ODT0A/B pins will drive low Note: Must always be set to 00

13.5 Registers (cont'd)

Quad Rank control register

Device: NodeID **Function: 3** **Offset: 88h**			
Bit	**Attr**	**Default**	**Description**
5	RWST	0	A15A and A15B pin muxing 0 = A15A and A15B functionality 1 = ECCA2 and ECCA6 functionality. Only permitted when the Quad Rank Enable bit is set to a 1.
4	RWST	0	CKE0B and CKE1B pin muxing 0 = CKE0B and CKE1B functionality 1 = ODT1 and ODT2 functionality. Only permitted when the Quad Rank Enable bit is set to a 1.
3:2	RWST	00	RS1 Select. This field determines what the AMB uses for the dBA0 (DRAM BA0) 00 = BA2 used for dBA0 01 = Row Address 14 used for dBA0 10 = Row Address 15 used for dBA0 11 = DS2 (Quad Rank Mode A)
1			Reserved
0	RWST	0	Quad Rank Enable 0 = normal single or dual rank operation 1 = Quad Rank DIMM

DCALCSR

Bits 23 and 24 added to support the additional ranks.

Device: NodeID **Function: 4** **Offset: 40h**			
Bit	**Attr**	**Default**	**Description**
24:21	RW	0000	This field corresponds to the chip select outputs: CS[3:0]. Setting a bit in this field will cause the corresponding CS pin to drive low when commands are issued on the DDR bus. This field Applies to NOP, Refresh, Precharge all, and MRS/ EMRS commands. Bit 21 is for CS0, Bit 22 if for CS1, Bit 23 is for CS2, and bit 24 is for CS3.

13.5 Registers (cont'd)

DCALADDR: DCAL Address Register

Bits 5:4 are added to XOR with ECCA2 and ECCA6

Device: NodeID **Function: 4** **Offset: 44h**	
Bit	**Description**
31:16	DRAM Address Bus 15:0
15:6	Reserved
5	XOR for ECCA6. This bit is set to a 1 to have the ECC A6 bit the opposite of A6. Only used if quad rank and A15A/B pin muxing enabled (function 3 offset 0x88 bits 0 and 5), otherwise ignored.
4	XOR for ECCA2. This bit is set to a 1 to have the ECC A2 bit the opposite of A2. Only used if quad rank and A15A/B pin muxing enabled (function 3 offset 0x88 bits 0 and 5), otherwise ignored.
3	Reserved
2:0	DRAM Bank Address bus 2:0

Shaded rows are unchanged.

A6/A2 and ECCA6/ECCA2 functionality of DCALADDR if Quad rank mode and A15A/B pin muxing enabled (function 3 offset 0x88 bits 0 and 5).

Signal	**DCALADDR bits**
A6A, A6B	Bit 22 (existing definition)
ECCA6	Bit 22 XOR Bit 5
A2A, A2B	Bit 18 (existing definition)
ECCA2	Bit 18 XOR Bit 4

MEMBIST register change

13.5 Registers (cont'd)

Only the two bits within the register that change are shown. Changes are shown in bold.

Device: NodeID **Function: 3** **Offset: 40h**			
Bit	**Attr**	**Default**	**Description**
21:20	RW	00	CS: CS[3:0] selection in MemBIST mode 01: Select Rank 0 10: Select Rank 1 00: Select Rank 2 11: Select Rank 3

NOTE The encoding of this register is based on the existing definition.

13.6 Fast Reset

During a fast reset, the AMB performs a set of operations to close all pages and put the DRAMs into self refresh. The AMB will access two ranks in parallel for all commands during this time.

NOTE For QRx8 this will be activating the same number of DRAMs at a time as is already done for DRx4 DIMMs.

For Quad Rank FB-DIMM MEMBIST and Transparent Mode support, please refer to JESD82-28A "Fully Buffered DIMM Design for Test, Design for Validation (DFx)".


JEDEC

Exhibit G

JEDEC STANDARD

FBDIMM: Architecture and Protocol

JESD206

SPECIAL DISCLAIMER: JEDEC has received information that certain patents or patent applications may be relevant to this standard, and, as of the publication date of this standard, no statements regarding an assurance or refusal to license such patents or patent applications have been provided.

http://www.jedec.org/download/search/FBDIMM/Patents.xls

JEDEC does not make any determination as to the validity or relevancy of such patents or patent applications. Prospective users of the standard should act accordingly.

JANUARY 2007

JEDEC SOLID STATE TECHNOLOGY ASSOCIATION

JEDEC

EIA Electronic Industries Alliance

Other Restrictions

Only one outstanding configuration read or write register transaction is allowed on the channel. A configuration register read begins with the command and ends with the data being returned to the host. A configuration write begins with the command and ends when the read data would have been returned if the command were a Read Config Reg. This is the same point that an Alert Frame would be generated if there were a CRC error on the Write Config Reg command. Allowing only one outstanding configuration transaction on the bus allows for proper replay of the Write Config Reg command following an Alert Frame.

A Soft Channel Reset requires NOP commands in all other command slots in the previous DRAM clock, the current DRAM clock, and the next 4 DRAM clocks.

Only one In-band Debug event may be sent within a DRAM clock.

The host controller is responsible for state and timing of the CKE pins vs. DRAM commands based on the DRAM specifications. A DRAM command and CKE command may target the same DIMM on the same DRAM clock provided that the DRAM specifications are met.

Examples:

A DRAM command may be issued to rank 1 on the same DRAM clock as a DRAM CKE per Rank command that changes the CKE of rank 0 while retaining a 1 on the CKE of rank 1.

A DRAM command may be issued to DIMM 2 on the same DRAM clock as a DRAM CKE Command per DIMM that targets all DIMMs, but retains the state the CKEs of DIMM 2 as 1.

4.2.3 Command Encoding

Commands are encoded into the 24 bit C[23:0] fields of Command frames. Table 4-39 defines the bit mapping of an example DRAM configuration and the channel commands into the C[23:0] field.

Table 4-39 — Command Encoding

DRAM Cmnds	23	22	21	20	f9	f8	17	16	15	14	f3	12	11	10	9	8	7	6	5	4	3	2	1	0
Activate	DS2	DS1	DS0	1	DRAM Addr		RS	DRAM Bank & Address																
Write	DS2	DS1	DS0	0	1	f	RS	DRAM Bank & Address																
Read	DS2	DS1	DS0	0	1	0	RS	DRAM Bank & Address																
Precharge All	DS2	DS1	DS0	0	0	1	RS	X	X	X	X	f	f	1	X	X	X	X	X	X	X	X	X	X
Precharge Single	DS2	DS1	DS0	0	0	1	RS	DRAM Bank				1	f	0	X	X	X	X	X	X	X	X	X	X
Auto (CBR) Refresh	DS2	DS1	DS0	0	0	1	RS	X	X	X	X	f	0	1	X	X	X	X	X	X	X	X	X	X
Enter Self Refresh	DS2	DS1	DS0	0	0	1	RS	X	X	X	X	1	0	0	X	X	X	X	X	X	X	X	X	X
Exit Self Refresh / Exit Power Down	DS2	DS1	DS0	0	0	1	RS	X	X	X	X	0	1	1	X	X	X	X	X	X	X	X	X	X
Enter Power Down	DS2	DS1	DS0	0	0	f	RS	X	X	X	X	0	1	0	X	X	X	X	X	X	X	X	X	X
reserved	X	X	X	0	0	1	X	X	X	X	X	0	0	X	X	X	X	X	X	X	X	X	X	X

Note: The values in "X" fields in non-reserved commands above may be driven onto the DRAM device pins

Channel Cmnds							OP3	OP2	OP1	OP0														
Debug: In-band Events	EV7	EV6	EV5	0	0	0	1	1	f	f	1	EV4	EV3	EV2	EV1	EV0	PV7	PV6	PV5	PV4	PV3	PV2	PV1	PV0
Debug: Relative Timing	PH5	PH4	PH3	0	0	0	1	1	1	1	0	PH2	PH1	PH0	RT9	RT8	RT7	RT6	RT5	RT4	RT3	RT2	RT1	RT0
Debug: Exposed Info	EX16	EX15	EX14	0	0	0	1	1	1	0	EX13	EX12	EX11	EX10	EX9	EX8	EX7	EX6	EX5	EX4	EX3	EX2	EX1	EX0
reserved	X	X	X	0	0	0	1	1	0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
reserved	X	X	X	0	0	0	1	0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
DRAM CKE per DIMM	DS2	DS1	DS0	0	0	0	0	1	1	1	BCST	X	X	X	X	X	DE7	DE6	DE5	DE4	DE3	DE2	DE1	DE0
DRAM CKE per Rank	DS2	DS1	DS0	0	0	0	0	1	1	0	BCST	X	X	X	X	X	D3 R1	D3 R0	D2 R1	D2 R0	D1 R1	D1 R0	D0 R1	D0 R0
Write Config Reg	DS2	DS1	DS0	0	0	0	0	f	0	f	DS3	TID	X	A10	A9	A8	A7	A6	A5	A4	A3	A2	0	0
Read Config Reg	DS2	DS1	DS0	0	0	0	0	1	0	0	DS3	X	X	A10	A9	A8	A7	A6	A5	A4	A3	A2	0	0
reserved	X	X	X	0	0	0	0	0	1	1	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Soft Channel Reset	X	X	X	0	0	0	0	0	f	0	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Sync	X	X	X	0	0	0	0	0	0	1	X	SD1	SD0	X	X	X	X	IER	ERC	ELOs	X	X	R1	R0
Channel NOP	X	X	X	0	0	0	0	0	0	0	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Note: All unused encodings are reserved. "X" values should be driven by the host to zero and ignored by the AMB

4.2.4 DRAM Commands

DRAM commands are generated by the host to access the DRAM devices behind each AMB buffer. The host has access to the DRAM devices as if the devices were directly connected to the host. The DS[2:0] field directs the command to one of the eight possible DRAM DIMMs on the FBD channel. The AMB decodes the DRAM commands and generates the control signals to the DRAM devices. The command delivery on the DRAM address and control signals (excluding CKE) use 1n command timing. 1n command timing means that the commands are present on the DRAM pins for a single clock cycle. The exact mapping of the control signals delivered to the DRAM devices are defined in the *FBD AMB Specification.*

AMB buffers may support more than one DRAM technology. The details of the mapping of bank and address bits from the commands to the DRAM devices are specified in the *FBD AMB Specification.* An example mapping is shown in Table 4-40. For complete details of the DRAM command encoding refer to the JEDEC SDRAM data sheets.

In the following table, the RS (Rank Select) bit specifies to the AMB which memory ranks located behind the buffer should be accessed. The other labels correspond to the familiar labels in the SDRAM data sheets. Rows labeled with an "*" are speculative and may change as the JEDEC SDRAM data sheets mature.

Bit position 10 is used in the command encoding of the Precharge Single and Precharge All commands to allow the command bit to be mapped directly onto the DRAM address bit 10 to match the DRAM AP bit usage.

Table 4-40 — DRAM Command Mapping Examples

DDR2 Config		20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
256Mb (64Mbx4)	Row	1	X	X	RS	X	X	B1	B0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1KB page	Col	0	1	r/w	RS	X	X	B1	B0	X	A11	AP	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
512Mb (128Mbx4)	Row	1	X	X	RS	A13	X	B1	B0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1KB page	Col	0	1	r/w	RS	X	X	B1	B0	X	A11	AP	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1Gb (256Mbx4)	Row	1	X	X	RS	A13	B2	B1	B0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1KB page	Col	0	1	r/w	RS	X	B2	B1	B0	X	A11	AP	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
2Gb (512Mbx4)	Row	1	X	A14	RS	A13	B2	B1	B0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1KB page	Col	0	1	r/w	RS	X	B2	B1	B0	X	A11	AP	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
4Gb (1Gbx4) *	Row	1	A15	A14	RS	A13	B2	B1	B0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1KB page	Col	0	1	r/w	RS	X	B2	B1	B0	X	A11	AP	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
4Gb (1Gbx4) *	Row	1	X	A14	RS	A13	B2	B1	B0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
2KB page	Col	0	1	r/w	RS	X	B2	B1	B0	A12	A11	AP	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0

DRAM Read and Write commands always transfer complete bursts of data determined by the Burst Length field programmed into the DRAM MRS registers. A burst length of 4 will transfer 36 bytes and a burst length of 8 will transfer 72 bytes to/from each ECC DIMM. Non-ECC memory DIMMs support the Data Mask function.

Write accesses transfer the data from the write data FIFO located inside the AMB device on the DIMM. A register instructs the AMB when to drive the data after the Write command. The DDR2 specific Off-Chip Driver (OCD) Impedance Adjust command (EMRS access with A[9:7] = 100) also transfers data from the write data FIFO to the DRAM devices.

The host is responsible for memory ordering, FBD channel scheduling, and error handling.

4.2.5 Channel Commands

Channel commands include the Sync command, miscellaneous DRAM commands, configuration register read and write commands, and miscellaneous maintenance commands. Channel commands may include a DS[2:0] field to specify which DIMM the command is addressing, a 4-bit operation code field to define the command type, and an 11-bit address field. Table 4-41 defines the encoding of the Channel commands. The individual configuration registers are defined in the *FBD AMB Specification* Register chapter.

Exhibit K

Howard G. Pollack (CA Bar No. 162897/pollack@fr.com)
Shelley K. Mack (CA Bar No. 209596/mack@fr.com)
Robert J. Kent (CA Bar No. 250905/rjkent@fr.com)
FISH & RICHARDSON P.C.
500 Arguello Street, Suite 500
Redwood City, CA 94063
Telephone: (650) 839-5070
Facsimile: (650) 839-5071

Attorneys for Plaintiff
GOOGLE INC.

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

(OAKLAND DIVISION)

GOOGLE INC., Plaintiff, v. NETLIST, INC., Defendant. AND RELATED COUNTERCLAIMS.	Case No. C 08-04144 SBA **PLAINTIFF GOOGLE INC.'S RESPONSES TO NETLIST'S INTERROGATORIES, SET NO. TWO [NOS. 6-9]**

Pursuant to Rule 33 of the Federal Rules of Civil Procedure, Plaintiff Google Inc. ("Google") hereby responds to Defendant Netlist, Inc.'s ("Netlist") Interrogatories, Set No. 2, as follows. These responses are based upon information presently available and are therefore made without prejudice to Google's right to use or rely upon subsequently discovered information. As permitted by the Federal Rules of Civil Procedure, these responses may be changed, modified, or supplemented. In responding to Netlist's Interrogatories, Google does not waive any objections on the grounds of privilege, competency, relevance, materiality, authenticity, or admissibility of the information contained in these responses. Google also expressly reserves the right to object later to the admissibility of any of this information into evidence on any permissible grounds, including grounds not identified herein.

GENERAL OBJECTIONS

Google's responses are subject to the following General Objections, which Google specifically incorporates into its responses to each of Netlist's interrogatories, whether or not such General Objection is expressly referenced. The incorporation by reference of any one of these General Objections shall not be construed to exclude the incorporation of any other General Objection. Moreover, Google does not waive its right to amend its objections.

1. Google objects to the Interrogatories and, in particular, to the Definitions and Instructions, to the extent they seek to impose obligations upon Google beyond those imposed by the Federal Rules of Civil Procedure, the Local Rules of the United States District Court for the Northern District of California, and/or the Court's scheduling Orders in this case. Google will respond to the Interrogatories only to the extent required by these Rules and Order.

2. Google objects to the Interrogatories insofar as they are vague, ambiguous, indefinite, overbroad, unduly burdensome, duplicative, cumulative, indefinite as to time or scope, unintelligible, or otherwise unclear as to the precise information sought.

3. Google objects to the Interrogatories insofar as they seek information that is neither relevant to a claim or defense of any party, nor reasonably calculated to lead to the discovery of admissible evidence.

4. Google objects to the Interrogatories to the extent that they seek documents protected by the attorney-client privilege or by the work-product doctrine, protected by any other applicable privilege or immunity, prepared in connection with settlement discussions, prepared in anticipation of adversarial proceedings such as litigation or for trial, prepared in connection with any applicable joint defense agreement, or not otherwise within the scope of permissive discovery under the Federal Rules of Civil Procedure and applicable Local Rules.

5. Google objects to the Interrogatories on the ground and to the extent they call for information that Google is under an obligation to third parties to not disclose.

6. Google objects to the Interrogatories on the ground and to the extent they seek to obtain information not in Google's possession, custody, or control.

7. Google objects to the Interrogatories on the ground and to the extent they seek information already in Netlist's possession or information that is a matter of public record or that is otherwise equally available to Netlist.

8. Google objects to the Interrogatories to the extent they seek identification of "all documents" or "all persons" on the basis that they are overbroad and unduly burdensome.

9. Google objects to the Interrogatories to the extent they call for a legal opinion or conclusion, particularly to the extent the Interrogatories include claim terms whose meanings are disputed by the parties and which have not yet been construed by the Court. Google neither expresses nor intends to express any legal opinion or conclusion by responding to Netlist's Interrogatories.

10. Google objects to the Interrogatories to the extent that they fail to specify a relevant time period for which information is requested, and/or to the extent the specified period is irrelevant.

11. Google objects to the Interrogatories to the extent they are premature under any relevant discovery and/or scheduling orders.

12. Google objects to the Interrogatories to the extent that they use terms that are not defined or understood, or are vaguely and/or ambiguously defined, and therefore fail to identify with reasonable particularity the information sought. Google will not speculate as to the meaning to ascribe to such terms.

13. Google objects to the Interrogatories to the extent they seek identification of "all persons," on the basis that this renders the interrogatories overbroad and unduly burdensome.

14. Google objects to the definition of "Google," "you," or "your" (Definition No. 2) on the grounds and to the extent that it purports to encompass non-Google entities. For purposes of responding to these interrogatories, Google will interpret the terms "Google," "you," or "your" to refer to Google Inc., including without limitation all of its corporate locations and all predecessors, subsidiaries, parents, and affiliates, and all past or present directors, officers, agents, representatives, employees, consultants, attorneys, and others acting on behalf of Google Inc.

///

15. Google objects to the definitions of "identify" and "identity" (Definition No. 7) as ambiguous, overly broad, unduly burdensome, and not reasonably calculated to lead to the discovery of admissible evidence. Google further objects to these definitions to the extent they purport to seek information not within Google's personal knowledge or control, or documents or things not within Google's possession, custody or control. Google will provide information that is presently available from a reasonable search of its own files likely to contain relevant or responsive documents and from a reasonable inquiry of its present employees.

16. Google objects to the definitions of "JEDEC Mode C" and "AMB Quad Rank Support" (Definition Nos. 17 and 20) as vague and ambiguous. Although Netlist purports to define these terms by reference to a JEDEC standard, Google further objects to their use in these Interrogatories to the extent such use is inconsistent with that standard.

17. Google objects to the definitions of "memory density multiplication" and "memory rank multiplication" (Definition Nos. 18 and 19) as over broad and unduly burdensome to the extent these terms refer to "any technology" that can be used to increase computer memory capacity, including concepts and techniques not claimed in the patent-in-suit.

18. Google objects to the definition of "memory rank decoding" (Definition No. 21) as over broad and unduly burdensome to the extent it includes "any technology," whether or not claimed by the patent-in-suit. Google further objects to this definition to the extent it includes claim terms whose meaning is disputed by the parties and which have not yet been construed by the Court.

RESPONSES TO INTERROGATORIES

INTERROGATORY NO. 6:

For each asserted claim of the '386 Patent for which Google asserts non-infringement or seeks a declaration of non-infringement, describe in detail all factual and legal bases for Google's contention that the manufacture, use, offer for sale, and/or sale of accused products does not infringe that claim, including an explanation of why Google contends that the claim is not infringed literally or under the doctrine of equivalents by such manufacture, use, offer to sell, or sale on a claim element by claim element basis (e.g., in a "claim chart"); identify all documents

tending to support each such contention; and identify the person(s) employed by Google most knowledgeable concerning the facts underlying each such contention.

RESPONSE TO INTERROGATORY NO. 6:

Google incorporates each of the foregoing General Objections as if set forth fully in response to this Interrogatory. Google further objects to this Interrogatory to the extent it calls for information protected by the attorney-client privilege, the work product doctrine, or any other applicable exemption from discovery. Google further objects to this Interrogatory as prematurely calling for a legal conclusion to the extent it would require Google to take a position on the construction of certain claims on which the Court in this case has not yet ruled. Google further objects to this Interrogatory as over broad and unduly burdensome.

Subject to and without waiving the foregoing objections, Google responds as follows: Google does not infringe any claim of the '386 patent because one or more elements required to be present by each claim is missing from Google's accused products, both literally and under the doctrine of equivalents. For example, the accused products do not include a structure that meets the "logic element" limitation because they nowhere include the functionality that is claimed in that limitation. Netlist has not offered any evidence to directly contradict this assertion; at the time of this response, Netlist has not even made any specific allegations of the supposed infringement of its patents by any Google product. Its only infringement contentions to date have consisted entirely of general allegations concerning a *proposed* industry standard.

Google has produced supporting documents and things during the course of this litigation in the form of specifications, source code, and a server inspection.

Google identifies Mr. Robert Sprinkle as a person employed by Google who is knowledgeable about the structure and function of the accused products.

Google reserves the right to supplement this response at an appropriate time, after further discovery and after a claim construction Order in this case.

INTERROGATORY NO. 7:

State the date on which Google first became aware of the '386 patent, the patent application that issued as the '386 Patent, any patent application to which the '386 Patent claims

priority, and/or any Netlist patent application disclosing and/or claiming memory density multiplication, memory rank decoding, and/or memory rank multiplication; describe the circumstances leading to such first awareness, including the identity of the person(s) involved, the identity of all documents which refer or relate to such first awareness, and/or the circumstances leading to such first awareness.

RESPONSE TO INTERROGATORY NO. 7:

Google incorporates each of the foregoing General Objections as if set forth fully in response to this Interrogatory. Google further objects to this Interrogatory to the extent it calls for information protected by the attorney-client privilege, the work product doctrine, or any other applicable exemption from discovery. Google objects to this Interrogatory as calling for the production of information that is neither relevant nor likely to lead to the discovery of admissible evidence to the extent it requests information concerning patents other than the '386 patent in suit. Google will respond concerning the patent in suit only. Google further objects to this Interrogatory as vague and ambiguous as to at least the terms "memory density multiplication," "memory rank decoding," and "memory rank multiplication." Google further objects to this Interrogatory as over broad and unduly burdensome to the extent it would require an investigation into the aforementioned irrelevant patents concerning vague and ambiguous subject matter, which have no bearing on this case.

Subject to and without waiving the foregoing objections, Google responds as follows: Google was first made aware of the '386 patent in suit by an e-mail from Mr. Phileasher Tanner of JEDEC to various JEDEC mailing list recipients, including Mr. Rob Sprinkle and Mr. Andrew Swing of Google, on or about Jan. 10, 2008, forwarding a Netlist patent disclosure letter concerning the patent. This e-mail, and the attached letter, were produced by Google in this matter as GNET034096-97 and GNET269919-20.

INTERROGATORY NO. 8:

Identify all Google personnel who communicated with and/or received information from Netlist concerning memory density multiplication, memory rank multiplication, JEDEC Mode C, DxD technology, memory rank decoding, AMB Quad Rank Support and/or the subject matter

disclosed and/or claimed in the '386 Patent; provide the dates on which such communications occurred and/or such information was received; identify the substance of such communications and/or information, and identify all documents referring to or including such communications and/or information.

RESPONSE TO INTERROGATORY NO. 8:

Google incorporates each of the foregoing General Objections as if set forth fully in response to this Interrogatory. Google further objects to this Interrogatory to the extent it calls for information protected by the attorney-client privilege, the work product doctrine, or any other applicable exemption from discovery. Google further objects to this Interrogatory as requesting information that is neither relevant nor likely to led to the discovery of admissible evidence to the extent it requests information concerning various technologies for which Netlist has articulated no connection to the '386 patent in suit. Google further objects to this Interrogatory as assuming facts not in evidence, to the extent it assumes that such topics were ever discussed between Netlist and anyone at Google. Google further objects to this Interrogatory as vague and ambiguous as to at least the terms "memory density multiplication," "memory rank multiplication," "DxD technology," and "memory rank decoding." Google further objects to this Interrogatory as over broad and unduly burdensome to the extent it would require an investigation of communications "concerning" these vague and ambiguous terms.

Subject to and without waiving the foregoing objections, Google responds as follows: pursuant to Federal Rule 33(d), Google identifies the following documents, already produced in this action, that relate to and describe the interactions of which Google is aware and, to the best of its current knowledge and understanding, relate to the subject matter of this Interrogatory: GNET011948; GNET023456; GNET023542-44; GNET257906-07; GNET258209; and GNET269258. This answer is based only on the knowledge of Google the party, and excludes information that outside counsel may have derived from documents produced by Netlist and designated as "Confidential – Attorneys' Eyes Only" which therefore cannot be shown to Google's employees.

///

INTERROGATORY NO. 9:

For each request for admission that Google did not admit in Netlist's First Set of Requests for Admission of Plaintiff Google, Inc., served September 10, 2009, please explain why Google did not admit the request, and identify all documents that support the basis for Google's response to the request and persons with knowledge of the basis for Google's response to the request.

RESPONSE TO INTERROGATORY NO. 9:

Google incorporates each of the foregoing General Objections as if set forth fully in response to this Interrogatory. Google further objects to this Interrogatory to the extent it calls for information protected by the attorney-client privilege, the work product doctrine, or any other applicable exemption from discovery. Google further objects to this Interrogatory as over broad, unduly burdensome, and duplicative to the extent it requests Google to re-state information that it has previously provided, or is concurrently providing, elsewhere.

Subject to and without waiving the foregoing objections, Google responds as follows: Google's responses and objections to Netlist's First Set of Requests for Admission are fully compliant with the requirements of Federal Rule 36, and as such, those responses and objections adequately disclose the reasons for Google's denials and partial denials. Google incorporates those responses and objections here by reference.

Dated: October 27, 2009

FISH & RICHARDSON P.C.

By: 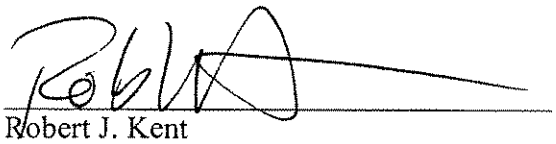
Robert J. Kent

Attorneys for Plaintiff
GOOGLE INC.

50675858.doc

PROOF OF SERVICE

I am employed in the County of San Mateo. My business address is Fish & Richardson P.C., 500 Arguello Street, Suite 500, Redwood City, California 94063. I am over the age of 18 and not a party to the foregoing action.

I am readily familiar with the business practice at my place of business for collection and processing of correspondence for personal delivery, for mailing with United States Postal Service, for facsimile, and for overnight delivery by Federal Express, Express Mail, or other overnight service.

On October 27, 2009, I caused a copy of the following document(s):

PLAINTIFF GOOGLE INC.'S RESPONSES TO NETLIST'S INTERROGATORIES, SET NO. TWO [NOS. 6-9]

to be served on the interested parties in this action by placing a true and correct copy thereof, enclosed in a sealed envelope, and addressed as follows:

Erica J. Pruetz Email: ejpruetz@pruetzlaw.com Adrian M. Pruetz Email: ampruetz@pruetzlaw.com Pruetz Law Group LLP 200 N. Sepulveda Blvd., Suite 1525 El Segundo, CA 90245 Telephone: (310) 765-7650 Facsimile: (310) 765-7641	Attorneys for Defendant and Counterclaimant NETLIST, INC.
Enoch H. Liang Email: ehl@ltlcounsel.com Steven R. Hansen Email: srh@ltlcounsel.com Lee Tran & Liang APLC 601 S. Figuroa Street, Suite 4025 Los Angeles, CA 90017 Telephone: (213) 612-3737 Facsimile: (213) 612-3773	Attorneys for Defendant and Counterclaimant NETLIST, INC.

[X]	**MAIL:**	Such correspondence was deposited, postage fully paid, with the United States Postal Service on the same day in the ordinary course of business.
[ ]	**PERSONAL:**	Such envelope was delivered by hand to the offices of the addressee.
[ ]	**FACSIMILE:**	Such document was faxed to the facsimile transmission machine with the facsimile machine number stated above. Upon completion of the transmission, the transmitting machine issued a transmission report showing the transmission was complete and without error.
[X]	**ELECTRONIC MAIL:**	Such document was transmitted by electronic mail to the addressees' email addresses as stated above.

☐ **FEDERAL EXPRESS:** Such correspondence was deposited on the same day in the ordinary course of business with a facility regularly maintained by Federal Express.

☐ **EXPRESS MAIL:** Such correspondence was deposited on the same day in the ordinary course of business with a facility regularly maintained by the United States Postal Service.

☐ **OVERNIGHT DELIVERY:** Such correspondence was given on the same day in the ordinary course of business to an authorized courier or a driver authorized by that courier to receive documents.

I declare that I am employed in the office of a member of the bar of this Court at whose direction the service was made.

I declare under penalty of perjury that the above is true and correct. Executed on October 27, 2009, at Redwood City, California.

Cheryl Marchesi-Sherwood

50675858.doc

Cheryl Sherwood

From:	Cheryl Sherwood
Sent:	Tuesday, October 27, 2009 5:43 PM
To:	Adrian M. Pruetz (ampruetz@pruetzlaw.com); Enoch H. Liang (ehl@ltlcounsel.com); Erica J. Pruetz (ejpruetz@pruetzlaw.com); Steven R. Hansen (srh@ltlcounsel.com)
Cc:	Robert Kent
Subject:	Google/Netlist: 2009-10-27 Google's Response to Netlist's Requests for Admissions 1-26 and Interrogatories 6-9
Attachments:	2009-10-27 Google Resp to Netlist 1st Set of RFAs Nos 1-26.pdf; 2009-10-27 Google Resp to Netlist 2nd Set of ROGs Nos 6-9.pdf

Counsel,

Please see the attached service copies of Google's discovery responses to Netlist. A paper copy will follow by U.S. mail.

Sincerely,

Cheryl Marchesi-Sherwood
Secretary to Howard G. Pollack and Robert J. Kent
FISH & RICHARDSON P.C.
500 Arguello Street, Suite 500
Redwood City, CA 94063
Direct: (650) 839-5003
Fax: (650) 839-5071
Email: sherwood@fr.com

Exhibit L

Howard G. Pollack (CA Bar No. 162897/pollack@fr.com)
Shelley K. Mack (CA Bar No. 209596/mack@fr.com)
Robert J. Kent (CA Bar No. 250905/rjkent@fr.com)
FISH & RICHARDSON P.C.
500 Arguello Street, Suite 500
Redwood City, CA 94063
Telephone: (650) 839-5070
Facsimile: (650) 839-5071

Attorneys for Plaintiff
GOOGLE INC.

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

(OAKLAND DIVISION)

GOOGLE INC., Plaintiff, v. NETLIST, INC., Defendant. AND RELATED COUNTERCLAIMS.	Case No. C 08-04144 SBA **PLAINTIFF GOOGLE INC.'S RESPONSES TO NETLIST'S REQUEST FOR ADMISSIONS SET NO. ONE [NOS. 1-26]**

Pursuant to Rule 36 of the Federal Rules of Civil Procedure, Plaintiff Google Inc. ("Google") hereby responds to Defendant Netlist, Inc.'s ("Netlist") Request for Admissions, Set No. 1, as follows. These responses are based upon information presently available and are therefore made without prejudice to Google's right to use or rely upon subsequently discovered information. As permitted by the Federal Rules of Civil Procedure, these responses may be changed, modified, or supplemented. In responding to Netlist's Requests for Admission, Google does not waive any objections on the grounds of privilege, competency, relevance, materiality, authenticity, or admissibility of the information contained in these responses. Google also expressly reserves the right to object later to the admissibility of any of this information into evidence on any permissible grounds, including grounds not identified herein.

PRELIMINARY STATEMENT

Netlist's Requests, generally, request Google to admit or deny whether the accused memory modules meet certain limitations, portions of limitations, or incorrect interpretations of limitations of the asserted claims. The Court has not yet construed any limitation of any asserted claim. The Court has not even held a hearing on these issues as of the date of these responses. Accordingly, Netlist's requests are entirely premature. After the Court construes the disputed limitations, Google will consider amending its responses.

Google recognizes its obligation, under Federal Rule 36(a)(4), to make a "reasonable inquiry" before denying or partially denying a request based on lack of knowledge or information. For each Request below which is fully or partially denied on the basis that it calls for a legal conclusion and/or implicates disputed claim terms, Google reviewed the claim construction positions of the parties, including all relevant briefing, to determine that this objection was well-founded and that the Request in question in fact could not be answered without offering a legal conclusion as to claim construction before the Court's hearing and order on that issue. Under the circumstances, this is the most extensive inquiry that could be performed.

The following responses are given without prejudice to Google's right to produce evidence of any facts which it may later discover. Google reserves the right to supplement the following responses and to change any and all of its responses as additional facts are ascertained, analyses are made, legal research is completed, contentions are made, or as a result of the Court's legal determination of issues.

OBJECTIONS TO THE INSTRUCTIONS

Google objects to the Instructions to the extent Netlist seeks to impose obligations on Google that are beyond the scope of or inconsistent with the Federal Rules of Civil Procedure, the Local Rules of the United States District Court for the Northern District of California, and/or the Court's Scheduling Order in this case. Google will respond to the Requests to the extent possible, and subject to its objections set forth herein.

Google further objects to the Instructions to the extent they seek to require to Google to produce information not in its possession, custody, or control. Google further objections to the

Instructions as vague and ambiguous as to at least the term "investigators." Google will respond to the Requests using information available to it after an investigation that is reasonable under the circumstances.

OBJECTIONS TO THE DEFINITIONS

Google objects to the definition of the terms "Google," "you," and "your" to the extent these definitions encompass entities other than plaintiff Google Inc. and to the extent Netlist requests, through these definitions, information not within Google's possession, custody, or control. Google responds on its own behalf only. Google's responses to these requests are made without prejudice to Google's right to produce relevant information obtained from third parties in the future.

Google objects to the definitions of "JEDEC Mode C," "JEDEC Mode A," "Mode C," and "Mode A" as vague and ambiguous. Although Netlist professes to use those terms as defined in JEDEC Standard number JESD82-20A, Google objects to their use in these Requests to the extent that use is incompatible or inconsistent with the way the terms are used within that standard. Google objects to the definitions of "Southbound Link," "Rank Select Bit," "Address Bit," "Row Address Bit," "Column Address Bit," "Chip Select Bit," "Command Bit," "Activate Command," "Write Command," "Read Command," "Precharge Command," and "Refresh Command" as vague and ambiguous. Although Netlist professes to use those terms as defined in JEDEC Standards documents, Google objects to their use in these Requests to the extent that use is incompatible or inconsistent with the way the terms are used within those standards.

GENERAL OBJECTIONS

Google's responses are subject to the following General Objections, which Google incorporates into its responses to each of Netlist's requests, whether or not such General Objection is expressly referenced. The incorporation by reference of any one of these General Objections shall not be construed to exclude the incorporation of any other General Objection. Moreover, Google does not waive its right to amend its objections.

///

///

1. Google objects to the requests insofar as they are vague, ambiguous, indefinite, overbroad, unduly burdensome, duplicative, cumulative, indefinite as to time or scope, unintelligible, or otherwise unclear as to the precise information sought.

2. In particular, Google objects to the term "bit," and variants, as used by Netlist in the Requests. While Netlist ostensibly imports the definition of "bit" and related terms ("Rank Select Bit," "Address Bit," etc.) from JEDEC standards documents, these terms are not expressly defined in those documents and instead are only defined, if at all, by contextual use in relation to other terms. In addition, the relation of these terms to disputed claim terms is ambiguous, and even contradictory, as used in the JEDEC standards and in the Requests. For instance, the term "bit" is nowhere expressly defined in either the Requests or in the JEDEC standards, although the term "bit lane" is defined in document JESD206, where it is said to mean "[a] differential pair of signals in one direction," JESD206 at p. 1, Table 1-1 – which indicates that a bit may be derived from multiple signals. However, as used in the Requests, e.g. where Netlist asks about "Input Command Bits encoding" various commands, it appears that the Requests presume a correspondence between a signal and a *series* of bits. Because Netlist defines these terms only by reference to ambiguous documents, and further because Netlist clearly implies a connection between these terms and various disputed claim terms, any Request using the term "bit" or any variant is vague, ambiguous, and prematurely calls for a legal conclusion before the disputed claim terms have been construed by the Court.

3. Google objects to the requests insofar as they seek information that is neither relevant to a claim or defense of any party, nor reasonably calculated to lead to the discovery of admissible evidence.

4. Google objects to the requests to the extent that they seek documents protected by the attorney-client privilege or by the work-product doctrine, protected by any other applicable privilege or immunity, prepared in connection with settlement discussions, prepared in anticipation of adversarial proceedings such as litigation or for trial, prepared in connection with any applicable joint defense agreement, or not otherwise within the scope of permissive discovery under the Federal Rules of Civil Procedure and applicable Local Rules.

5. Google objects to the requests on the ground and to the extent they call for information that Google is under an obligation to third parties to not disclose.

6. Google objects to the requests on the ground and to the extent they seek to obtain information not in Google's possession, custody, or control.

7. Google objects to the requests as overly burdensome on the ground and to the extent they seek information already in Netlist's possession or information that is a matter of public record or that is otherwise equally available to Netlist.

8. Google objects to the requests to the extent they call for a legal opinion or conclusion. Google neither expresses nor intends to express any legal opinion or conclusion by responding to Netlist's requests.

9. Google objects to the requests to the extent that they fail to specify a relevant time period for which information is requested, and/or to the extent the specified period is irrelevant.

10. Google objects to the Requests to the extent they are premature under any relevant discovery and/or scheduling orders, and due to the fact that there has been no claim construction hearing or order in this case.

11. Google objects to the Requests to the extent that they use terms that are not defined or understood, or are vaguely and/or ambiguously defined, and therefore fail to identify with reasonable particularity the information sought. Google will not speculate as to the meaning to ascribe to such terms.

RESPONSES TO REQUEST FOR ADMISSIONS

REQUEST FOR ADMISSION NO. 1:

Google uses 4-Rank Fully Buffered Dual-In-Line Memory Modules in certain of its servers ("Google's 4-Rank FBDIMMs").

RESPONSE TO REQUEST FOR ADMISSION NO. 1:

Google incorporates by reference each of the General Objections. Google further objects to this request as vague and ambiguous as to the term "4-Rank Fully Buffered Dual-In-Line Memory Modules," which is not defined in the requests either explicitly or via reference to a standard.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: as Google understands the term "4-Rank FBDIMM," Google admits that it uses 4-Rank FBDIMMs. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 2:

The server that Google provided to Netlist for inspection on August 19, 2009 is representative of Google's servers that include Google's 4-Rank FBDIMMs.

RESPONSE TO REQUEST FOR ADMISSION NO. 2:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to the term "representative."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google admits that the server presented for inspection on August 19, 2009 is functionally representative of servers using the allegedly infringing 4-rank FBDIMM memory modules in Google's data centers, in that it allowed Netlist to operate the allegedly infringing 4-rank FBDIMM memory module in a manner functionally representative of the memory module as used in servers in Google's data centers. To the extent that Netlist uses the term "representative" in any other sense, Google is unable to admit or deny the remainder of this Request. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 3:

In certain of Google's servers, Google operates Google's 4-Rank FBDIMMs in JEDEC Mode C.

RESPONSE TO REQUEST FOR ADMISSION NO. 3:

Google incorporates by reference each of the General Objections. In addition, insofar as the term "Mode C" has the same meaning as in the JEDEC Standard JESD82-20A, it invokes the terms "chip select signal" by implication, as those terms are in turn used to define "Mode C" in the standards documents. The definitions of "chip select signal" in the context of the patent are currently subject to debate by the parties, as is the relevance of the JEDEC standards in determining this meaning. Google further objects to this Request as vague and ambiguous as to the term "Mode C."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google admits that certain FBDIMMs used in certain of its servers follow the Mode C serial channel communication protocol set forth in the JEDEC standard for the respective DRAM used on the DIMM. To the extent not admitted, Google lacks sufficient information to admit or deny this Request. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 4:

Google's 4-Rank FBDIMMs include a plurality of DRAM chips coupled to a printed circuit board.

RESPONSE TO REQUEST FOR ADMISSION NO. 4:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to the terms "DRAM chips" and "printed circuit board."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: without acceding to Netlist's definitions of the aforementioned vague, ambiguous, and/or disputed terms, Google admits that certain of its memory modules include DRAM chips coupled to a printed circuit board. To the extent not admitted, Google lacks sufficient information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 5:

Google's 4-Rank FBDIMMs include an Advanced Memory Buffer ("Google's AMB").

RESPONSE TO REQUEST FOR ADMISSION NO. 5:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to the term "Advanced Memory Buffer" (AMB).

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: without acceding to Netlist's definitions of the aforementioned vague, ambiguous, and/or disputed terms, Google admits that the FBDIMMs used by Google include what it understands to be an Advanced Memory Buffer. To the extent not admitted, Google lacks sufficient information

to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 6:

Certain of Google's AMBs include a hardware circuit that receives bits as input ("Input Bits") and which performs at least one predefined function on the Input Bits.

RESPONSE TO REQUEST FOR ADMISSION NO. 6:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "hardware circuit" and "predefined function." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms. Google further objects to this Request as prematurely calling for a legal conclusion before the Court construes the claims, as it includes terms alleged by Netlist to define the disputed term "logic element."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 7:

Certain of Google's AMBs include a hardware circuit that performs a predefined function on Input Bits to generate output bits.

RESPONSE TO REQUEST FOR ADMISSION NO. 7:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "hardware circuit," "predefined function," and "output bits." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms. Google further objects to this Request as prematurely calling for a legal conclusion before the Court construes the claims, as it includes terms alleged by Netlist to define the disputed term "logic element."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: denied.

REQUEST FOR ADMISSION NO. 8:

DRAM chips on Google's 4-Rank FBDIMMs are arranged in ranks.

RESPONSE TO REQUEST FOR ADMISSION NO. 8:

Google incorporates by reference each of the General Objections. Google further objects to this request as vague and ambiguous as to the term "arranged."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Per the parties' stipulated construction of the term "rank," the allocation of DRAM chips into ranks is not a matter of physical arrangement, but rather of electrical connection and logical relationship. Based on that construction, Google admits that its 4-Rank FBDIMMs include DRAM chips organized in ranks. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 9:

DRAM chips on Google's 4-Rank FBDIMMs are arranged in rows.

RESPONSE TO REQUEST FOR ADMISSION NO. 9:

Google incorporates by reference each of the General Objections.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: without acceding to Netlist's definitions any disputed claim terms, Google admits that some of the DRAM chips on certain of its FBDIMMs are physically laid out in rows. To the extent not admitted, Google denies this request.

REQUEST FOR ADMISSION NO. 10:

In certain of Google's servers, at least one Google AMB is electrically coupled to the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 10:

Google incorporates by reference each of the General Objections. Google further objects to this request as vague and ambiguous as to at least the terms "Google AMB," "electrically coupled" and "memory controller."

///

///

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: without acceding to Netlist's definitions of the aforementioned vague, ambiguous, and/or disputed terms, as Google understands it, this Request is admitted.

REQUEST FOR ADMISSION NO. 11:

In certain of Google's servers, at least one Google AMB receives bits ("Google's AMB Input Bits") from the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 11:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "receives" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge or information to admit or deny this Request at this time. Google reserves the right to supplement its response at an appropriate time.

REQUEST FOR ADMISSION NO. 12:

In certain of Google's servers, a Southbound Link is electrically coupled to at least one Google AMB and to the server memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 12:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "electrically coupled" and "memory controller."

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: without acceding to Netlist's definitions of the aforementioned vague, ambiguous, and/or disputed terms, as Google understands it, this Request is admitted.

REQUEST FOR ADMISSION NO. 13:

In certain of Google's servers, at least one Google AMB receives DRAM Address Bits from the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 13:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Address Bits" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to admit or deny this Request at this time. Google reserves the right to supplement its response at an appropriate time.

REQUEST FOR ADMISSION NO. 14:

In certain of Google's servers, at least one Google AMB receives DRAM Row Address Bits from the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 14:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Row Address Bits" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to admit or deny this Request at this time. Google reserves the right to supplement its response at an appropriate time.

REQUEST FOR ADMISSION NO. 15:

In certain of Google's servers, at least one Google AMB receives DRAM Column Address Bits from the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 15:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Column Address Bits" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

///

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to admit or deny this Request at this time. Google reserves the right to supplement its response at an appropriate time.

REQUEST FOR ADMISSION NO. 16:

In certain of Google's servers, at least one Google AMB receives DRAM Bank Address Bits from the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 16:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Bank Address Bits" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to admit or deny this Request at this time. Google reserves the right to supplement its response at an appropriate time.

REQUEST FOR ADMISSION NO. 17:

In certain of Google's servers, at least one Google AMB receives a number of Rank Select Bits ("AMB Input Rank Select Bits") from the server's memory controller.

RESPONSE TO REQUEST FOR ADMISSION NO. 17:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Rank Select Bits" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to admit or deny this Request at this time. Google reserves the right to supplement its response at an appropriate time.

REQUEST FOR ADMISSION NO. 18:

In certain of Google's servers, at least one Google AMB receives a number of AMB Input Rank Select Bits and generates a number of Rank Select Bits ("AMB Output Rank Select Bits")

wherein the number of AMB Output Rank Select Bits is greater than the number of AMB Input Rank Select Bits.

RESPONSE TO REQUEST FOR ADMISSION NO. 18:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the term "Google AMB," "Rank Select Bits." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: denied. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 19:

In certain of Google's servers, at least one Google AMB receives Chip Select Bits that are collectively capable of activating no more than two ranks of DRAM chips (AMB Input Chip Select Bits).

RESPONSE TO REQUEST FOR ADMISSION NO. 19:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Chip Select Bits," "collectively capable of activating," and "capable of activating no more than two ranks." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: as phrased, Google lacks sufficient information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 20:

In certain of Google's servers, at least one Google AMB receives Google's AMB Input Chip Select Bits and generates Chip Select Bits that are collectively capable of activating four ranks of DRAM chips.

///

RESPONSE TO REQUEST FOR ADMISSION NO. 20:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Chip Select Bits" and "collectively capable of activating." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: denied. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 21:

In certain of Google's servers, at least one Google AMB receives DRAM Command Bits from the server's memory controller ("Google's AMB Input Command Bits").

RESPONSE TO REQUEST FOR ADMISSION NO. 21:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google AMB," "Command Bits" and "memory controller." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 22:

Certain of Google's AMB Input Command Bits encode DRAM Activate Commands.

RESPONSE TO REQUEST FOR ADMISSION NO. 22:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google's AMB," "Command Bits," "encode," and "Activate Commands." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

///

///

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 23:

Certain of Google's AMB Input Command Bits encode DRAM Write Commands.

RESPONSE TO REQUEST FOR ADMISSION NO. 23:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google's AMB," "Command Bits," "encode," and "Write Commands." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 24:

Certain of Google's AMB Input Command Bits encode DRAM Precharge Commands.

RESPONSE TO REQUEST FOR ADMISSION NO. 24:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google's AMB," "Command Bits," "encode," and "Precharge Commands." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 25:

Certain of Google's AMB Input Command Bits encode DRAM Refresh Commands.

RESPONSE TO REQUEST FOR ADMISSION NO. 25:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google's AMB," "Command

Bits," "encode," and "Refresh Commands." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

REQUEST FOR ADMISSION NO. 26:

Certain of Google's AMB Input Command Bits encode DRAM Read Commands.

RESPONSE TO REQUEST FOR ADMISSION NO. 26:

Google incorporates by reference each of the General Objections. Google further objects to this Request as vague and ambiguous as to at least the terms "Google's AMB," "Command Bits," "encode," and "Read Commands." Google further specifically objects to this Request on the basis of General Objection No. 2, above, concerning the "bit" terms.

Subject to, without waiving, and based upon the foregoing objections, Google responds as follows: Google lacks sufficient knowledge and information to either admit or deny this Request at this time. Google reserves the right to supplement or amend its response at an appropriate time.

Dated: October 27, 2009

FISH & RICHARDSON P.C.

By: _____
Robert J. Kent

Attorneys for Plaintiff
GOOGLE INC.

50675868.doc

PROOF OF SERVICE

I am employed in the County of San Mateo. My business address is Fish & Richardson P.C., 500 Arguello Street, Suite 500, Redwood City, California 94063. I am over the age of 18 and not a party to the foregoing action.

I am readily familiar with the business practice at my place of business for collection and processing of correspondence for personal delivery, for mailing with United States Postal Service, for facsimile, and for overnight delivery by Federal Express, Express Mail, or other overnight service.

On October 27, 2009, I caused a copy of the following document(s):

PLAINTIFF GOOGLE INC.'S RESPONSES TO NETLIST'S REQUEST FOR ADMISSIONS SET NO. ONE [NOS. 1-26]

to be served on the interested parties in this action by placing a true and correct copy thereof, enclosed in a sealed envelope, and addressed as follows:

Erica J. Pruetz Email: ejpruetz@pruetzlaw.com Adrian M. Pruetz Email: ampruetz@pruetzlaw.com Pruetz Law Group LLP 200 N. Sepulveda Blvd., Suite 1525 El Segundo, CA 90245 Telephone: (310) 765-7650 Facsimile: (310) 765-7641	Attorneys for Defendant and Counterclaimant NETLIST, INC.
Enoch H. Liang Email: ehl@ltlcounsel.com Steven R. Hansen Email: srh@ltlcounsel.com Lee Tran & Liang APLC 601 S. Figuroa Street, Suite 4025 Los Angeles, CA 90017 Telephone: (213) 612-3737 Facsimile: (213) 612-3773	Attorneys for Defendant and Counterclaimant NETLIST, INC.

[X]	**MAIL:**	Such correspondence was deposited, postage fully paid, with the United States Postal Service on the same day in the ordinary course of business.
[ ]	**PERSONAL:**	Such envelope was delivered by hand to the offices of the addressee.
[ ]	**FACSIMILE:**	Such document was faxed to the facsimile transmission machine with the facsimile machine number stated above. Upon completion of the transmission, the transmitting machine issued a transmission report showing the transmission was complete and without error.
[X]	**ELECTRONIC MAIL:**	Such document was transmitted by electronic mail to the addressees' email addresses as stated above.

☐ **FEDERAL EXPRESS:** Such correspondence was deposited on the same day in the ordinary course of business with a facility regularly maintained by Federal Express.

☐ **EXPRESS MAIL:** Such correspondence was deposited on the same day in the ordinary course of business with a facility regularly maintained by the United States Postal Service.

☐ **OVERNIGHT DELIVERY:** Such correspondence was given on the same day in the ordinary course of business to an authorized courier or a driver authorized by that courier to receive documents.

I declare that I am employed in the office of a member of the bar of this Court at whose direction the service was made.

I declare under penalty of perjury that the above is true and correct. Executed on October 27, 2009, at Redwood City, California.

Cheryl Marchesi-Sherwood

50675868.doc

Cheryl Sherwood

From: Cheryl Sherwood

Sent: Tuesday, October 27, 2009 5:43 PM

To: Adrian M. Pruetz (ampruetz@pruetzlaw.com); Enoch H. Liang (ehl@ltlcounsel.com); Erica J. Pruetz (ejpruetz@pruetzlaw.com); Steven R. Hansen (srh@ltlcounsel.com)

Cc: Robert Kent

Subject: Google/Netlist: 2009-10-27 Google's Response to Netlist's Requests for Admissions 1-26 and Interrogatories 6-9

Attachments: 2009-10-27 Google Resp to Netlist 1st Set of RFAs Nos 1-26.pdf; 2009-10-27 Google Resp to Netlist 2nd Set of ROGs Nos 6-9.pdf

Counsel,

Please see the attached service copies of Google's discovery responses to Netlist. A paper copy will follow by U.S. mail.

Sincerely,

Cheryl Marchesi-Sherwood
Secretary to Howard G. Pollack and Robert J. Kent
FISH & RICHARDSON P.C.
500 Arguello Street, Suite 500
Redwood City, CA 94063
Direct: (650) 839-5003
Fax: (650) 839-5071
Email: sherwood@fr.com

Exhibit M

TIMOTHY T. SCOTT (SBN 126971/tscott@kslaw.com)
GEOFFREY M. EZGAR (SBN 184243/ gezgar@kslaw.com)
LEO SPOONER III (SBN 241541/lspooner@kslaw.com)
KING & SPALDING LLP
333 Twin Dolphin Drive, Suite 400
Redwood Shores, CA 94065
Telephone: (650) 590-0700
Facsimile: (650) 590-1900

SCOTT T. WEINGAERTNER (*pro hac vice*/sweingaertner@kslaw.com)
ROBERT F. PERRY (rperry@kslaw.com)
ALLISON ALTERSOHN (*pro hac vice*/aaltersohn@kslaw.com)
SUSAN KIM (*pro hac vice*/skim@kslaw.com)
MARK H. FRANCIS (*pro hac vice*/mfrancis@kslaw.com)
DANIEL MILLER (*pro hac vice*/dmiller@kslaw.com)
King & Spalding LLP
1185 Avenue of the Americas
New York, NY 10036-4003
Telephone: (212) 556-2100
Facsimile: (212) 556-2222

Attorneys for Plaintiff
GOOGLE INC.

UNITED STATES DISTRICT COURT FOR THE

NORTHERN DISTRICT OF CALIFORNIA

OAKLAND DIVISION

GOOGLE INC. Plaintiff, v. NETLIST, INC., Defendant.	Civil Action No. C08 04144 SBA **EXPERT REPORT OF WILLIAM HOFFMAN REGARDING INVALIDITY OF U.S. PATENT 7,289,386**

I. Introduction

1. My name is William Hoffman, and I have written this Expert Report at the request of Google Inc. ("Google") for consideration by the U.S. District Court, Northern District of

115. To the extent that the above references are not deemed to be anticipatory, it is my opinion that the asserted claims of the '386 Patent are invalid as obvious in view of the Dell '074 Reference in combination with the Wong '868 Reference. My detailed claim analysis is provided in Exhibit Y of this Expert Report.

116. A person of ordinary skill in the art would have combined these references because they are both directed to addressing schemes in memory modules.

F. Wong '868 Reference in View of Wong '281 Reference

117. To the extent that the above references are not deemed to be anticipatory, it is my opinion that the asserted claims of the '386 Patent are invalid as obvious in view of the Wong '868 Reference in combination with the Wong '281 Reference. My detailed claim analysis is provided in Exhibit Z of this Expert Report.

118. A person of ordinary skill in the art would have combined these references because they are both directed to addressing schemes in memory modules, are assigned to the same entity, and share common inventors.

G. Amidi Reference in View of Dell '074 Reference

119. To the extent that the above references are not deemed to be anticipatory, it is my opinion that the asserted claims of the '386 Patent are invalid as obvious in view of the Amidi Reference in combination with the Dell '074 Reference. My detailed claim analysis is provided in Exhibit AA of this Expert Report.

120. A person of ordinary skill in the art would have combined these references because they are both directed to addressing schemes in memory modules.

X. Invalidity Under Section 112

A. Violation of the Written Description Requirement and Lack of Enablement

121. Claim 1 of the '386 Patent recites in relevant part "the set of input control signals corresponding to a second number of memory devices smaller than the first number of memory devices." This claim limitation requires that the control signals be for a number of logical memory devices less than the number of actual memory devices. For example, the claim language would apply to a memory module have four memory devices that receives input control signals for only two devices.

122. The '386 Patent only describes embodiments where the input signals are for fewer than the actual number of ranks. For example, Figure 1A of the '386 Patent, reproduced below, shows that two received chip select signals, CS_0 and CS_1, are mapped to four chip select signals, CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B}.

FIG. 1A

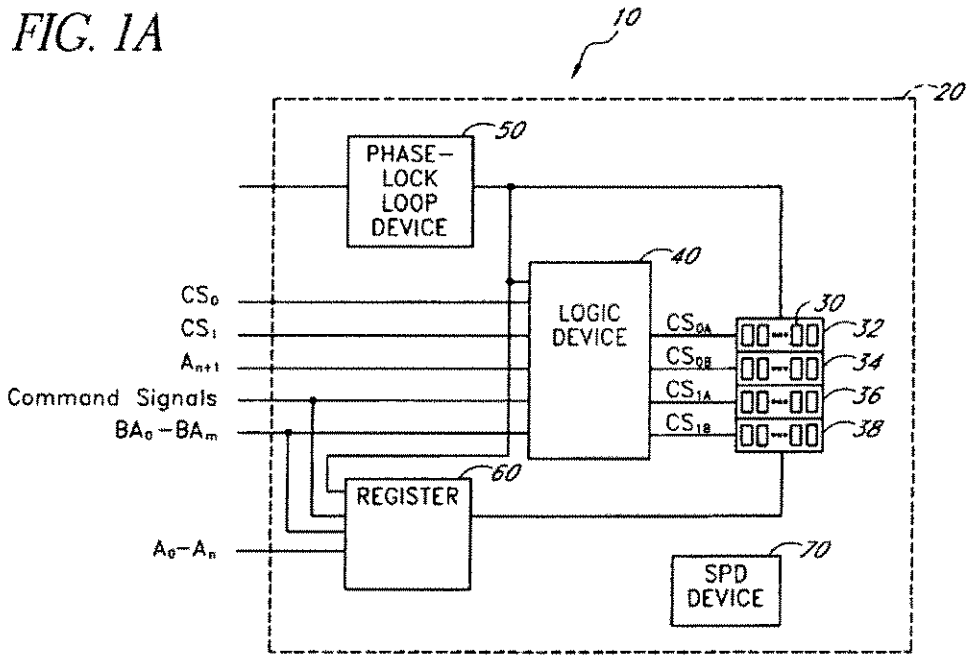

123. In describing the preferred embodiment, the '386 Patent explains:

> For example, in the exemplary embodiment as schematically illustrated by FIG. 1A, the first number of ranks is four while the second number of ranks is two. In the exemplary embodiment of FIG. 1B, the first number of ranks is two while the second number of ranks is one. Thus, in certain embodiments, even though the memory module 10 actually has the first number of ranks of memory

devices 30, the memory module 10 simulates a virtual memory module by operating as having the second number of ranks of memory devices 30. In certain embodiments, the memory module 10 simulates a virtual memory module when the number of memory devices 30 of the memory module 10 is larger than the number of memory devices 30 per memory module for which the computer system is configure to utilize.
('386 Patent, Ex. A, 7:14-29.)

In the exemplary embodiment schematically illustrated by FIG. 1A, the memory module 10 has four ranks of memory devices 30 and the computer system is configured for two ranks of memory devices per memory module.
('386 Patent, Ex. A, 7:45-48.)

124. The specification of the '386 Patent describes how to implement a memory module that has more memory sub-arrays than the computer system is expecting. In fact, the specification shows that the Serial Presence Detect device ("SPD") only tells the computer system that there are fewer devices than are actually present.

In certain embodiments, the SPD device 70 comprises data which characterize the memory module **10** as having fewer ranks of memory devices than the memory module **10** actually has, with each of these ranks having more memory density. For example, for a memory module **10** compatible with certain embodiments described herein having two ranks of memory devices 30, the SPD device 70 comprises data which characterizes the memory module **10** as having one rank of memory devices with twice the memory density per rank. Similarly, for a memory module **10** compatible with certain embodiments described herein having four ranks of memory devices 30, the SPD device 70 comprises data which characterizes the memory module **10** as having two ranks of memory devices with twice the memory density per rank. In addition, in certain embodiments, the SPD device 70 comprises data which characterize the memory module 10 as having fewer memory devices than the memory module **10** actually has, with each of these memory devices having more memory density per memory device. For example, for a memory module **10** compatible with certain embodiments described herein, the SPD device 70 comprises data which characterizes the memory module **10** as having one-half the number of memory devices that the memory module **10** actually has, with each of these memory devices having twice the memory density per memory device.

('386 Patent, Ex. A, 10:31-55.)

125. I understand that Netlist has asserted that this claim element is met by Google products that practice Mode C of JEDEC Standard JESD 82-20A. (*See* Netlist Amended Preliminary

Infringement Contentions, attached as Exhibit BB.) More specifically, I understand that Netlist has asserted that this claim element is met by four-rank memory modules that receive from the memory controller a first rank select in the field RS0 and a second rank select (RS1) in the field BA0. As a result, I understand that Netlist has asserted that a four-rank device that receives as input two rank select signals would infringe.

126. Two bits, such as those contained in two rank select signals, may be used to identify four devices. As a result, the accused four-rank devices receive input control signals corresponding to the number of actual devices, four, not some smaller number of devices, such as two.

127. To the extent that Netlist's infringement contentions are directed to memory modules that receive control signals corresponding to the number of actual memory devices, not some lesser number of logical memory devices, there is no support in the specification to enable this implementation. There is simply a logical contradiction in Netlist's application in view of the specification. The specification does not enable a system in which there is an input addressing a number of ranks that is less than the actual number of ranks while also at the same time being the same as the actual number of ranks.

128. In addition, to the extent that Netlist's infringement contentions are directed to memory modules that receive control signals corresponding to the number of actual memory devices, not some lesser number of logical memory devices, there is no supporting written description. As seen in the examples above, the specification does not describe any addressing scheme where the computer system addresses the actual number of ranks.

XI. Omission of Reference Material to Patentability

A. Dell '827 Reference

129. The '386 Patent did not issue until October 30, 2007. I understand that a patentee has an obligation to disclose all material references of which it is aware.

130. Netlist and its attorney, Bruce Itchkawitz, were aware of the Dell '827 Reference as early as September 4, 2007, almost a month an a half before the issuance of the '386 Patent. (See Information Disclosure Statement, attached as Exhibit CC.)

131. The Dell '827 Reference is material to claims 1 and 11 of the '386 Patent, as can be seen in my analysis above and in the attached claim chart, Exhibit L.

XII. Reservation of Right to Supplement

132. I reserve the right to modify or supplement my opinions, if necessary, based on further review and analysis of evidence in this case, including review and analysis of any information that may be provided to me subsequent to the filing of this report. I also reserve the right to modify or supplement my opinions based on any changes to the claim constructions I have used in my analysis for this matter.

133. In addition, I understand that Netlist has filed a motion to amend its infringement contentions to include additional claims. I further reserve the right to supplement my opinions to include review and analysis of invalidity with regard to any claims that the Court may allow Netlist to add.

XII. Conclusion

134. Based on my review and analysis to date, I hold the opinion that the '386 Patent is invalid in view of at least the aforementioned prior art.

DATED: April 15, 2010

By: William Hoffman
William Hoffman